

# **Learning**

## **From An Optimization Viewpoint**

by

**Karthik Sridharan**

Submitted to :

Toyota Technological Institute at Chicago  
6045 S. Kenwood Ave, Chicago, IL, 60637

For the degree of Doctor of Philosophy in Computer Science

Thesis Committee :

**Nathan Srebro (Thesis Supervisor),**

**David McAllester,**

**Arkadi Nemirovski,**

**Alexander Razborov**



In memory of my dear father, Raghavan Sridharan . . .

# Acknowledgements

Starting to write this acknowledgement I realized that I owed thanks to many people who have directly or indirectly influenced my research and have helped me complete this dissertation.

First and foremost, I would like to thank my advisor Nati Srebro for his encouragement, guidance and patience throughout my PhD. One too many times I have walked into Nati's office with vague questions and still unformed ideas and every time I walked out with crisp and formalized ideas and clear thoughts on how to approach problems. Nati was the one who inspired me to look at learning alternatively as an optimization problem, which culminated into the central idea of this thesis. Besides research skills, I also learned (or at least started the process of learning) other skills including scientific writing, presenting my ideas succinctly and choosing the right problems to work on. Perhaps the most important thing I learned from Nati is that asking the right questions and formalizing these questions precisely is as important as, if not more than, finding answers.

I would like to thank Sham Kakade with whom I worked during the first three years of my PhD on many exciting projects and learned many things from. I am greatly indebted to Shai Shalev-Shwartz; apart from being a wonderful collaborator, Shai has guided me in many ways and greatly influence my research and thinking. I express my heartfelt gratitude to Alexander Rakhlin, Ohad Shamir and Ambuj Tewari. Besides being my close collaborators and friends, they gave me timely advice and help, and taught me many things regarding research and other aspects. I would like to also thank my other research collaborators who influenced and guided my research: Shai Ben-David, Kamalika Chaudhuri, Andrew Cotter, Ofer Dekel, Dean Foster, Rina Foygel, Claudio Gentile, Karen Livescuc and David Loker.

I would like to thank my committee members David McAllester, Arkadi Nemirovski and Alexander Razborov for all their help in guiding and shaping this thesis work. I enjoyed the many insightful discussions with David; I will especially remember the ones we had fairly early in the morning at TTI-Chicago when it used to be the start of Davids day and most likely the end of mine. I am grateful to Arkadi Nemirovski for his valuable suggestions that helped resolve issues I was struggling with, and also make my thesis better . I am grateful to Alexander Razborov for his valuable guidance and advice. Overall I feel privileged and honored to have a thesis committee that guided and moulded my research and the way I think about problems.

I greatly acknowledge the most wonderful staff at TTI-Chicago who made my stay at Chicago so memorable and pleasant. I would like to specially thank Carole Flemming, Gary Hamburg, Liv Leader and Christina Novak for helping me with numerous things throughout my PhD. I also thank Avleen Bijral, Wonseok Chae, Heejin Choi, Andrew Cotter, Patrick Donovan, Taehwan Kim, Jianzhu Ma, Arild Brandrud Nss, Jian Peng, Ankan Saha, Allie Shapiro, Hao Tang, Hoang Trinh, Zhiyong Wang, Xing Xu, Payman Yadollahpour, Jian Yao and Feng Zhao for their friendship and making my time at TTI-Chicago fun. Thanks also to the other students and faculty at TTI-Chicago with whom I have had insightful discussions and pleasant conversations.

Last, but not the least, I would like to express my deepest gratitude to my family and relatives, without whose support I would not have been able to complete my PhD. To Aditi, Pranav, Lavan, Sreeram and Amma, thanks for all the love and encouragement you have given me. I thank my parents for their unconditional support and confidence they have shown in me. I dedicate this dissertation to my late father, Raghavan Sridharan, whose greatest dream was to see me complete my PhD.

# Abstract

Optimization has always played a central role in machine learning and advances in the field of optimization and mathematical programming have greatly influenced machine learning models. However the connection between optimization and learning is much deeper : one can phrase statistical and online learning problems directly as corresponding optimization problems. In this dissertation I take this viewpoint and analyze learning problems in both the statistical and online learning frameworks from an optimization perspective. In doing so we develop a deeper understanding of the connections between statistical and online learning and between learning and optimization.

The dissertation can roughly be divided into two parts. In the first part we consider the question of learnability and possible learning rates for general statistical and online learning problems without regard to tractability issues. In the second part we restrict ourselves to convex learning problems and address the issue of tractability for both online and statistical learning problems by considering the oracle complexity of these learning problems.

- I. We first consider the question of learnability and possible learning rates for statistical learning problems under the general learning setting. The notion of learnability was first introduced by Valiant (1984) for the problem of binary classification in the realizable case. Vapnik (1995) introduced the general learning setting as a unifying framework for the general problem of statistical learning from empirical data. In this framework the learner is provided with a sample of instances drawn i.i.d. from some distribution unknown to the learner. The goal of the learner is to pick a hypothesis with low expected loss based on the sample received. The question of learnability is well studied and fully characterized for binary classification using the Vapnik Chervonenkis (VC) theory and for real valued supervised learning problems using the theory of uniform convergence with tools like Rademacher complexity, covering numbers and fat-shattering dimension etc. However we show that for the general learning setting the traditional approach of using uniform convergence theory to characterize learnability fails. Specifically we phrase the learning problem as a stochastic optimization problem and construct an example of a convex problem where Stochastic Approximation (SA) approach provides successful learning guarantee but Empirical Risk Minimization (ERM) (or equivalently Sample Average Approximation (SAA) approach) fails to give any meaningful learning guarantee. This example establishes that for general learning problems the concept of uniform convergence fails to capture learnability and ERM/SAA ap-

proaches can fail to provide successful learning algorithms. To fill this void in the theory of statistical learnability in the general setting we instead turned to the concept of stability of learning algorithms to fully characterize learnability in the general setting. We specifically show that a problem is learnable if and only if there exists a stable approximate minimizer of average loss over the sample. Using this notion of stability, we also provide a universal learning procedure that guarantees success whenever the problem is learnable.

Next we consider the problem of online learning in the general setting. Online learning is a continual and sequential learning process where instances provided to the learner round by round and can be chosen adversarially (as opposed to stochastically). The goal of the learner for an online learning problem is to minimize regret with respect to the single best hypothesis that can be chosen in hindsight. Most of the work on online learning problems so far have been algorithmic and problem specific. The usual approach has been to build algorithm for the specific problem at hand and prove regret guarantees for this algorithm which in turn implies learnability with associated learning rates. Unlike the statistical learning framework there is a dearth of generic tools that can be used to establish learnability and rates for online learning problems in general. Only recently, Ben-David, Pal and Shalev-Shwartz (2009) showed that the Littlestone dimension (introduced by Littlestone (1988)) is an online analog to the VC dimension and fully characterizes learnability for online binary classification problems. However the question of characterizing online learnability for even the real valued supervised learning problems in the online framework was open. In this dissertation, analyzing the so called value of the online learning game, we provide online analogs to classical tools from statistical learning theory like Rademacher complexity, covering numbers, fat-shattering dimension etc. While these tools can be used to provide upper bounds for more general online learning problems, the results mirror uniform convergence theory for online learning problems. Hence analogous to the statistical learning case, we used these tools to fully characterize learnability and rates for real valued online supervised learning problems. We also provide a generic algorithm for the real valued online supervised learning problem. However unlike the statistical learning case, we don't yet have a full characterization of learnability for general online learning problems and leave this as open problem for future work.

- II. In the first part of the dissertation we focused on the question of learnability and rates for statistical and online learning problems without paying attention to tractability or efficiency of the learning algorithms that could be used. Even the generic algorithms provided in the first part are in general not at all tractable. In the second part of the dissertation we address the issue of building tractable or efficient learning algorithms by first focusing our attention specifically on convex learning problems. In the second part, for general classes of convex learning problems, we provide appropriate mirror descent updates that are guaranteed to be successful for online and statistical learning of these convex problems. Further, using and extending results from the geometry of Banach spaces, we show that the mirror descent method (with appropriate prox-function and step-size) is near optimal for online convex learning problems and for most reasonable cases, is also near optimal for statistical convex

learning problems. Further noting that when used for statistical convex learning, the mirror descent algorithm is a first-order (gradient based),  $O(1)$  memory, single pass algorithm, we conclude that mirror descent method is also near optimal in terms of number of gradient accesses and for many commonly encountered problems, optimal also in terms of computational time. We next consider the problem of (offline) convex optimization and to capture the notion of efficiency of optimization procedures for these problems, we use the notion of oracle complexity of the problem introduced by Nemirovski and Yudin (1978). We show that for a general class of convex optimization problems, oracle complexity of the problem can be lower bounded by the so called fat-shattering dimension of the associated linear class. Thus we establish a strong connection between offline convex optimization problems and statistical convex learning problems. We further show that for a large class of infinite dimensional (or high dimensional) optimization problems, mirror descent is in fact near optimal in terms of oracle efficiency even for these offline convex optimization problems.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Learning and Optimization . . . . .	2
1.2	Overview of the Thesis . . . . .	3
1.2.1	Part I : Statistical and Online Learning : Learnability and Rates . . .	3
1.2.2	Part II : Convex Problems : Oracle Efficient Learning/Optimization	4
1.3	Main Contributions . . . . .	5
1.4	Bibliographic Notes . . . . .	6
<b>I</b>	<b>Statistical and Online Learning : Learnability and Rates</b>	<b>7</b>
<b>2</b>	<b>Preliminary Setup and Notations</b>	<b>8</b>
2.1	General Learning Problem Setup . . . . .	8
2.2	More Definitions and Notations . . . . .	9
<b>3</b>	<b>Statistical Learning/Optimization</b>	<b>12</b>
3.1	The Statistical Learning Problem and Learnability . . . . .	13
3.2	Background . . . . .	16
3.2.1	Learnability and Uniform Convergence . . . . .	16
3.2.2	Various Complexity Measures and Uniform Convergence . . . . .	18
3.2.3	Learnability and Stability . . . . .	20



3.3	Failure of Uniform Convergence and ERM/SAA Approaches . . . . .	21
3.3.1	Learning without Uniform Convergence : Stochastic Convex Optimization . . . . .	21
3.3.2	Learnability via Stability : Role of Regularization . . . . .	24
3.3.3	Contradiction to Vapnik? . . . . .	26
3.4	Stability of Learning Rules . . . . .	27
3.4.1	Comparison with Existing Literature and Other Notions of Stability	28
3.5	Characterizing Learnability : Main Results . . . . .	32
3.6	Randomization, Convexification, and a Generic Learning Rule . . . . .	34
3.6.1	Stronger Results with Randomized Learning Rules . . . . .	34
3.6.2	A Generic Learning Rule . . . . .	36
3.7	Detailed Results and Proofs . . . . .	37
3.7.1	Detailed Proof of Main Result (Section 3.5) . . . . .	37
3.7.2	Other Proofs . . . . .	45
3.8	Discussion . . . . .	50
<b>4</b>	<b>Online Learning/Optimization</b>	<b>52</b>
4.1	The Online Learning Problem . . . . .	53
4.2	Online Learnability and the Value of the Game . . . . .	54
4.3	Sequential Rademacher Complexity . . . . .	56
4.3.1	Structural Results . . . . .	58
4.4	Sequential Covering Number and Combinatorial Parameters . . . . .	60
4.4.1	A Combinatorial Upper Bound . . . . .	62
4.4.2	Finite Class Lemma and the Chaining Method . . . . .	63
4.5	Martingale Uniform Convergence . . . . .	64
4.6	Characterizing Learnability of Supervised Learning Problem . . . . .	66
4.6.1	Generic Algorithm for Supervised Learning Problem . . . . .	68

4.7	Examples . . . . .	70
4.7.1	Example: Margin Based Regret . . . . .	70
4.7.2	Example : Neural Networks . . . . .	70
4.7.3	Example: Decision Trees . . . . .	71
4.7.4	Example: Online Transductive Learning . . . . .	72
4.7.5	Example: Isotron . . . . .	73
4.8	Detailed Proofs and More Results . . . . .	74
4.8.1	Proofs . . . . .	74
4.8.2	Exponentially Weighted Average (EWA) Algorithm on Countable Experts . . . . .	100
4.9	Discussion . . . . .	103
<b>II Convex Problems : Oracle Efficient Learning/Optimization</b>		<b>104</b>
<b>5</b>	<b>Convex Learning and Optimization Problem Setup</b>	<b>105</b>
5.1	Convex Problems . . . . .	105
5.2	Various Convex Learning/Optimization Problems . . . . .	106
5.3	Discussion . . . . .	108
<b>6</b>	<b>Mirror Descent Methods</b>	<b>109</b>
6.1	The Mirror Descent Update . . . . .	109
6.2	Online Mirror Descent . . . . .	110
6.3	Stochastic Mirror Descent . . . . .	113
6.4	Mirror Descent for Offline Optimization . . . . .	115
6.5	Detailed Proofs . . . . .	116
6.6	Discussion . . . . .	122
<b>7</b>	<b>Optimality of Mirror Descent for Online Convex Learning Problem</b>	<b>124</b>

7.1	Value of the Linear Game . . . . .	125
7.2	Value and Martingale Type . . . . .	126
7.3	Martingale Type and Uniform Convexity . . . . .	128
7.4	Main Result : Optimality of Online Mirror Descent . . . . .	129
7.4.1	Smooth Loss Case . . . . .	130
7.4.2	Uniformly Convex Loss Case . . . . .	132
7.5	Examples . . . . .	133
7.5.1	Example : $\ell_p$ non-dual pairs . . . . .	133
7.5.2	Example : Non-dual Schatten norm pairs in finite dimensions . . .	134
7.5.3	Example : Non-dual group norm pairs in finite dimensions . . . . .	134
7.5.4	Example : Max Norm . . . . .	134
7.5.5	Example : Interpolation Norms . . . . .	135
7.6	Detailed Proofs . . . . .	136
7.7	Discussion . . . . .	150
<b>8</b>	<b>Optimality of Mirror Descent for Statistical Convex Learning Problems</b>	<b>151</b>
8.1	Lower Bounds for Statistical Learning Rates . . . . .	152
8.1.1	Lower Bounds for Smooth Losses . . . . .	154
8.2	Optimal Rates and Rademacher Type . . . . .	154
8.3	Main Result : Optimality of Stochastic Mirror Descent . . . . .	155
8.3.1	Banach Lattices . . . . .	157
8.3.2	Decoupling Inequalities . . . . .	159
8.3.3	Optimality of Mirror Descent in Terms of Efficiency . . . . .	161
8.4	Examples . . . . .	162
8.4.1	Example : $\ell_p$ non-dual pairs . . . . .	162
8.4.2	Example : Non-dual Schatten norm pairs in finite dimensions . . .	162
8.4.3	Example : Non-dual group norm pairs in finite dimensions . . . . .	163

8.4.4	Computational Efficiency Issues . . . . .	163
8.5	Detailed Proofs . . . . .	163
8.6	Discussion . . . . .	174
<b>9</b>	<b>Optimality of Mirror Descent for Offline Convex Optimization</b>	<b>175</b>
9.1	Oracle-based Offline Convex Optimization . . . . .	176
9.2	Lower Bounding Oracle Complexity: Connections to Statistical Convex Learning . . . . .	178
9.3	Main Result : Optimality of Mirror Descent for Offline Convex Optimization	180
9.4	Statistical Learning With Distributed Oracles . . . . .	181
9.5	Detailed Proofs . . . . .	183
9.6	Discussion . . . . .	186
<b>10</b>	<b>Conclusion and Future Work</b>	<b>187</b>
10.1	Open Problems . . . . .	187
10.1.1	Online Optimization and Stability . . . . .	187
10.1.2	Upper Bounding Oracle Complexity in Terms of Fat-Shattering Dimension . . . . .	187
10.2	Further Directions . . . . .	188
10.3	Summary . . . . .	188
	<b>Bibliography</b>	<b>189</b>
<b>A</b>	<b>Relating Various Complexity Measures : Statistical Learning</b>	<b>197</b>
A.1	The Refined Dudley Integral: Bounding Rademacher Complexity with $L_2$ Covering Numbers . . . . .	197
A.2	Bounding $L_\infty$ covering number by Fat-shattering Dimension . . . . .	199
A.3	Relating Fat-shattering Dimension and Rademacher complexity . . . . .	201

# List of Figures

3.1	An example of a “trivial” learning situation. Each line represents some $\mathbf{h} \in \mathcal{H}$ , and shows the value of $\ell(\mathbf{h}, z)$ for all $z \in \mathcal{Z}$ . The hypothesis $\tilde{\mathbf{h}}$ dominates any other hypothesis (e.g., $\ell(\tilde{\mathbf{h}}; z) < \ell(\mathbf{h}; z)$ uniformly for all $z$ ), and thus the problem is learnable without uniform convergence or any other property of $\mathcal{H}$ . . . . .	17
7.1	Relationship between the various constants . . . . .	129
8.1	Relationship between the various constants in statistical setting . . . . .	155
8.2	Relationship between the various constants in both online and statistical settings. Red inequality is not always true but true for most commonly encountered problems. . . . .	161

# List of Tables

2.1 Summary of Notations . . . . . 11

# Chapter 1

## Introduction

There are two main paradigms under which machine learning problems have been commonly studied; these are the statistical learning framework and the online learning framework. In the statistical learning framework, the data source is said to be iid, that is data (or instances) are assumed to be drawn independently from some fixed but unknown (to the learner) source distribution. The goal of the learner in this setting is to pick a hypothesis that minimizes expected error on future examples drawn from the source distribution. The statistical learning framework is well suited for applications like object recognition, natural language processing and other applications where the data generation process is unchanging with time. For instance when we consider a problem of image classification between cat and dog images, it is not unreasonable to assume the data source does not vary over time and each time we sample repeatedly from the same distribution over input images. Roughly speaking cats and dogs would look the same even 10 years from now and so we can think of the distribution from which we draw the image label pair to be fixed.

The online learning framework on the other hand is a learning scenario where the learner is faced with a data source that changes over time or even reacts to the learners choices. The online learning framework is often studied as a multiple round game where the data source could even be adversarial. The instances produced by the adversary at any round could depend on the past hypotheses chosen by the learner. The online learning framework is better suited for more interactive learning tasks like spam mail detection and stock market prediction where decisions of the learner could affect future instances the learner receives. For instance in a spam mail filter application, while the company providing the spam filter service works on filtering out spam mails, the spammers themselves try to outwit the system by adapting to how the spam detection software works. In a sense there is a continuous game played between the company providing spam filtering service and the spammers where each one tries to outwit the other. Such a scenario is naturally captured by the online learning framework.

In this dissertation we consider learning in both online and statistical learning frameworks. We view learning problems in both the frameworks from an optimization viewpoint and establish connections between learning and optimization. Specifically, in the first part of the dissertation we consider very general learning problems

in both the statistical and online learning frameworks and focus on the question of when a given learning problem is learnable and at what rates. Of course at that generality we restrict our focus to only questions of learnability and rates for the learning problems and do not pay attention to tractability/efficiency of possible learning algorithms that can be used for these problems.

In the second half of the thesis, we restrict our focus to convex learning problems and beyond just characterizing optimal rates for these learning problems, we also provide simple first order methods based on mirror descent and show that they are near optimal both in terms of rate of convergence and in terms of their efficiency. Specifically we establish that for a large family of general convex learning problems, mirror descent algorithm is universal and near optimal for the online learning framework. We also show that for most reasonable cases, mirror descent is near optimal for these convex learning problems even in the statistical learning framework. Further, the fact that mirror descent is an order  $O(1)$  memory, simple first order method makes it near optimal even in terms of efficiency for many of these learning problems. We also establish that for several high dimensional problems, mirror descent could be near optimal (in terms of oracle efficiency) for (offline) convex optimization problems. In the second half of the thesis we establish some interesting connections between convex learning problems and convex optimization. We establish that for most reasonable cases online convex learning and statistical convex learning are equally hard (or easy). The main contention of the second part of the thesis is that for a large family of convex learning problems, whether it is online or statistical learning framework, the simple one pass, first order method of mirror descent is optimal both in terms of rates and efficiency.

While the work in this dissertation is of a theoretical in nature, we believe that the results and viewpoints provided will drive the choice of learning algorithms considered for various problems and models chosen for learning. Further the second part of the thesis has more direct practical implications and advocates the use of simple first order online learning methods for convex learning problems even when our end goal is statistical learning.

## 1.1 Learning and Optimization

Optimization has always played a central role in shaping machine learning algorithms and models. The typical approach taken to tackle learning problems has been to pick a suitable set of models or hypotheses or predictors, pick appropriate empirical cost over the training sample and finally solve the optimization problem of picking the hypothesis from the set of hypotheses that minimize the empirical cost over training sample. The learning algorithm itself in this case corresponds to the optimization algorithm that solves the minimization problem of picking that hypothesis from the set of hypotheses that minimizes empirical cost over training sample. Theoretical guarantees on learning rates are typically provided by decomposing error term into two terms. The first term, the estimation error that accounts for error we would incur if we picked the hypothesis that minimizes empirical cost over training sample chosen. The second term is the optimization error term that accounts for the sub-optimality of the optimization algorithm in minimizing the empirical cost chosen.



Overall, we see that optimization plays a critical role in the design and analysis of algorithms for learning problems. However the connection between optimization and learning is much deeper. Learning problems under both statistical and online learning frameworks can be directly framed as optimization problems and analysis of the sub-optimality of algorithms the . In this dissertation I will take this viewpoint that learning problems can be viewed as optimization problems and using this point of view establish connections between learning in statistical and online frameworks and optimization. Specifically in the second half of the dissertation, focusing on convex learning problems, we show that mirror descent method originally introduced for convex optimization problems by Nemirovski and Yudin [1], is always near optimal for online learning problems, near optimal for most reasonable statistical learning problems and even near optimal for several high dimensional offline convex optimization problems. We use concepts from the geometry of Banach spaces to establish these optimality results for mirror descent. We also establish some interesting connections between convex learning and convex optimization.

## 1.2 Overview of the Thesis

The dissertation can be split into two parts; the first part focuses on the question of learnability and learning rates for fairly general class of learning problems in both the statistical and online learning frameworks. In the first part, we do not concern ourselves with questions of efficiency or tractability of the learning methods/algorithms for solving these problems. In the second part of this dissertation, we restrict ourselves to convex learning problems and show that for most reasonable problems, the first order method of mirror descent is near optimal for both statistical and online convex learning problems. We also establish connections between convex learning and (offline) convex optimization and establish that for several high dimensional (offline) convex optimization problems, mirror descent is optimal in terms of efficiency for (offline) convex optimization problem. The notion of efficiency we use in this case is the notion of oracle complexity introduced by Nemirovski and Yudin (1978).

### 1.2.1 Part I : Statistical and Online Learning : Learnability and Rates

The first part of this dissertation contains two chapters. The first one, Chapter 3, is dedicated to the question of learnability in the statistical learning framework. In this chapter, through an example that is an instance of a stochastic convex optimization problem, we show that the concept of uniform convergence (and hence tools like Rademacher complexity, fat-shattering dimension) that are commonly used to analyze learning rates in the supervised learning settings fail in the general. Thus we establish that the folklore of uniform convergence is necessary and sufficient for learnability is not true in general. This of course opens up the question : “How can one characterize learnability of general learning problems, in the statistical framework?”. We provide an answer to this question in this chapter by turning to the concept of stability of learning algorithms. We show that learnability in the general setting can be characterized by existence of learning algorithm that is stable and an approximate empirical minimizer. We further proceed to show that if we are allowed to consider randomized learning rules, then we can provide a “Universal Learning Algorithm” which

has non-trivial learning rate whenever the problem is learnable.

Chapter 4, is the second chapter in the first part of the dissertation. It deals with the question of learnability and rates for online learning problems. While the question of learnability is well studied in the statistical framework, the question of learnability in the online setting has relatively less explored. Most of the work on online learning problems so far have been algorithmic and problem specific. The usual approach has been to build algorithm for the specific problem at hand and prove regret guarantees for this algorithm which in turn implies learnability with associated learning rates. Unlike the statistical learning framework there is a dearth of generic tools that can be used to establish learnability and rates for online learning problems in general. Only recently, Ben-David, Pal and Shalev-Shwartz (2009) showed that the Littlestone dimension (introduced by Littlestone (1988)) is an online analog to the VC dimension and fully characterizes learnability for on-line binary classification problems. In general there have been no generic tools like Rademacher complexity, covering numbers and fat-shattering dimension that are present for analyzing statistical learning problems. In this chapter we explore the the question of learnability and optimal rates for online learning by first formalizing them as value of the online learning game. We then build complexity measures analogous to those in the statistical framework like Rademacher complexity, covering numbers and fat-shattering dimension and show that these tools can be used to bound learning rates for online learning problems. These tools can be seen as tools for studying uniform convergence for general stochastic processes (non iid). We go ahead and show that these tools can even characterize learnability and learning rates of real valued supervised learning problems in the online learning framework. Based on these complexity measures, we also provide a generic algorithm for the supervised learning problem in the online learning framework that has diminishing regret whenever the problem is online learnable.

## **1.2.2 Part II : Convex Problems : Oracle Efficient Learning/Optimization**

In the first part of this thesis while we even provide generic learning algorithms for fairly general class of problems in both statistical and online learning frameworks, we never concerned ourselves with any form of tractability of these learning algorithms. However, in the second part of this dissertation, focusing on convex learning problems, we aim at providing optimal and efficient learning algorithms for both the statistical and online learning frameworks. In doing so we also explore connections between learning and convex optimization. There are five chapters in this part. The first chapter of this part, Chapter 5, introduces the basic set up of the convex learning and optimization problems we consider in this part, describes the online and statistical learning protocols/frameworks for these convex problems and introduces the oracle model for offline convex optimization. The next chapter, Chapter 6 introduces the mirror descent methods (see [1]) for the statistical and online convex learning problems described in the previous chapter and provides upper bounds on learning rate for them. Upper bounds on rate of optimization of the mirror descent method for offline convex optimization problems are also provided. Chapter ?? deals with online convex learning problems and in the chapter we show that mirror descent algorithm is universal near optimal for online convex learning problems. In Chapter 8 statistical convex learning problems are considered and lower bounds on learning rates for these problems are provided. Using these lower bounds we further show that for most

reasonable cases, the mirror descent method for statistical learning (stochastic mirror descent method) is near optimal even for statistical learning. In the final chapter of the second part of this dissertation, Chapter 9, (offline) convex optimization problems are considered. As mentioned earlier, we use oracle complexity of the learning problem (ie. minimum number of calls to any local oracle needed by any method to achieve desired accuracy) as a measure of efficiency of the optimization procedure. We show that the oracle complexity of convex optimization problems can be lower bounded by fat-shattering dimension of the associated linear function class (a classic concept from statistical learning theory). Using this lower bound and results from previous chapters we show that for certain classes of convex optimization problems (high dimensional), mirror descent method is near optimal even for offline convex optimization problem. These results are also further used to argue that for certain statistical convex learning problems, mirror descent method is near optimal even when the learner has access to a parallel computation oracles (i.e. oracle query on entire sample is considered as one oracle call in the model). This in turn implies that for these convex learning problems, parallelization does not help.

## 1.3 Main Contributions

### 1. Part I :

#### (a) Statistical learning

- Illustrate limitations of uniform convergence and ERM/SAA
- Provide characterization of statistical learning in general through stability
- Provide universal randomized learning rule

#### (b) Online learning

- Introduce analogs of various complexity measures of hypothesis class for online learning, provide tools martingale uniform convergence theory
- Characterize Online learnability for supervised learning problems
- Generic algorithm for online supervised learning

### 2. Part II :

#### (a) Online convex learning problems

- Characterize learnability and rates for various convex learning problems through notion of martingale type
- Show universality and near optimality of online mirror descent for online convex learning problems

#### (b) Statistical convex learning problems

- Establish lower bounds for oracle complexity and learning rates for various convex learning problems in statistical framework using Rademacher complexity of linear class

- Show that for most reasonable statistical convex learning problems, mirror descent algorithm is near optimal both in terms of learning rates and number of oracle (gradient) access
- (c) Offline convex optimization problems
- Generic lower bound on convex optimization problem by fat-shattering dimension of associated linear class
  - Show that for a large class of large dimensional convex optimization problems, mirror descent is near optimal even for offline convex optimization

## 1.4 Bibliographic Notes

Results in Chapter 3 are from joint work with Shai Shalev-Shwartz, Ohad Shamir and Nati Srebro. Early versions of the results can be found in [2, 3]. See [4] for later version which is closer to the one presented in the chapter. The results in Chapter 4 are from joint work with Alexander Rakhlin and Ambuj Tewari and can be found in [5]. In the second part of the dissertation, few of the result from Chapter 6 can be found in [6]. The results in Chapter 7 is from joint work with Nathan Srebro and Ambuj Tewari. Relating basic concept of martingale type and certain online convex learning problems was first done in [7]. In [6] the result of universality and near optimality of mirror descent method from Chapter 7 is provided. Chapters 8 and 9 are joint work with Nati Srebro.

## **Part I**

# **Statistical and Online Learning : Learnability and Rates**

## Chapter 2

# Preliminary Setup and Notations

In this chapter we provide the basic setup, some preliminary definitions and notations used throughout this dissertation.

### 2.1 General Learning Problem Setup

In all the learning problems we consider instances provided to the learner are chosen from the instance set  $\mathcal{Z}$ . The learner in picks hypotheses from set  $\bar{\mathcal{H}}$ . The instantaneous loss incurred by learner on instance  $z \in \mathcal{Z}$  for picking hypothesis  $\mathbf{h} \in \bar{\mathcal{H}}$  is given by  $\ell(\mathbf{h}, z)$  where  $\ell : \bar{\mathcal{H}} \times \mathcal{Z} \mapsto \mathbb{R}$  is the “loss” or cost function. The exact setup of the learning problem and the goal of the learner depends on whether we consider the statistical learning framework or the online learning framework and will formally be specified later on in the corresponding chapters or sections. However, irrespective of the exact framework, the rough goal of the learner for the learning problems we consider, is to pick hypotheses that are competitive with respect to the best hypothesis from a set of hypotheses  $\mathcal{H} \subseteq \bar{\mathcal{H}}$ . Notice that usually, in most of the learning problems considered in literature, the setting considered is one where learner also picks hypothesis from set  $\mathcal{H}$ . In contrast we shall consider the so called “Improper learning setting” here and allow learner to pick hypothesis from a set  $\bar{\mathcal{H}}$  while the goal is to compete with the best hypothesis from a smaller set  $\mathcal{H}$ . Whenever  $\bar{\mathcal{H}} = \mathcal{H}$  we call this setting as the proper learning setting. This framework is sufficiently general to include a large portion of the learning and optimization problems we are aware of, such as:

- **Binary Classification:** Let  $\mathcal{Z} = \mathcal{X} \times \{0, 1\}$ , let  $\bar{\mathcal{H}}$  be a set of functions  $\mathbf{h} : \mathcal{X} \mapsto \{0, 1\}$ , and let  $\ell(\mathbf{h}; (\mathbf{x}, y)) = \mathbf{1}_{\{\mathbf{h}(\mathbf{x}) \neq y\}}$ . Here, loss function is simply the 0 – 1 loss, measuring whether the binary hypothesis  $\mathbf{h}$  misclassified the example  $(\mathbf{x}, y)$ .
- **Regression:** Let  $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$  where  $\mathcal{X}$  and  $\mathcal{Y}$  are bounded subsets of  $\mathbb{R}^d$  and  $\mathbb{R}$  respectively. Let  $\bar{\mathcal{H}}$  be a set of bounded functions  $\mathbf{h} : \mathcal{X} \mapsto \mathbb{R}$ , and let  $\ell(\mathbf{h}; (\mathbf{x}, y)) = (\mathbf{h}(\mathbf{x}) - y)^2$ . Here, the loss function is

simply the squared loss.

- **Large Margin Classification in a Reproducing Kernel Hilbert Space (RKHS):** Let  $\mathcal{Z} = \mathcal{X} \times \{0, 1\}$ , where  $\mathcal{X}$  is a bounded subset of an RKHS, let  $\bar{\mathcal{H}}$  be another bounded subset of the same RKHS, and let  $\ell(\mathbf{h}; (\mathbf{x}, y)) = \max\{0, 1 - y \langle \mathbf{x}, \mathbf{h} \rangle\}$ . Here, the loss function is the well known hinge loss function, and our goal is to perform margin-based linear classification in the RKHS.
- **K-Means Clustering in Euclidean Space:** Let  $\mathcal{Z} = \mathbb{R}^n$ , let  $\bar{\mathcal{H}}$  be all subsets of  $\mathbb{R}^n$  of size  $k$ , and let  $\ell(\mathbf{h}; z) = \min_{\mathbf{c} \in \mathbf{h}} \|\mathbf{c} - z\|^2$ . Here, each  $\mathbf{h}$  represents a set of  $k$  centroids, and the loss  $\ell$  measures the Euclidean distance squared between an instance  $z$  and its nearest centroid, according to the hypothesis  $\mathbf{h}$ .
- **Density Estimation:** Let  $\mathcal{Z}$  be a subset of  $\mathbb{R}^n$ , let  $\bar{\mathcal{H}}$  be a set of bounded probability densities on  $\mathcal{Z}$ , and let  $\ell(\mathbf{h}; z) = -\log(\mathbf{h}(z))$ . Here, loss function  $\ell$  is simply the negative log-likelihood of an instance  $z$  according to the hypothesis density  $\mathbf{h}$ .
- **Convex Learning Problems:** Let  $\mathcal{Z}$  be an arbitrary measurable set, let  $\bar{\mathcal{H}}$  be a closed, convex subset of a vector space, and for each  $z \in \mathcal{Z}$ , let the function  $\ell(\mathbf{h}; z)$  be convex w.r.t. its first argument.

We shall use the letter  $S$  to denote a sample of instances, that is a sample  $S \in \bigcup_{n \in \mathbb{N}} \mathcal{Z}^n$  is a sequence of instances. For instance  $S = (z_1, \dots, z_n)$  is a sample of size  $n$ , note that the order and multiplicity of instances may be important. Given a sample  $S$  we shall denote the empirical average loss over this sample  $S$  as :

$$L_S(\mathbf{h}) = \frac{1}{|S|} \sum_{i=1}^{|S|} \ell(\mathbf{h}, z_i)$$

We shall also use the notation  $\hat{L}(\mathbf{h})$  to refer  $L_S(\mathbf{h})$  whenever the sample  $S$  is well understood under the context. Further, given a distribution  $\mathcal{D}$  on instance space  $\mathcal{Z}$ , we shall use the notation

$$L_{\mathcal{D}}(\mathbf{h}) = \mathbb{E}_{z \sim \mathcal{D}} [\ell(\mathbf{h}, z)]$$

Further we shall use the notation  $L(\mathbf{h})$  to refer to  $L_{\mathcal{D}}(\mathbf{h})$  whenever the distribution is understood under the context.

## 2.2 More Definitions and Notations

An important object that we will encounter especially while analyzing online learning are trees. Unless specified, all trees considered in this paper are either *rooted complete binary* trees. While it is useful to have the tree picture in mind when reading the paper, it is also necessary to precisely define trees as mathematical objects. We opt for the following definition.

**Definition 1 (Trees).** Given some set  $\mathcal{Z}$ , a  $\mathcal{Z}$ -valued tree  $z$ , of depth  $n$  is a sequence  $(z_1, \dots, z_n)$  of  $n$  mappings  $z_i : \{\pm 1\}^{i-1} \mapsto \mathcal{Z}$ . The *root* of the tree  $z$  is the constant function  $z_1 \in \mathcal{Z}$ .

A tree of infinite depth is defined exactly as above as an infinite sequence  $(z_n)_{n \in \mathbb{N}}$ .

Armed with this definition, we can talk about various operations on trees. For a function  $f : \mathcal{Z} \mapsto \mathcal{U}$ ,  $f(z)$  denotes the  $\mathcal{U}$ -valued tree defined by the mappings  $(f \circ z_1, \dots, f \circ z_n)$ . Analogously, for  $f : \mathcal{Z} \times \mathcal{Z} \mapsto \mathcal{U}$ , the  $\mathcal{U}$ -valued tree  $f(z, z')$  is defined as mappings  $(f(z_1, z'_1), \dots, f(z_n, z'_n))$ . In particular, this defines the usual binary arithmetic operations on real-valued trees. Furthermore, for a class of functions  $\mathcal{F}$  and a tree  $z$ , the projection of  $\mathcal{F}$  onto  $z$  is  $\mathcal{F}(z) = \{f(z) : f \in \mathcal{F}\}$ .

**Definition 2 (Path).** A *path* of length  $n$  is a sequence  $\epsilon = (\epsilon_1, \dots, \epsilon_{n-1}) \in \{\pm 1\}^{n-1}$ .

We shall abuse notation by referring to  $z_i(\epsilon_1, \dots, \epsilon_{i-1})$  by  $z_i(\epsilon)$ . Clearly  $z_i$  only depends on the first  $i - 1$  elements of  $\epsilon$ . We will also refer to  $\epsilon = (\epsilon_1, \dots, \epsilon_n) \in \{\pm 1\}^n$  as a path in a tree of depth  $T$  even though the value of  $\epsilon_T$  is inconsequential. Next we define the notion of subtrees.

**Definition 3 (Subtrees).** The *left subtree*  $z^\ell$  of  $z$  at the root is defined as  $n - 1$  mappings  $(z_1^\ell, \dots, z_{n-1}^\ell)$  with  $z_i^\ell(\epsilon) = z_{i+1}(\{-1\} \times \epsilon)$  for  $\epsilon \in \{\pm 1\}^{n-1}$ . The *right subtree*  $z^r$  is defined analogously by conditioning on the first coordinate of  $z_{i+1}$  to be  $+1$ .

Given two subtrees  $z, v$  of the same depth  $n - 1$  and a constant mapping  $z_1$ , we can *join* the two subtrees to obtain a new set of mappings  $(x_1, \dots, x_n)$  as follows. The root is the constant mapping  $z_1$ . For  $i \in \{2, \dots, n\}$  and  $\epsilon \in \{\pm 1\}^n$ ,  $x_i(\epsilon) = z_{i-1}(\epsilon)$  if  $\epsilon_1 = -1$  and  $w_i(\epsilon) = v_{i-1}(\epsilon)$  if  $\epsilon_1 = +1$ .

We will also need to talk about the values given by the tree  $z$  over all the paths. Formally, let  $\text{Img}(z) = z(\{\pm 1\}^n) = \{x_t(\epsilon) : t \in [n], \epsilon \in \{\pm 1\}^n\}$  be the image of the mappings of  $z$ .

Table 2.1 contains a list of the basic notations used.



$\mathbb{N}$	The set of natural numbers
$\mathbb{R}$	The set of real numbers
$\mathbb{R}_+$	The set of non-negative real numbers
$[n]$	The set $\{1, 2, \dots, n\}$
$\mathbf{1}_{\{A\}}$	1 if predicate $A$ holds and 0 otherwise
$[a]_+$	$\max\{0, a\}$
$x, w$	Scalars
$\mathbf{x}, \mathbf{w}$	Vectors
$\mathbf{X}, \mathbf{W}$	Matrices
$\mathbf{x}[i]$	$i^{\text{th}}$ element of vector $\mathbf{x}$
$\mathbf{X}[i, j]$	$i \times j$ th entry of matrix $\mathbf{X}$
$\mathcal{B}$	Vector space
$\mathcal{B}^*$	Dual of vector space $\mathcal{B}$
$\langle \mathbf{x}, \mathbf{w} \rangle$	Linear functional $\mathbf{x} \in \mathcal{B}^*$ applied to $\mathbf{w} \in \mathcal{B}$
$f, g$	Functions
$\Delta(\mathcal{X})$	Set of Borel probability measures on set $\mathcal{X}$
$f^*$	The Fenchel conjugate of function $f$
$\nabla f(\mathbf{w})$	A sub-gradient of $f$ at $\mathbf{w}$
$\ \cdot\ _*$	Dual norm of the norm $\ \cdot\ $
$\mathbb{E}[Z]$	Expectation of random variable $Z$
$\mathbb{P}[A]$	Probability that event $A$ occurs

Table 2.1: Summary of Notations

## Chapter 3

# Statistical Learning/Optimization

In this chapter we consider the problem of statistical learning in the general learning problems introduced by [8] where we would like to minimize a population risk functional (stochastic objective)

$$L(\mathbf{h}) = \mathbb{E}_{z \sim \mathcal{D}} [\ell(\mathbf{h}; z)] \quad (3.0.1)$$

based on i.i.d. sample  $z_1, \dots, z_n$  drawn from  $\mathcal{D}$  over some target hypothesis class  $\mathcal{H}$ . The distribution  $\mathcal{D}$  over instance space  $\mathcal{Z}$  is unknown to the learner. Notice that this is basically a stochastic optimization problem. In this chapter we are mainly concerned with the question of statistical “learnability”. That is, when can Eq. (3.0.1) be minimized to within arbitrary precision based only on a finite sample  $z_1, \dots, z_n$ , as  $n \rightarrow \infty$ ?


For supervised classification and regression problems, it is well known that a problem is learnable if and only if the empirical risks

$$L_S(\mathbf{h}) = \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{h}, z_i) \quad (3.0.2)$$


for all  $\mathbf{h} \in \mathcal{H}$  converge uniformly to the population risk ([9, 10]). If uniform convergence holds, then the empirical risk minimizer (ERM) is *consistent*, i.e. the population risk of the ERM converges to the optimal population risk, and the problem is learnable using the ERM. We therefore have:

- A necessary and sufficient condition for learnability, namely uniform convergence of the empirical risks. Furthermore, this can be shown to be equivalent to a combinatorial condition: having finite VC-dimension in the case of classification, and having finite fat-shattering dimensions in the case of regression.
- A complete understanding of *how* to learn: since learnability is equivalent to learnability by ERM, we can focus our attention solely on empirical risk minimizers.

The situation, for supervised classification and regression, can be depicted as follows:



In this chapter we start by showing that the situation for general learning problems in the statistical learning framework is actually much more complex. In particular, in Subsection 3.3 we show an example of a learning problem which is learnable (using Stochastic Approximation approach), but is *not* learnable using empirical risk minimization and uniform convergence fails. We discuss how notion stability (through regularization) plays an important role in learnability of the problem. Having shown that uniform convergence fails to characterize learnability for general learning problems we then approach the question of characterizing learnability in general for statistical learning problems. To do so in Section 3.4 we first introduce definitions of stability of learning rules we consider using for characterizing learnability. In section ??, we show that any problem that is learnable, is always learnable with some learning rule which is an “asymptotically ERM”. Moreover, such an AERM must be stable (under a suitable notion of stability). Namely, we provide following characterization of learnability for general statistical learning problems and this can be considered as the highlight of this chapter :



Note that this characterization holds even for learnable problems with no uniform convergence. In this sense, stability emerges as a strictly more powerful notion than uniform convergence for characterizing learnability. Finally in Section 3.6 we show how we can get stronger results by considering randomized learning algorithms and we go on to provide a generic learning rule that is guaranteed to be successful whenever the problem is learnable.

Section 3.7 contains the details of the proofs and technical results used in the chapter. Following that we conclude with Bibliographic notes and Discussion.

### 3.1 The Statistical Learning Problem and Learnability

In the statistical learning problem, instances are drawn i.i.d. from some fixed distribution  $\mathcal{D}$  unknown to the learner. Given a sample  $S = z_1, \dots, z_n$  of size  $n$ , the goal of the learner then is to pick a hypothesis  $\mathbf{h} \in \mathcal{H}$  based only on the sample  $S$  that has small expected loss  $L_{\mathcal{D}}(\mathbf{h})$ . This problem of learning can in turn be phrased as a stochastic optimization problem where our goal is the minimization problem

$$\min_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h})$$

The term  $L_{\mathcal{D}}(\mathbf{h})$  is often referred to as the risk of choosing hypothesis  $\mathbf{h}$ . However given a target hypothesis class  $\mathcal{H}$  our goal is only to do as well as the best hypothesis in this target class and so the sub-optimality (also

referred to as ) of any hypothesis  $\mathbf{h} \in \bar{\mathcal{H}}$  is given by

$$L_{\mathcal{D}}(\mathbf{h}) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}).$$

Of course notice that for proper learning case, the problem is exactly that of stochastic optimization.

To formally refer to the strategy or algorithm used by the learner to pick hypotheses based on sample provided, we now define the notion of learning algorithm for statistical learning problems.

**Definition 4.** A “Statistical Learning Rule”  $\mathbf{A} : \bigcup_{n \in \mathbb{N}} \mathcal{Z}^n \mapsto \bar{\mathcal{H}}$  is a mapping from sequences of instances in  $\mathcal{Z}$  to the set of Hypothesis  $\bar{\mathcal{H}}$ .

We shall refer to any “Learning Rule”  $\mathbf{A}$  that only outputs hypothesis in the set  $\mathcal{H}$  instead of entire  $\bar{\mathcal{H}}$  as a “Proper Learning Rule”.

Learnability deals with the question of when it is even possible (information theoretically) to drive the sub-optimality for a given problems to 0 with increase in sample size. Note that since we are given a randomly drawn sample we shall requires all results in expectation over draw of sample. Later on we illustrate how such a result can be converted to a result that hold with high probability over the sample.

**Definition 5.** We say that a problem is “Statistically Learnable” given a target hypothesis class  $\mathcal{H}$  with rate  $\epsilon_{\text{cons}}(n)$  if there exists a Statistical Learning Rule,  $\mathbf{A}$ , such that :

$$\sup_{\mathcal{D}} \mathbb{E}_{S \sim \mathcal{D}^n} \left[ L_{\mathcal{D}}(\mathbf{A}(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] \leq \epsilon_{\text{cons}}(n)$$

Further as long as  $\epsilon_{\text{cons}}(n) \rightarrow 0$  we simply say that the problem is statistically learnable/optimizable.

Whenever a problem is learnable, we will refer to any learning rule  $\mathbf{A}$  such that

$$\sup_{\mathcal{D}} \mathbb{E}_{S \sim \mathcal{D}^n} \left[ L_{\mathcal{D}}(\mathbf{A}(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] \rightarrow 0$$

as a universally consistent learning rule. This definition of learnability, requiring a uniform rate for all distributions, is the relevant notion for studying learnability of a hypothesis class. It is a direct generalization of agnostic PAC-learnability ([11]) to Vapnik’s General Setting of Learning as studied by [12] and others.

A closely related notion to learnability is that of the sample complexity of a problem. We first define the sample complexity of a given learning rule and then proceed to define sample complexity of a learning problem.

**Definition 6.** Given a learning rule  $\mathbf{A}$ , and  $\epsilon > 0$ , the sample complexity of the rule for  $\epsilon$  is defined as

$$n(\epsilon, \mathbf{A}, \mathcal{Z}) = \inf \left\{ n \in \mathbb{N} \mid \sup_{\mathcal{D}} \mathbb{E}_{S \sim \mathcal{D}^n} \left[ L_{\mathcal{D}}(\mathbf{A}(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] \leq \epsilon \right\}$$

Sample complexity of the learning problem is defined as  $n(\epsilon, \mathcal{Z}) = \inf_{\mathbf{A}} n(\epsilon, \mathbf{A}, \mathcal{Z})$ .

Obviously a problem is learnable if and only if for each  $\epsilon > 0$ ,  $n(\epsilon) < \infty$  and for any universally consistent learning rule  $\mathbf{A}$  and any  $\epsilon > 0$ ,  $n(\epsilon, \mathbf{A}) < \infty$ .

The notion of sample complexity talked about above is the worst case one in the sense that we want learnability with uniform rates, irrespective of the distribution chosen. Sometimes one might have prior knowledge or restrictions on distributions that can occur and so one might only need rates to hold uniformly over a specific family of distributions. To capture this notion we define below sample complexity of a learning problem given a specific family of distributions  $\mathcal{D}$  over instances  $\mathcal{Z}$ .

**Definition 7.** Given a family  $\mathcal{D}$  of Borel distributions over instance space  $\mathcal{Z}$  and an  $\epsilon > 0$ , the sample complexity of any rule  $\mathbf{A}$  is defined as

$$n^{\mathcal{D}}(\epsilon, \mathbf{A}) = \inf \left\{ n \in \mathbb{N} \mid \sup_{\mathcal{D} \in \mathcal{D}} \mathbb{E}_{S \sim \mathcal{D}^n} \left[ L_{\mathcal{D}}(\mathbf{A}(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] \leq \epsilon \right\}$$

Further the sample complexity of the learning problem over this family of distributions  $\mathcal{D}$  is defined as  $n^{\mathcal{D}}(\epsilon, \mathcal{Z}) = \inf_{\mathbf{A}} n^{\mathcal{D}}(\epsilon, \mathbf{A}, \mathcal{Z})$ .

A concept closely related to statistical learning that plays an important role in characterizing learnability is the notion of generalization.

**Definition 8.** A learning rule  $\mathbf{A}$  *generalizes* with rate  $\epsilon_{\text{gen}}(n)$  under distribution  $\mathcal{D}$  if for all  $n \in \mathbb{N}$ ,

$$\mathbb{E}_{S \sim \mathcal{D}^n} [|L(\mathbf{A}(S)) - L_S(\mathbf{A}(S))|] \leq \epsilon_{\text{gen}}(n). \quad (3.1.1)$$

The rule  $\mathbf{A}$  is said to *universally generalize* with rate  $\epsilon_{\text{gen}}(n)$  if it generalizes with rate  $\epsilon_{\text{gen}}(n)$  under all distributions  $\mathcal{D}$  over  $\mathcal{Z}$ .

We note that other authors sometimes define “consistent”, and thus also “learnable” as a combination of our notions of “consistent” and “generalizing”.

## **Empirical Risk Minimization (Sample Average Approximation Approach) and Related Notions :**

Perhaps the most common approach used for learning problems is Empirical Risk Minimization (ERM) as it is referred to in machine learning terminology or Sample Average Approximation (SAA) approach as is widely referred to in the stochastic optimization terminology. The basic idea is for the learning algorithm to return hypothesis in target class  $\mathcal{H}$  that minimizes average loss over sample. Specifically a learning rule  $\mathbf{A}_{\text{ERM}}$  is an *ERM (Empirical Risk Minimizer)* if it minimizes the average loss, that is :

$$L_S(\mathbf{A}_{\text{ERM}}(S)) = L_S(\hat{\mathbf{h}}_S) = \inf_{\mathbf{h} \in \mathcal{H}} L_S(\mathbf{h}). \quad (3.1.2)$$

where we use  $L_S(\hat{\mathbf{h}}_S) = \inf_{\mathbf{h} \in \mathcal{H}} L_S(\mathbf{h})$  to refer to the minimal empirical loss. But since there might be several hypotheses minimizing the empirical risk,  $\hat{\mathbf{h}}_S$  does not refer to a specific hypotheses and there might be many rules which are all ERM's. While Empirical Risk Minimization (or equivalently the Sample Average Approximation approach) is a widely used learning rule, we now introduce a closely related concept of an Asymptotic Empirical Risk Minimizer which will play an important role in characterizing learnability in general problems.

**Definition 9.** We say that a rule  $\mathbf{A}$  is an *AERM (Asymptotic Empirical Risk Minimizer)* with rate  $\epsilon_{\text{erm}}(n)$  under distribution  $\mathcal{D}$  if:

$$\mathbb{E}_{S \sim \mathcal{D}^n} \left[ \hat{L}(\mathbf{A}(S)) - \hat{L}(\hat{\mathbf{h}}_S) \right] \leq \epsilon_{\text{erm}}(n) \quad (3.1.3)$$

Further, a learning rule is *universally an AERM* with rate  $\epsilon_{\text{erm}}(n)$ , if it is an AERM with rate  $\epsilon_{\text{erm}}(n)$  under all distributions  $\mathcal{D}$  over  $\mathcal{Z}$ . A

Yet another closely related notion of an approximate ERM is the following notion of always AERM.

**Definition 10.** A learning rule  $\mathbf{A}$  is an *always AERM* with rate  $\epsilon_{\text{erm}}(n)$ , if for *any* sample  $S$  of size  $n$ , it holds that

$$\hat{L}(\mathbf{A}(S)) - \hat{L}(\hat{\mathbf{h}}_S) \leq \epsilon_{\text{erm}}(n) \quad (3.1.4)$$

## 3.2 Background

### 3.2.1 Learnability and Uniform Convergence

As discussed in the introduction, a central notion for characterizing learnability is uniform convergence. Formally, we say that uniform convergence holds for a learning problem, if the empirical risks of hypotheses in the hypothesis class converges to their population risk uniformly, with a distribution-independent rate:

$$\sup_{\mathcal{D}} \mathbb{E}_{S \sim \mathcal{D}^n} \left[ \sup_{\mathbf{h} \in \mathcal{H}} |L(\mathbf{h}) - L_S(\mathbf{h})| \right] \xrightarrow{n \rightarrow \infty} 0. \quad (3.2.1)$$

It is straightforward to show that if uniform convergence holds, then a problem can be learned with the ERM learning rule.

For binary classification problems (where  $\mathcal{Z} = \mathcal{X} \times \{0, 1\}$ , each hypothesis is a mapping from  $\mathcal{X}$  to  $\{0, 1\}$ , and  $\ell(\mathbf{h}; (\mathbf{x}, y)) = \mathbf{1}_{\{\mathbf{h}(\mathbf{x}) \neq y\}}$ ), [13] showed that the finiteness of a simple combinatorial measure known as the VC-dimension implies uniform convergence. Furthermore, it can be shown that binary classification problems with infinite VC-dimension are not learnable in a distribution-independent sense. This establishes the condition of having finite VC-dimension, and thus also uniform convergence, as a necessary and sufficient condition for learnability.

Such a characterization can also be extended to regression, such as regression with squared loss, where  $\mathbf{h}$  is now a real-valued function, and  $\ell(\mathbf{h}; (\mathbf{x}, y)) = (\mathbf{h}(\mathbf{x}) - y)^2$ . The property of having finite fat-shattering di-




Figure 3.1: An example of a “trivial” learning situation. Each line represents some  $\mathbf{h} \in \mathcal{H}$ , and shows the value of  $\ell(\mathbf{h}, z)$  for all  $z \in \mathcal{Z}$ . The hypothesis  $\tilde{\mathbf{h}}$  dominates any other hypothesis (e.g.,  $\ell(\tilde{\mathbf{h}}; z) < \ell(\mathbf{h}; z)$  uniformly for all  $z$ ), and thus the problem is learnable without uniform convergence or any other property of  $\mathcal{H}$ .

mension at all finite scales now replaces the property of having finite VC-dimension, but the basic equivalence still holds: a problem is learnable if and only if uniform convergence holds ([10], see also [14], Chapter 19). These results are usually based on clever reductions to binary classification. However, the General Learning Setting that we consider is much more general than classification and regression, and includes setting where a reduction to binary classification is impossible.

To justify the necessity of uniform convergence even in the General Learning Setting, Vapnik attempted to show that in this setting, learnability with the ERM learning rule is equivalent to uniform convergence ([15]). Vapnik noted that this result does not hold, due to “trivial” situations. In particular, consider the case where we take an arbitrary learning problem (with hypothesis class  $\mathcal{H}$ ), and add to  $\mathcal{H}$  a single hypothesis  $\tilde{\mathbf{h}}$  such that  $\ell(\tilde{\mathbf{h}}, z) < \inf_{\mathbf{h} \in \mathcal{H}} \ell(\mathbf{h}, z)$  for all  $z \in \mathcal{Z}$  (see figure 3.1 below). This learning problem is now trivially learnable, with the ERM learning rule which always picks  $\tilde{\mathbf{h}}$ . Note that no assumptions whatsoever are made on  $\mathcal{H}$  - in particular, it can be arbitrarily complex, with no uniform convergence or any other particular property. Note also that such a phenomenon is not possible in the binary classification setting, where  $\ell(h; (\mathbf{x}, y)) = \mathbf{1}_{\{h(\mathbf{x}) \neq y\}}$ , since on any  $(x, y)$  we will have hypotheses with  $\ell(h; (\mathbf{x}, y)) = \ell(\tilde{h}; (\mathbf{x}, y))$  and thus if  $\mathcal{H}$  is very complex (has infinite VC dimension) then on every training set there will be many hypotheses with zero empirical error.

To exclude such “trivial” cases, Vapnik introduced a stronger notion of consistency, termed as “strict consistency”, which in our notation is defined as

$$\forall c \in \mathbb{R}, \quad \inf_{\mathbf{h}: L(\mathbf{h}) \geq c} L_S(\mathbf{h}) \xrightarrow{n \rightarrow \infty} \inf_{\mathbf{h}: L(\mathbf{h}) \geq c} L(\mathbf{h}),$$

where the convergence is in probability. The intuition is that we require the empirical risk of the ERM to

converge to the lowest possible risk, even after discarding all the “good” hypotheses whose risk is smaller than some threshold. Vapnik then showed that such strict consistency of the ERM is in fact equivalent to (one-sided) uniform convergence, of the form

$$\sup_{\mathbf{h} \in \mathcal{H}} (L(\mathbf{h}) - L_S(\mathbf{h})) \xrightarrow{n \rightarrow \infty} 0 \quad (3.2.2)$$

in probability. Note that this equivalence holds for every distribution separately, and does not rely on universal consistency of the ERM.

These results seem to imply that up to “trivial” situations, a uniform convergence property indeed characterizes learnability, at least using the ERM learning rule. However, as we will see later on, the situation is in fact not that simple.

### 3.2.2 Various Complexity Measures and Uniform Convergence

We begin by defining the class of functions  $\mathcal{F} \subseteq \mathbb{R}^{\mathcal{Z}}$  we will often refer to the loss class, as :

$$\mathcal{F}(\mathcal{H}, \mathcal{Z}) := \{z \mapsto \ell(\mathbf{h}, z) : \mathbf{h} \in \mathcal{H}\} \quad (3.2.3)$$

We shall often drop the arguments  $\mathcal{H}$  and  $\mathcal{Z}$  and simply use  $\mathcal{F}$  to mean  $\mathcal{F}(\mathcal{H}, \mathcal{Z})$ . Notice that the function class is indexed by hypotheses in class  $\mathcal{H}$ . Note that the notion uniform convergence in Equation 3.2.1 can now be rewritten as :

$$\sup_{\mathcal{D}} \mathbb{E}_{S \sim \mathcal{D}^n} \left[ \sup_{f \in \mathcal{F}} \left( \mathbb{E}[f] - \widehat{\mathbb{E}}_S[f] \right) \right] \xrightarrow{n \rightarrow \infty} 0.$$

Various complexity measures have been introduced in machine learning and empirical process theory literature to bound rates of uniform convergence and thus get upper bounds on learning rates for ERM/SA algorithms. We introduce and discuss about a few of these below. See Appendix ?? for the relationships among these various complexity measures.

**Definition 11** (Statistical Rademacher Complexity). The empirical Rademacher Complexity of a function class  $\mathcal{F} \subset \mathbb{R}^{\mathcal{Z}}$  given a sample  $S = \{z_1, \dots, z_{|S|}\}$  is defined as :

$$\widehat{\mathcal{R}}_S(\mathcal{F}) := \mathbb{E}_{\epsilon} \left[ \sup_{f \in \mathcal{F}} \frac{1}{|S|} \sum_{i=1}^{|S|} \epsilon_i f(z_i) \right]$$

Further for any  $n \in \mathbb{N}$ , we define the worst-case statistical Rademacher complexity as :

$$\mathcal{R}_n^{\text{iid}}(\mathcal{F}) = \sup_{z_1, \dots, z_n \in \mathcal{Z}} \mathbb{E}_{\epsilon} \left[ \sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \epsilon_i f(z_i) \right] \quad (3.2.4)$$



It can be easily shown that for any distribution  $\mathcal{D}$ ,

$$\mathbb{E}_{S \sim \mathcal{D}^n} \left[ \sup_{f \in \mathcal{F}} \left( \mathbb{E}[f] - \widehat{\mathbb{E}}_S[f] \right) \right] \leq 2 \mathbb{E}_{S \sim \mathcal{D}^n} \left[ \widehat{\mathcal{R}}_S(\mathcal{F}) \right] \leq 2 \mathcal{R}_n^{\text{iid}}(\mathcal{F})$$

and so the empirical Rademacher complexity and so also the worst case statistical Rademacher complexity provide upper bound on the rate of uniform convergence over any function class  $\mathcal{F}$ .

Another tool from empirical process theory that is often used to upper bound rates of uniform convergence is covering numbers of the function class. The statistical covering number of a function class is defined below.

**Definition 12** (Statistical Covering Number). A set  $V \subset \mathbb{R}^{\mathcal{Z}}$  is an  $\alpha$ -cover (with respect to  $\ell_p$ -norm) of a function class  $\mathcal{F} \subseteq \mathbb{R}^{\mathcal{Z}}$  on a sample  $S = z_1, \dots, z_n$  if,

$$\forall f \in \mathcal{F}, \forall \epsilon \in \{\pm 1\}^n \exists v \in V \text{ s.t. } \left( \frac{1}{n} \sum_{t=1}^n |v_t - f(z_t)|^p \right)^{1/p} \leq \alpha$$

The *statistical covering number* of a function class  $\mathcal{F}$  on a given sample  $S = z_1, \dots, z_n$  is defined as

$$\mathcal{N}_p^{\text{iid}}(\alpha, \mathcal{F}, S) = \min\{|V| : V \text{ is an } \alpha\text{-cover w.r.t. } \ell_p\text{-norm of } \mathcal{F} \text{ on } S\}.$$

Further define  $\mathcal{N}_p^{\text{iid}}(\alpha, \mathcal{F}, n) = \sup_{S \in \mathcal{Z}^n} \mathcal{N}_p^{\text{iid}}(\alpha, \mathcal{F}, S)$ , the maximal  $\ell_p$  covering number of  $\mathcal{F}$  over samples of size  $n$ .

Pollard's bound [16] and Dudley integral bound [17] can be used to bound rates of uniform convergence in terms of covering numbers. See A.1 in the appendix for a bound on statistical Rademacher complexity in terms of covering numbers through a refined Dudley bound. Yet another combinatorial tool that can be used to bound rates of uniform convergence is the so called fat-shattering dimension defined below.

**Definition 13** (Fat-Shattering Dimension). A sample  $S = z_1, \dots, z_d$  is said to be  $\alpha$ -shattered by a function class  $\mathcal{F} \subseteq \mathbb{R}^{\mathcal{X}}$ , if there exists  $s_1, \dots, s_d \in \mathbb{R}$  such that

$$\forall \epsilon \in \{\pm 1\}^d, \exists f \in \mathcal{F} \text{ s.t. } \forall t \in [d], \epsilon_t(f(z_t) - s_t) \geq \alpha/2$$

The sequence  $s_1, \dots, s_d$  is called the *witness to shattering*. The *statistical fat-shattering dimension*  $\text{fat}_\alpha^{\text{iid}}(\mathcal{F}, \mathcal{Z})$  at scale  $\alpha$  is the largest  $d$  such that  $\mathcal{F}$   $\alpha$ -shatters some sample  $S$  of size  $d$ .

One can upper bound covering numbers in terms of fat-shattering dimension. See Section A.2 of Appendix A for a bound on covering number in terms of fat shattering dimension of function class.

### 3.2.3 Learnability and Stability

Instead of focusing on the hypothesis class, and ensuring uniform convergence of the empirical risks of hypothesis in this class, an alternative approach is to directly control the variance of the learning rule. Here, it is not the complexity of the hypothesis class which matters, but rather the way that the learning rule explores this hypothesis class. This alternative approach leads to the notion of stability in learning. It is important to note that stability is a property of a learning rule, not of the hypothesis class.

In the context of modern learning theory<sup>1</sup>, the use of stability can be traced back at least to the work of [21], which noted that the sensitivity of a learning algorithm with regard to small changes in the sample controls the variance of the leave-one-out estimate. The authors used this observation to obtain generalization bounds (w.r.t. the leave-one-out estimate) for the  $k$ -nearest neighbor algorithm. It is interesting to note that a uniform convergence approach for analyzing this algorithm simply cannot work, because the “hypothesis class” in this case has unbounded complexity. These results were later extended to other “local” learning algorithms (see [22] and references therein). In addition, practical methods have been developed to introduce stability into learning algorithms, in particular the Bagging technique introduced by [23].

Over the last decade, stability was studied as a generic condition for learnability. [24] showed that an algorithm operating on a hypothesis class with finite VC dimension is also stable (under a certain definition of stability). [25] introduced a strong notion of stability (denoted as *uniform stability*) and showed that it is a sufficient condition for learnability, satisfied by popular learning algorithms such as regularized linear classifiers and regressors in Hilbert spaces (including several variants of SVM). [26] introduced several weaker variants of stability, and showed how they are sufficient to obtain generalization bounds for algorithms stable in their sense.

The works cited above mainly considered stability as a *sufficient* condition for learnability. A more recent line of work ([27],[28]) studied stability as a *necessary* condition for learnability. However, the line of argument is specific to settings where uniform convergence holds and is necessary for learning. With this assumption, it is possible to show that the ERM algorithm is stable, and thus stability is also a necessary condition for learning. However, as we will see later on in this chapter, uniform convergence is in fact not necessary for learning in the General Learning Setting, and stability plays there a key role which has nothing to do with uniform convergence.

Finally, it is important to note that the results cited above make use of many different definitions of stability, which unfortunately are not always comparable. All of them measure stability as the amount of change in the algorithm’s output as a function of small changes to the sample on which the algorithm is run. However, “amount of change to the output” and “small changes to the sample” can be defined in many different ways. “Amount of change to the output” can mean change in risk, change in loss with respect to particular examples, or supremum of change in loss over all examples. “Small changes to the sample” usually mean either deleting one example or replacing it with another one (and even here, one can talk about removing/replacing one instance at random, or in some arbitrary manner). Finally, this measure of change can be measured with

---

<sup>1</sup>In a more general mathematical context, stability has been around for much longer. The necessity of stability for so-called inverse problems to be well posed was first recognized by [18]. The idea of regularization (that is, introducing stability into ill-posed inverse problems) became widely known through the works of [19] and [20]. We return to the notion of regularization later on.

respect to any arbitrary sample, in expectation over samples drawn from the underlying distribution; or in high probability over samples. .

### 3.3 Failure of Uniform Convergence and ERM/SAA Approaches

In this section, we study a special case of the General Learning Setting, where there is a real gap between learnability and uniform convergence, in the sense that there are non-trivial problems where no uniform convergence holds (not even in a local sense), but they are still learnable. Moreover, some of these problems are learnable with an ERM (again, without any uniform convergence), and some are not learnable with an ERM, but rather with a different mechanism. We also discuss why this peculiar behavior does not formally contradict Vapnik’s results on the equivalence of strict consistency of the ERM and uniform convergence, as well as the important role that regularization seems to play here, but in a different way than in standard theory.

#### 3.3.1 Learning without Uniform Convergence : Stochastic Convex Optimization

A stochastic convex optimization problem is a special case of the General Learning Setting discussed above, with the added constraints that the objective function  $\ell(\mathbf{h}; z)$  is Lipschitz-continuous and convex in  $\mathbf{h}$  for every  $z$ , and that  $\mathcal{H}$  is closed, convex and bounded. We will focus here on problems where  $\mathcal{H}$  is a subset of a Hilbert space. A special case is the familiar linear prediction setting, where  $z = (\mathbf{x}, y)$  is an instance-label pair, each hypothesis  $\mathbf{h}$  belongs to a subset  $\mathcal{H}$  of a Hilbert space, and  $\ell(\mathbf{h}; \mathbf{x}, y) = \ell(\langle \mathbf{h}, \phi(\mathbf{x}) \rangle, y)$  for some feature mapping  $\phi$  and a loss function  $\ell : \mathbb{R} \times \mathcal{Y} \rightarrow \mathbb{R}$ , which is convex w.r.t. its first argument.

The situation in which the stochastic dependence on  $\mathbf{h}$  is linear, as in the preceding example, is fairly well understood. When the domain  $\mathcal{H}$  and the mapping  $\phi$  are bounded, we have uniform convergence, in the sense that  $|L(\mathbf{h}) - \hat{L}(\mathbf{h})|$  is uniformly bounded over all  $\mathbf{h} \in \mathcal{H}$  (see [29]). This uniform convergence of  $\hat{L}(\mathbf{h})$  to  $L(\mathbf{h})$  justifies choosing the empirical minimizer  $\hat{\mathbf{h}}_S = \arg \min_{\mathbf{h}} \hat{L}(\mathbf{h})$ , and guarantees that the expected value of  $L(\hat{\mathbf{h}}_S)$  converges to the optimal value  $L^* = \inf_{\mathbf{h}} L(\mathbf{h})$ .

Even if the dependence on  $\mathbf{h}$  is not linear, it is still possible to establish uniform convergence (using covering number arguments) provided that  $\mathcal{H}$  is finite dimensional. Unfortunately, when we turn to infinite dimensional hypothesis spaces, uniform convergence might not hold and the problem might not be learnable with empirical minimization. Surprisingly, it turns out that this does not imply that the problem is unlearnable. We will show that using a regularization mechanism, it is possible to devise a learning algorithm for any stochastic convex optimization problem, even when uniform convergence does not hold. This mechanism is fundamentally related to the idea of stability, and will be a good starting point for our more general treatment of stability and learnability in the next section of this chapter.

We now turn to discuss our first concrete example. Consider the convex stochastic optimization problem

given by

$$\ell_{Eq. (3.3.1)}(\mathbf{h}; (\mathbf{x}, \alpha)) = \|\alpha * (\mathbf{h} - \mathbf{x})\| = \sqrt{\sum_i \alpha^2[i] (\mathbf{h}[i] - \mathbf{x}[i])^2}, \quad (3.3.1)$$

where for now we let  $\mathcal{H}$  to be the  $d$ -dimensional unit sphere  $\mathcal{H} = \{\mathbf{h} \in \mathbb{R}^d : \|\mathbf{h}\| \leq 1\}$ , we let  $z = (\mathbf{x}, \alpha)$  with  $\alpha \in [0, 1]^d$  and  $\mathbf{x} \in \mathcal{H}$ , and we define  $u * v$  to be an element-wise product. We will first consider a sequence of problems, where  $d = 2^n$  for any sample size  $n$ , and establish that we cannot expect a convergence rate which is independent of the dimensionality  $d$ . We then formalize this example in infinite dimensions.

One can think of the problem in Eq. (3.3.1) as that of finding the ‘‘center’’ of an unknown distribution over  $\mathbf{x} \in \mathbb{R}^d$ , where we also have stochastic per-coordinate ‘‘confidence’’ measures  $\alpha[i]$ . We will actually focus on the case where some coordinates are missing, namely that  $\alpha[i] = 0$ .

Consider the following distribution over  $(\mathbf{x}, \alpha)$ :  $\mathbf{x} = 0$  with probability one, and  $\alpha$  is uniform over  $\{0, 1\}^d$ . That is,  $\alpha[i]$  are i.i.d. uniform Bernoulli. For a random sample  $(\mathbf{x}_1, \alpha_1), \dots, (\mathbf{x}_n, \alpha_n)$  if  $d > 2^n$  then we have that with probability greater than  $1 - e^{-1} > 0.63$ , there exists a coordinate  $j \in 1 \dots d$  such that all confidence vectors  $\alpha_i$  in the sample are zero on the coordinate  $j$ , that is  $\alpha_i[j] = 0$  for all  $i = 1..n$ . Let  $\mathbf{e}_j \in \mathcal{H}$  be the standard basis vector corresponding to this coordinate. Then

$$\hat{L}_{Eq. (3.3.1)}(\mathbf{e}_j) = \frac{1}{n} \sum_{i=1}^n \|\alpha_i * (\mathbf{e}_j - 0)\| = \frac{1}{n} \sum_{i=1}^n |\alpha_i[j]| = 0,$$

where  $\hat{L}_{Eq. (3.3.1)}(\cdot)$  denotes the empirical risk w.r.t. the function  $\ell_{Eq. (3.3.1)}(\cdot)$ . On the other hand, letting  $L_{Eq. (3.3.1)}(\cdot)$  denote the actual risk w.r.t.  $\ell_{Eq. (3.3.1)}(\cdot)$ , we have

$$L_{Eq. (3.3.1)}(\mathbf{e}_j) = \mathbb{E}_{\mathbf{x}, \alpha} [\|\alpha * (\mathbf{e}_j - 0)\|] = \mathbb{E}_{\mathbf{x}, \alpha} [|\alpha[j]|] = 1/2.$$

Therefore, for any  $n$ , we can construct a convex Lipschitz-continuous objective in a high enough dimension such that with probability at least 0.63 over the sample,  $\sup_{\mathbf{h}} |L_{Eq. (3.3.1)}(\mathbf{h}) - \hat{L}_{Eq. (3.3.1)}(\mathbf{h})| \geq 1/2$ . Furthermore, since  $\ell(\cdot; \cdot)$  is non-negative, we have that  $\mathbf{e}_j$  is an empirical minimizer, but its expected value  $L_{Eq. (3.3.1)}(\mathbf{e}_j) = 1/2$  is far from the optimal expected value  $\min_{\mathbf{h}} L_{Eq. (3.3.1)}(\mathbf{h}) = L_{Eq. (3.3.1)}(0) = 0$ .

To formalize the example in a sample-size independent way, take  $\mathcal{H}$  to be the unit sphere of an infinite-dimensional Hilbert space with orthonormal basis  $\mathbf{e}_1, \mathbf{e}_2, \dots$ , where for  $\mathbf{v} \in \mathcal{H}$ , we refer to its coordinates  $\mathbf{v}[j] = \langle \mathbf{v}, \mathbf{e}_j \rangle$  w.r.t this basis. The confidences  $\alpha$  are now a mapping of each coordinate to  $[0, 1]$ . That is, an infinite sequence of reals in  $[0, 1]$ . The element-wise product operation  $\alpha * \mathbf{v}$  is defined with respect to this basis and the objective function  $\ell_{Eq. (3.3.1)}(\cdot)$  of Eq. (3.3.1) is well defined in this infinite-dimensional space.

We again take a distribution over  $z = (\mathbf{x}, \alpha)$  where  $\mathbf{x} = 0$  and  $\alpha$  is an infinite i.i.d. sequence of uniform Bernoulli random variables (that is, a Bernoulli process with each  $\alpha_i$  uniform over  $\{0, 1\}$  and independent of all other  $\alpha_j$ ). Now, for any finite sample there is almost surely a coordinate  $j$  with  $\alpha_i[j] = 0$  for all  $i$ , and so we a.s. have an empirical minimizer  $\hat{L}_{Eq. (3.3.1)}(\mathbf{e}_j) = 0$  with  $L_{Eq. (3.3.1)}(\mathbf{e}_j) = 1/2 > 0 = L_{Eq. (3.3.1)}(0)$ .

As a result, we see that the empirical values  $\hat{L}_{Eq. (3.3.1)}(\mathbf{h})$  do not converge uniformly to their expectations, and empirical minimization is not guaranteed to solve the problem. Moreover, it is possible to construct a

sharper counterexample, in which the *unique* empirical minimizer  $\hat{\mathbf{h}}_S$  is far from having optimal expected value. To do so, we augment  $\ell_{Eq. (3.3.1)}(\cdot)$  by a small term which ensures its empirical minimizer is unique, and far from the origin. Consider:

$$\ell_{Eq. (3.3.2)}(\mathbf{h}; (\mathbf{x}, \alpha)) = \ell_{Eq. (3.3.1)}(\mathbf{h}; (\mathbf{x}, \alpha)) + \epsilon \sum_i 2^{-i} (\mathbf{h}[i] - 1)^2 \quad (3.3.2)$$

where  $\epsilon = 0.01$ . The objective is still convex and  $(1 + \epsilon)$ -Lipschitz. Furthermore, since the additional term is strictly convex, we have that  $\ell_{Eq. (3.3.2)}(\mathbf{h}; z)$  is strictly convex w.r.t.  $\mathbf{h}$  and so the empirical minimizer is unique.

Consider the same distribution over  $z$ :  $\mathbf{x} = 0$  while  $\alpha[i]$  are i.i.d. uniform zero or one. The empirical minimizer is the minimizer of  $\hat{L}_{Eq. (3.3.2)}(\mathbf{h})$  subject to the constraints  $\|\mathbf{h}\| \leq 1$ . Identifying the solution to this constrained optimization problem is tricky, but fortunately not necessary. It is enough to show that the optimum of the *unconstrained* optimization problem  $\mathbf{h}_{UC}^* = \arg \min \hat{L}_{Eq. (3.3.2)}(\mathbf{h})$  (without constraining  $\mathbf{h} \in \mathcal{H}$ ) has norm  $\|\mathbf{h}_{UC}^*\| \geq 1$ . Notice that in the unconstrained problem, whenever  $\alpha_i[j] = 0$  for all  $i = 1..n$ , only the second term of  $\ell_{Eq. (3.3.2)}$  depends on  $\mathbf{h}[j]$  and we have  $\mathbf{h}_{UC}^*[j] = 1$ . Since this happens a.s. for some coordinate  $j$ , we can conclude that the solution to the constrained optimization problem lies on the boundary of  $\mathcal{H}$ , that is  $\|\hat{\mathbf{h}}_S\| = 1$ . But for such a solution we have

$$L_{Eq. (3.3.2)}(\hat{\mathbf{h}}_S) \geq \mathbb{E}_\alpha \left[ \sqrt{\sum_i \alpha[i] \hat{\mathbf{h}}_S^2[i]} \right] \geq \mathbb{E}_\alpha \left[ \sum_i \alpha[i] \hat{\mathbf{h}}_S^2[i] \right] = \sum_i \hat{\mathbf{h}}_S^2[i] \mathbb{E}_\alpha [\alpha[i]] = \frac{1}{2} \|\hat{\mathbf{h}}_S\|^2 = \frac{1}{2},$$

while  $L^* \leq L(0) = \epsilon$ .

In conclusion, no matter how big the sample size is, the unique empirical minimizer  $\hat{\mathbf{h}}_S$  of the stochastic convex optimization problem in Eq. (3.3.2) is a.s. much worse than the population optimum,  $L(\hat{\mathbf{h}}_S) \geq \frac{1}{2} > \epsilon \geq L^*$ , and certainly does not converge to it.

### Still Learnable Using Stochastic Approximation (SA) Approach :

So far we established that the learning problem specified in Equation 3.3.2 fails to satisfy uniform convergence and further ERM/SAA approach fails to provide any non-trivial guarantee. Specifically we saw that with the particular distribution over  $\alpha$  as the uniform distribution over  $\{\pm 1\}^{\mathbb{N}}$ , we have that  $L(\hat{\mathbf{h}}_S) \geq \frac{1}{2}$  but however  $L^* \leq \epsilon = 0.01$ . However we shall show that this problem is still learnable. Specifically we will see that using the so called Stochastic Approximation (SA) approach, we can get a guarantee over learning rate of order  $1/\sqrt{n}$  for the problem.

The first thing we already noticed was that the problem specified in Equation 3.3.2 is  $(1 + \epsilon)$ -Lipschitz. Further note that the hypothesis set  $\mathcal{H}$  is the unit ball in a Hilbert space and that the problem is convex in its first argument. In the second part of this dissertation we will formally define Stochastic Approximation (SA) approach for convex learning problems and in a more general way than usually described in literature. For now we will informally use the term Stochastic Approximation approach to refer to stochastic gradient

descent (or online gradient descent followed by averaging). To this end consider the learning algorithm for the problem which given sample  $S = \{(\mathbf{x}_1, \alpha_1), \dots, (\mathbf{x}_n, \alpha_n)\}$  is described below :

**Stochastic Gradient Descent for problem Eq. (3.3.2):**

**Initialize**  $\mathbf{h}_1 = \mathbf{0}$  and  $\eta = 1/\sqrt{|S|}$   
**for**  $t = 1$  **to**  $|S|$   
 $\mathbf{h}'_{t+1} = \mathbf{h}_t - \eta \nabla \ell_{Eq. (3.3.2)}(\mathbf{h}_t; (\mathbf{x}_t, \alpha_t))$   
 $\mathbf{h}_{t+1} = \begin{cases} \mathbf{h}'_{t+1} & \text{if } \|\mathbf{h}'_{t+1}\| \leq 1 \\ \frac{\mathbf{h}'_{t+1}}{\|\mathbf{h}'_{t+1}\|} & \text{otherwise} \end{cases}$   
**end for**  
**Return**  $\bar{\mathbf{h}}_S = \frac{1}{|S|} \sum_{t=1}^{|S|} \mathbf{h}_t$ .

Since the problem is convex and  $(1 + \epsilon)$ -Lipschitz and since  $\mathcal{H}$  is the unit ball in the Hilbert space it follows (for instance from [30] + online to batch conversion) that irrespective of which distribution  $\mathcal{D}$  over instances we use,

$$L_{\mathcal{D}}(\bar{\mathbf{h}}_S) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \leq \sqrt{\frac{2(1 + \epsilon)}{n}}$$

Thus we can conclude that the problem is learnable and in fact enjoys a rate of order  $\frac{1}{\sqrt{n}}$ , yet as we already say both uniform convergence and ERM (SAA approach) fails.

### 3.3.2 Learnability via Stability : Role of Regularization

At this point, we have seen an example in the stochastic convex optimization framework where uniform convergence does not hold, and the ERM algorithm fails. Yet we saw that the problem was learnable using SA approach. We will now see an alternate explanation for why the problem is learnable that mainly uses stability to explain the success. We will specifically consider an algorithm that minimizes a regularized average loss and show that this algorithm can guarantee a similar learning rate as the stochastic gradient descent approach. We will show that the regularization induces stability to the learning problem and this stability in turn assures learnability.

Given a stochastic convex optimization problem with an objective function  $\ell(\mathbf{h}; z)$ , consider a *regularized* version of it: instead of minimizing the expected risk  $\mathbb{E}_z [\ell(\mathbf{h}; z)]$  over  $\mathbf{h} \in \mathcal{H}$ , we will try to minimize

$$\mathbb{E}_z \left[ \ell(\mathbf{h}; z) + \frac{\lambda}{2} \|\mathbf{h}\|^2 \right]$$

for some  $\lambda > 0$ . Notice that this is simply a stochastic convex optimization problem w.r.t. the objective function  $\ell(\mathbf{h}; z) + \frac{\lambda}{2} \|\mathbf{h}\|^2$ . We will show that this regularized problem *is* learnable using the ERM algorithm (namely, by attempting to minimize  $\frac{1}{n} \sum_i \ell(\mathbf{h}; z_i) + \frac{\lambda}{2} \|\mathbf{h}\|^2$ ), by showing that the ERM algorithm is *stable*. By taking  $\lambda \rightarrow 0$  at an appropriate rate as the sample size increases, we are able to solve the original stochastic problem optimization problem, w.r.t.  $\ell(\mathbf{h}; z)$ .

The key characteristic of the regularized objective function we need is that it is  $\lambda$ -strongly convex. Formally, we say that a real function  $g(\cdot)$  over a domain  $\mathcal{H}$  in a Hilbert space is  $\lambda$ -strongly convex (where  $\lambda \geq 0$ ), if the function  $g(\cdot) - \frac{\lambda}{2} \|\cdot\|^2$  is convex. In this case, it is easy to verify that if  $\mathbf{h}$  minimizes  $g$  then

$$\forall \mathbf{h}', g(\mathbf{h}') - g(\mathbf{h}) \geq \frac{\lambda}{2} \|\mathbf{h}' - \mathbf{h}\|^2.$$

When  $\lambda = 0$ , strong convexity corresponds to standard convexity. In particular, it is immediate from the definition that  $\ell(\mathbf{h}; z) + \frac{\lambda}{2} \|\mathbf{h}\|^2$  is  $\lambda$ -strongly convex w.r.t.  $\mathbf{h}$  (assuming  $\ell(\mathbf{h}; z)$  is convex).

The arguments above are formalized in the following two theorems:

**Theorem 1.** Consider a stochastic convex optimization problem such that  $\ell(\mathbf{h}; z)$  is  $\lambda$ -strongly convex and  $L$ -Lipschitz with respect to  $\mathbf{h} \in \mathcal{H}$ . Let  $z_1, \dots, z_n$  be an i.i.d. sample and let  $\hat{\mathbf{h}}_S$  be the empirical minimizer. Then, we have that :

$$\mathbb{E}_S \left[ L(\hat{\mathbf{h}}_S) - \inf_{\mathbf{h} \in \mathcal{H}} L(\mathbf{h}) \right] \leq \frac{4L^2}{\lambda n}. \quad (3.3.3)$$

**Theorem 2.** Let  $f : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}$  be such that  $\mathcal{H}$  is bounded by  $B$  and  $\ell(\mathbf{h}, z)$  is convex and  $L$ -Lipschitz with respect to  $\mathbf{h}$ . Let  $z_1, \dots, z_n$  be an i.i.d. sample, let  $\lambda = \sqrt{\frac{16L^2}{B^2 n}}$  and let  $\hat{\mathbf{h}}_\lambda$  be the minimizer of

$$\hat{\mathbf{h}}_\lambda = \min_{\mathbf{h} \in \mathcal{H}} \left( \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{h}, z_i) + \frac{\lambda}{2} \|\mathbf{h}\|^2 \right) \quad (3.3.4)$$

Then, we have that

$$\mathbb{E}_S \left[ L(\hat{\mathbf{h}}_\lambda) - \inf_{\mathbf{h} \in \mathcal{H}} L(\mathbf{h}) \right] \leq 4 \sqrt{\frac{L^2 B^2}{n}} \left( 1 + \frac{8}{n} \right).$$

From the above theorem, we see that regularization is essential for convex stochastic optimization. It is important to note that even for the strongly convex optimization problem in Theorem 1, where the ERM algorithm does work, it is not due to uniform convergence. To see this, consider augmenting the objective function  $\ell_{Eq. (3.3.1)}(\cdot)$  from Eq. (3.3.1) with a strongly convex term:

$$\ell_{Eq. (3.3.5)}(\mathbf{h}; \mathbf{x}, \alpha) = \ell_{Eq. (3.3.1)}(\mathbf{h}; \mathbf{x}, \alpha) + \frac{\lambda}{2} \|\mathbf{h}\|^2. \quad (3.3.5)$$

The modified objective  $\ell_{Eq. (3.3.5)}(\cdot; \cdot)$  is  $\lambda$ -strongly convex and  $(1 + \lambda)$ -Lipschitz over  $\mathcal{H} = \{\mathbf{h} : \|\mathbf{h}\| \leq 1\}$  and thus satisfies the conditions of Theorem 1. Now, consider the same distribution over  $z = (\mathbf{x}, \alpha)$  used earlier:  $\mathbf{x} = 0$  and  $\alpha$  is an i.i.d. sequence of uniform zero/one Bernoulli variables. Recall that almost surely we have a coordinate  $j$  that is never ‘‘observed’’, namely such that  $\forall_i \alpha_i[j] = 0$ . Consider a vector  $t\mathbf{e}_j$  of magnitude  $0 < t \leq 1$  in the direction of this coordinate. We have that  $\hat{L}_{Eq. (3.3.5)}(t\mathbf{e}_j) = \frac{\lambda}{2} t^2$  (where  $\hat{L}_{Eq. (3.3.5)}(\cdot)$  is the empirical risk w.r.t.  $\ell_{Eq. (3.3.5)}(\cdot)$ ) but  $L_{Eq. (3.3.5)}(t\mathbf{e}_j) = \frac{1}{2} t + \frac{\lambda}{2} t^2$ . Hence, letting  $L_{Eq. (3.3.5)}(\cdot)$  denote the risk w.r.t.  $\ell_{Eq. (3.3.5)}(\cdot)$ , we have that  $L_{Eq. (3.3.5)}(t\mathbf{e}_j) - \hat{L}_{Eq. (3.3.5)}(t\mathbf{e}_j) = t/2$ . In particular, we can set  $t = 1$  and establish  $\sup_{\mathbf{h} \in \mathcal{H}} (L_{Eq. (3.3.5)}(\mathbf{h}) - \hat{L}_{Eq. (3.3.5)}(\mathbf{h})) \geq \frac{1}{2}$  regardless of the sample size.

We see then that the empirical averages  $\hat{L}_{Eq. (3.3.5)}(\mathbf{h})$  do not converge uniformly to their expectations.

Moreover, the example above shows that there is no uniform convergence even in a local sense, namely over all hypotheses whose risk is close enough to  $L^*$ , or those close enough to the minimizer of  $\ell_{Eq. (3.3.5)}(\mathbf{h}; \mathbf{x}, \alpha)$ .

### Regularization Vs Constrained Minimization

The technique of regularizing the objective function by adding a “bias” term is old and well known. In particular, adding  $\|\mathbf{h}\|^2$  is the so-called Tikhonov Regularization technique, which has been known for more than half a century (see [19]). However, the role of regularization in our case is very different than in familiar settings such as  $\ell_2$  regularization in SVMs and  $\ell_1$  regularization in LASSO. In those settings regularization serves to constrain our domain to a low-complexity domain (e.g., low-norm predictors), where we rely on uniform convergence. In fact, almost all learning guarantees that we are aware of can be expressed in terms of some sort of uniform convergence.

In our case, constraining the norm of  $\mathbf{h}$  does *not* ensure uniform convergence. Consider the example  $\ell_{Eq. (3.3.1)}(\cdot)$  we have seen earlier. Even over a restricted domain  $\mathcal{H}_r = \{\mathbf{h} : \|\mathbf{h}\| \leq r\}$ , for arbitrarily small  $r > 0$ , the empirical averages  $\hat{L}(\mathbf{h})$  do *not* uniformly converge to  $L(\mathbf{h})$ . Furthermore, consider replacing the regularization term  $\lambda \|\mathbf{h}\|^2$  with a constraint on the norm of  $\|\mathbf{h}\|$ , namely, solving the problem

$$\tilde{\mathbf{h}}_r = \arg \min_{\|\mathbf{h}\| \leq r} \hat{L}(\mathbf{h}) \quad (3.3.6)$$

We cannot solve the stochastic optimization problem by setting  $r$  in a distribution-independent way (i.e., without knowing the solution...). To see this, note that when  $\mathbf{x} = 0$  a.s. we must have  $r \rightarrow 0$  to ensure  $L(\tilde{\mathbf{h}}_r) \rightarrow L^*$ . However, if  $\mathbf{x} = \mathbf{e}_1$  a.s., we must set  $r \rightarrow 1$ . No constraint will work for all distributions over  $\mathcal{Z} = (\mathcal{X}, \alpha)$ ! This sharply contrasts with traditional uses of regularization, where learning guarantees are typically stated in terms of a constraint on the norm rather than in terms of a parameter such as  $\lambda$ , and adding a regularization term of the form  $\frac{\lambda}{2} \|\mathbf{h}\|^2$  is viewed as a proxy for bounding the norm  $\|\mathbf{h}\|$ .

### 3.3.3 Contradiction to Vapnik?

In Subsection 3.2.1, we discussed how Vapnik showed that uniform convergence is in fact necessary for learnability with the ERM. At first glance, this might seem confusing in light of the examples presented above, where we have problems learnable with the ERM without uniform convergence whatsoever.

The solution for this apparent paradox is that our examples are not “strictly consistent” in Vapnik’s sense. Recall that in order to exclude “trivial” cases, Vapnik defined strict consistency of empirical minimization as (in our notation):

$$\forall c \in \mathbb{R}, \quad \inf_{\mathbf{h}: L(\mathbf{h}) \geq c} L_S(\mathbf{h}) \longrightarrow \inf_{\mathbf{h}: L(\mathbf{h}) \geq c} L(\mathbf{h}), \quad (3.3.7)$$

where the convergence is in probability. This condition indeed ensures that  $L(\hat{\mathbf{h}}_S) \xrightarrow{P} L^*$ . Vapnik’s Key Theorem on Learning Theory [15, Theorem 3.1] then states that *strict* consistency of empirical minimization



is equivalent to one-sided<sup>2</sup> uniform convergence. In the example presented above, even though Theorem 1 establishes  $L_{Eq. (3.3.5)}(\hat{\mathbf{h}}_S) \xrightarrow{P} L^*$ , the consistency isn't "strict" by the definition above. To see this, for any  $c > 0$ , consider the vector  $t\mathbf{e}_j$  (where  $\forall_i \alpha_i[j] = 0$ ) with  $t = 2c$ . We have  $L_{Eq. (3.3.5)}(t\mathbf{e}_j) = \frac{1}{2}t + \frac{\lambda}{2}t^2 > c$  but  $\hat{L}_{Eq. (3.3.5)}(t\mathbf{e}_j) = \frac{\lambda}{2}t^2 = 2\lambda c^2$ . Focusing on  $\lambda = \frac{1}{2}$  we get:

$$\inf_{L_{Eq. (3.3.5)}(\mathbf{h}) \geq c} \hat{L}_{Eq. (3.3.5)}(\mathbf{h}) \leq c^2 \quad (3.3.8)$$

almost surely for any sample size  $n$ , violating the strict consistency requirement Eq. (3.3.7).

We emphasize that stochastic convex optimization is far from "trivial" in that there is no dominating hypothesis that will always be selected. Although for convenience of analysis we took  $\mathbf{x} = 0$ , one should think of situations in which  $\mathbf{x}$  is stochastic with an unknown distribution. This shows that uniform convergence is a sufficient, but not at all necessary, condition for consistency of empirical minimization in non-trivial settings.

### 3.4 Stability of Learning Rules

In the previous section, we have shown that in the General Learning Setting, it is possible for problems to be learnable without uniform convergence, in sharp contrast to previously considered settings. The key underlying mechanism which allowed us to learn is stability. In this section, we study the connection between learnability and stability in greater depth, and show that stability can in fact *characterize* learnability. Also, we will see how various "common knowledge facts", which we usually take for granted and are based on a "uniform convergence equivalent to learnability" assumption, do not hold in the General Learning Setting, and things can be much more delicate.

We will refer to settings where learnability is equivalent to uniform convergence as "supervised classification" settings. While supervised classification does not encompass all settings where this equivalence holds, most equivalence results refer to it either explicitly or implicitly (by reduction to a classification problem).

We start by giving the exact definition of the stability notions that we will use. As discussed earlier, there are many possible stability measures, some of which can be used to obtain results of a similar flavor to the ones below. The definition we use seems to be the most convenient for the goal of characterizing learnability in the General Learning Setting. In subsection 3.4.1, we provide a few illustrating examples to the subtle differences that can arise from slight variations in the stability measure.

Our two stability notions are based on replacing one of the training sample instances. For a sample  $S$  of size  $n$ , let  $S^{(i)} = \{z_1, \dots, z_{i-1}, z'_i, z_{i+1}, \dots, z_n\}$  be a sample obtained by replacing the  $i$ -th observation of  $S$  with some different instance  $z'_i$ . When not discussed explicitly, the nature of how  $z'_i$  is obtained should be obvious from context.

---

<sup>2</sup>"One-sided" meaning requiring only  $\sup(L(\mathbf{h}) - L_S(\mathbf{h})) \rightarrow 0$ , rather than  $\sup |L(\mathbf{h}) - L_S(\mathbf{h})| \rightarrow 0$ .

**Definition 14.** A rule  $\mathbf{A}$  is **uniform-RO stable**<sup>3</sup> with rate  $\epsilon_{\text{stable}}(n)$ , if for all possible  $S^{(i)}$  and any  $z' \in \mathcal{Z}$ ,

$$\frac{1}{n} \sum_{i=1}^n \left| \ell(\mathbf{A}(S^{(i)}); z') - \ell(\mathbf{A}(S); z') \right| \leq \epsilon_{\text{stable}}(n).$$

**Definition 15.** A rule  $\mathbf{A}$  is **average-RO stable** with rate  $\epsilon_{\text{stable}}(n)$  under distributions  $\mathcal{D}$  if

$$\left| \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{S \sim \mathcal{D}^n, (z'_1, \dots, z'_n) \sim \mathcal{D}^n} \left[ \ell(\mathbf{A}(S^{(i)}); z'_i) - \ell(\mathbf{A}(S); z'_i) \right] \right| \leq \epsilon_{\text{stable}}(n).$$

Note that this definition corresponds to assuming that the expected empirical risk of the learning rule converges to the expected risk - see Lemma 13.

We say that a rule is *universally stable* with rate  $\epsilon_{\text{stable}}(n)$ , if the stability property holds with rate  $\epsilon_{\text{stable}}(n)$  for all distributions.

**Claim 3.** *Uniform-RO stability with rate  $\epsilon_{\text{stable}}(n)$  implies average-RO stability with rate  $\epsilon_{\text{stable}}(n)$ .*

The following subsection contains a brief literature survey on various definitions of stability and discusses and illustrates the differences between these various notions. The main results in this chapter are only based on the two definitions of stability we introduced so far and the following subsection is mainly for pedantic purposes and the reader may choose to skip it.

### 3.4.1 Comparison with Existing Literature and Other Notions of Stability

The existing literature on stability in learning, briefly surveyed in Subsection 3.2.3, utilizes many different stability measures. All of them measure the amount of change in the algorithm’s output as a function of small changes to the sample on which the algorithm is run. However, they differ in how “output”, “amount of change to the output”, and “small changes to the sample” are defined. In Section ??, we used three stability measures. Roughly speaking, one measure (average-RO stability) is the expected change in the objective value on a particular instance, after that instance is replaced with a different instance. The second measure and third measure (uniform-RO stability and strongly-uniform-RO stability respectively) basically deal with the maximal possible change in the objective value with respect to a particular instance, by replacing a single instance in the training set. However, instead of measuring the objective value on a specific instance, we could have measured the change in the risk of the returned hypothesis, or any other distance measure between hypotheses. Instead of replacing an instance, we could have talked about adding or removing one instance from the sample, either in expectation or in some arbitrary manner. Such variations are common in the literature.

To relate our stability definitions to the ones in the literature, we note that our definitions of uniform-RO stability and strongly-uniform-RO stability are somewhat similar to uniform stability ([25]), which in our

---

<sup>3</sup>RO is short for “replace-one”.

notation is defined as  $\sup_{S,z} \max_i |\ell(\mathbf{A}(S; z)) - \ell(\mathbf{A}(S^{\setminus i}); z)|$ , where  $S^{\setminus i}$  is the training sample  $S$  with instance  $z_i$  removed. Compared to uniform-RO stability, here we measure maximal change over any particular instance, rather than average change over all instances in the training sample. Also, we deal with removing an instance rather than replacing it. Strongly-uniform-RO stability is more similar, with the only formal difference being removal vs. replacement of an instance. However, the results for uniform stability mostly assume deterministic learning rules, while in this work we have used strongly-uniform-RO stability solely in the context of randomized learning rules. For deterministic learning rules, the differences outlined above are sufficient to make uniform stability a strictly stronger requirement than uniform-RO stability, since it is easy to come up with learning problems and (non-symmetric) learning rules which are uniform-RO stable but not uniformly stable. Moreover, we show that uniform-RO stable AERM's characterize learnability, while it is well known that uniformly stable AERM's are not necessary for learnability (see [26]). For the same reason, our notion of strongly-uniform-RO stability is apparently too strong to characterize learnability when we deal with deterministic learning rules, as opposed to randomized learning rules.

Our definition of average-RO stable is similar to “average stability” defined in [27], which in our notation is defined as  $\mathbb{E}_{S \sim \mathcal{D}^n, z_1'} [\ell(\mathbf{A}(S^{(i)}); z_1) - \ell(\mathbf{A}(S); z_1)]$ . Compared to average-RO stability, the main difference is that the change in the objective value is measured with respect to  $z_1$  rather than an average over  $z_i$  for all  $i$ , and stems from the assumption there that the learning algorithm is symmetric. Notice however that we do not make such an assumption.

For an elaborate study on other stability notions and their relationships, see [26].

Unfortunately, many of the stability notions in the literature are incomparable, and even slight changes in the definition radically affect their behavior. WE shall now investigate these differences in more detail.

### LOO Stability vs. RO Stability

The stability definitions we saw introduced up to now were all based on the idea of replacing one instance in the training sample by another instance (e.g., “RO” or “replace-one” stability). An alternative set of definitions can be obtained based on *removing* one instance in the training sample (e.g., “LOO” or “leave-one-out” stability). Despite seeming like a small change, it turns out there is a considerable discrepancy in terms of the obtainable results, compared to RO stability. In this subsection, we wish to discuss these discrepancies, as well as show how small changes to the stability definition can materially affect its strength.

Specifically, we consider the following four LOO stability measures, each slightly weaker than the previous one. The first and last are similar to our notion of uniform-RO stability and average-RO stability respectively. However, we emphasize that RO stability and LOO stability are in general incomparable notions, as we shall see later on. Also, we note that some of these definitions appeared in previous literature. For instance, the notion of “all-i-LOO” below has been studied by several authors under different names [25, 28, 27]. The notation  $S^{\setminus i}$  below refer to a training sample  $S$  with instance  $z_i$  removed.

**Definition 16.** A rule  $\mathbf{A}$  is **uniform-LOO stable** with rate  $\epsilon_{\text{stable}}(n)$  if for all samples  $S$  of  $n$  points and for

all  $i$ :

$$\left| \ell(\mathbf{A}(S^{\setminus i}); z_i) - \ell(\mathbf{A}(S); z_i) \right| \leq \epsilon_{\text{stable}}(n).$$

**Definition 17.** A rule  $\mathbf{A}$  is **all-i-LOO stable** with rate  $\epsilon_{\text{stable}}(n)$  under distribution  $\mathcal{D}$  if for all  $i$ :

$$\mathbb{E}_{S \sim \mathcal{D}^n} \left[ \left| \ell(\mathbf{A}(S^{\setminus i}); z_i) - \ell(\mathbf{A}(S); z_i) \right| \right] \leq \epsilon_{\text{stable}}(n).$$

**Definition 18.** A rule  $\mathbf{A}$  is **LOO stable** with rate  $\epsilon_{\text{stable}}(n)$  under distribution  $\mathcal{D}$  if

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E}_{S \sim \mathcal{D}^n} \left[ \left| \ell(\mathbf{A}(S^{\setminus i}); z_i) - \ell(\mathbf{A}(S); z_i) \right| \right] \leq \epsilon_{\text{stable}}(n).$$

**Definition 19.** A rule  $\mathbf{A}$  is **on-average-LOO stable** with rate  $\epsilon_{\text{stable}}(n)$  under distribution  $\mathcal{D}$  if

$$\left| \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{S \sim \mathcal{D}^n} \left[ \ell(\mathbf{A}(S^{\setminus i}); z_i) - \ell(\mathbf{A}(S); z_i) \right] \right| \leq \epsilon_{\text{stable}}(n).$$

While some of the definitions above might look rather similar, we show below that each one is strictly stronger than the other. Example 3 is interesting in its own right, since it presents a learning problem and an AERM that is universally consistent, but not LOO stable. While this is possible in the General Learning Setting, in supervised classification every such AERM has to be LOO stable (this is essentially proven in [28]).

**Example 1.** There exists a learning problem with a universally consistent and all-i-LOO stable learning rule, but there is no universally consistent and uniform LOO stable learning rule.

*Proof.* This example is taken from [26]. Consider the hypothesis space  $\{0, 1\}$ , the instance space  $\{0, 1\}$ , and the objective function  $\ell(\mathbf{h}, z) = |h - z|$ .

It is straightforward to verify that an ERM is a universally consistent learning rule. It is also universally all-i-LOO stable, because removing an instance can change the hypothesis only if the original sample had an equal number of 0's and 1's (plus or minus one), which happens with probability at most  $O(1/\sqrt{n})$  where  $n$  is the sample size. However, it is not hard to see that the only uniform LOO stable learning rule, at least for large enough sample sizes, is a constant rule which always returns the same hypothesis  $h$  regardless of the sample. Such a learning rule is obviously not universally consistent.  $\square$

**Example 2.** There exists a learning problem with a universally consistent and LOO-stable AERM, which is not symmetric and is not all-i-LOO stable.

*Proof.* Let the instance space be  $[0, 1]$ , the hypothesis space  $[0, 1] \cup 2$ , and the objective function  $\ell(h, z) = \mathbf{1}_{\{h=z\}}$ . Consider the following learning rule  $\mathbf{A}$ : given a sample, check if the value  $z_1$  appears more than once in the sample. If no, return  $z_1$ , otherwise return 2.

Since  $L_S(2) = 0$ , and  $z_1$  returns only if this value constitutes  $1/n$  of the sample, the rule above is an AERM with rate  $\epsilon_{\text{erm}}(n) = 1/n$ . To see universal consistency, let  $\Pr z_1 = p$ . With probability  $(1-p)^{n-2}$ ,  $z_1 \notin \{z_2, \dots, z_n\}$ , and the returned hypothesis is  $z_1$ , with  $L(z_1) = p$ . Otherwise, the returned hypothesis is 2, with  $L(2) = 0$ . Hence  $\mathbb{E}_S [L(\mathbf{A}(S))] \leq p(1-p)^{n-2}$ , which can be easily verified to be at most  $1/(n-1)$ , so the learning rule is consistent with rate  $\epsilon_{\text{cons}}(n) \leq 1/(n-1)$ . To see LOO-stability, notice that our learning hypothesis can change by deleting  $z_i, i > 1$ , only if  $z_i$  is the only instance in  $z_2, \dots, z_n$  equal to  $z_1$ . So  $\epsilon_{\text{stable}}(n) \leq 2/n$  (in fact, LOO-stability holds even without the expectation). However, this learning rule is not all-i-LOO-stable. For instance, for any continuous distribution,  $|\ell(\mathbf{A}(S^{\setminus 1}), z_1) - \ell(\mathbf{A}(S), z_1)| = 1$  with probability 1, so it obviously cannot be all-i-LOO-stable with respect to  $i = 1$ .  $\square$

**Example 3.** There exists a learning problem with a universally consistent (and on-average-LOO stable) AERM, which is not LOO stable.

*Proof.* Let the instance space, hypothesis space and objective function be as in Example 1. Consider the following learning rule, based on a sample  $S = (z_1, \dots, z_n)$ : if  $\sum_i \mathbf{1}_{\{z_i=1\}}/n > 1/2 + \sqrt{\log(4)/2n}$ , return 1. If  $\sum_i \mathbf{1}_{\{z_i=1\}}/n < 1/2 - \sqrt{\log(4)/2n}$ , return 0. Otherwise, return  $\text{Parity}(S) = (z_1 + \dots + z_n) \bmod 2$ .

This learning rule is an AERM, with  $\epsilon_{\text{erm}}(n) = \sqrt{2\log(4)/n}$ . Since we have only two hypotheses, we have uniform convergence of  $L_S(\cdot)$  to  $L(\cdot)$  for any hypothesis. Therefore, our learning rule universally generalizes (with rate  $\epsilon_{\text{gen}}(n) = \sqrt{\log(4/\delta)/2n}$ ), and by Theorem 7, this implies that the learning rule is also universally consistent and on-average-LOO stable.

However, the learning rule is not LOO stable. Consider the uniform distribution on the instance space. By Hoeffding's inequality,  $|\sum_i \mathbf{1}_{\{z_i=1\}}/n - 1/2| \leq \sqrt{\log(4)/2n}$  with probability at least  $1/2$  for any sample size  $n$ . In that case, the returned hypothesis is the parity function (even when we remove an instance from the sample, assuming  $n \geq 3$ ). When this happens, it is not hard to see that for any  $i$ ,

$$\ell(\mathbf{A}(S), z_i) - \ell(\mathbf{A}(S^{\setminus i}), z_i) = \mathbf{1}_{\{z_i=1\}}(-1)^{\text{Parity}(S)}.$$

This implies that

$$\begin{aligned} & \mathbb{E} \left[ \frac{1}{n} \sum_{i=1}^n \left| \ell(\mathbf{A}(S^{\setminus i}); z_i) - \ell(\mathbf{A}(S); z_i) \right| \right] \\ & \geq \frac{1}{2} \mathbb{E} \left[ \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\{z_i=1\}} \left| \sqrt{\frac{\log(4)}{2n}} \geq \left| \sum_{i=1}^n \frac{\mathbf{1}_{\{z_i=1\}}}{n} - \frac{1}{2} \right| \right] \right] \\ & \geq \frac{1}{2} \left( \frac{1}{2} - \sqrt{\frac{\log(4)}{2n}} \right) \rightarrow \frac{1}{4}, \end{aligned} \tag{3.4.1}$$

which does not converge to zero with the sample size  $n$ . Therefore, the learning rule is not LOO stable.  $\square$

Note that the proof implies that on-average-LOO stability cannot be replaced even by something between

on-average-LOO stability and LOO stability. For instance, a natural candidate would be

$$\mathbb{E}_{S \sim \mathcal{D}^n} \left[ \left| \frac{1}{n} \sum_{i=1}^n \left( \ell(\mathbf{A}(S \setminus^i); z_i) - \ell(\mathbf{A}(S); z_i) \right) \right| \right], \quad (3.4.2)$$

where the absolute value is now over the entire sum, but inside the expectation. In the example used in the proof, Eq. (3.4.2) is still lower bounded by Eq. (3.4.1), which does not converge to zero with the sample size.

After showing that the hierarchy of definitions above is indeed strict, we turn to the question of what can be characterized in terms of LOO stability. In [31], we show a version of Theorem 5, which asserts that a problem is learnable if and only if there is an on-average-LOO stable AERM. However, on-average-LOO stability is qualitatively much weaker than the notion of uniform-RO stability used in Theorem 5 (see Definition 14). Rather, we would expect to prove a version of the theorem with the notion of uniform-LOO stability or at least LOO stability, which are more analogous to uniform-RO stability. However, the proof of Theorem 5 does not work for these stability definitions (technically, this is because the proof relies on the sample size remaining constant, which is true for replacement stability, but not when we remove an instance as in LOO stability). We do not know if one can prove a version of Theorem 5 with an LOO stability notion stronger than on-average-LOO stability.

On the plus side, LOO stability allows us to prove the following interesting result, specific to ERM learning rules.

**Theorem 4.** *For an ERM the following are equivalent:*

- *Universal LOO stability.*
- *Universal consistency.*
- *Universal generalization.*

In particular, the theorem implies that LOO stability is a necessary property for consistent ERM learning rules. This parallels Theorem 7, which dealt with AERM’s in general, and used RO stability. As before, we do not know how to obtain something akin to Theorem 7 with RO stability.

### 3.5 Characterizing Learnability : Main Results

Our overall goal is to characterize learnable problems (namely, problems for which there exists a universally consistent learning rule, as in Eq. (??)). That means finding some condition which is both *necessary* and *sufficient* for learnability. In the uniform convergence setting, such a condition is the stability of the ERM (under any of several possible stability measures, including both variants of RO-stability defined above). This is still sufficient for learnability in the General Learning Setting, but far from being necessary, as we have seen in Section ??.

The most important result in this section is a condition which is necessary and sufficient for learnability in the General Learning Setting:

**Theorem 5.** *A learning problem is learnable if and only if there exists a uniform-RO stable, universally AERM learning rule.*

*In particular, if there exists a  $\epsilon_{\text{cons}}(n)$ -universally consistent rule, then there exists a rule that is  $\epsilon_{\text{stable}}(n)$ -uniform-RO stable and universally  $\epsilon_{\text{erm}}(n)$ -AERM where:*

$$\begin{aligned}\epsilon_{\text{erm}}(n) &= 3\epsilon_{\text{cons}}(n^{1/4}) + \frac{8B}{\sqrt{n}}, \\ \epsilon_{\text{stable}}(n) &= \frac{2B}{\sqrt{n}}.\end{aligned}\tag{3.5.1}$$

*In the opposite direction, if a learning rule is  $\epsilon_{\text{stable}}(n)$ -uniform-RO stable and universally  $\epsilon_{\text{erm}}(n)$ -AERM, then it is universally consistent with rate*

$$\epsilon_{\text{cons}}(n) \leq \epsilon_{\text{stable}}(n) + \epsilon_{\text{erm}}(n)$$

Thus, while we have seen in Section ?? that the ERM rule might fail for learning problems which are in fact learnable, there is always an AERM rule which will work. In other words, when designing learning rules, we might need to look beyond empirical risk minimization, but not beyond AERM learning rules. On the downside, we must choose our AERM carefully, since not any AERM will work. This contrasts with supervised classification, where any AERM will work if the problem is learnable at all.

How do we go about proving this assertion? The easier part is showing sufficiency. Namely, that a stable AERM must be consistent (and generalizing). In fact, this holds both separately for any particular distribution  $\mathcal{D}$ s, and uniformly over all distributions:

**Theorem 6.** *If a rule is an AERM with rate  $\epsilon_{\text{erm}}(n)$  and average-RO stable (or uniform-RO stable) with rate  $\epsilon_{\text{stable}}(n)$  under  $\mathcal{D}$ , then it is consistent and generalizes under  $\mathcal{D}$  with rates*

$$\begin{aligned}\epsilon_{\text{cons}}(n) &\leq \epsilon_{\text{stable}}(n) + \epsilon_{\text{erm}}(n) \\ \epsilon_{\text{gen}}(n) &\leq \epsilon_{\text{stable}}(n) + 2\epsilon_{\text{erm}}(n) + \frac{2B}{\sqrt{n}}\end{aligned}$$

The second part of Theorem 5 follows as a direct corollary. We note that close variants of Theorem 6 has already appeared in previous literature (e.g., [28] and [27]).

The harder part is showing that a uniform-RO stable AERM is *necessary* for learnability. This is done in several steps.

First, we show that consistent AERMs have to be average-RO stable:

**Theorem 7.** *For an AERM, the following are equivalent:*

- *Universal average-RO stability.*
- *Universal consistency.*
- *Universal generalization.*

The exact conversion rate of Theorem 7 is specified in the corresponding proof (Sub-section 3.7), and are all polynomial. In particular, an  $\epsilon_{\text{cons}}$ -universal consistent  $\epsilon_{\text{erm}}$ -AERM is average-RO stable with rate

$$\epsilon_{\text{stable}}(n) \leq \epsilon_{\text{erm}}(n) + 3\epsilon_{\text{cons}}(n^{1/4}) + \frac{4B}{\sqrt{n}}.$$

Next, we show that if we seek universally consistent and generalizing learning rules, then we must consider only AERMs:

**Theorem 8.** *If a rule  $\mathbf{A}$  is universally consistent with rate  $\epsilon_{\text{cons}}(n)$  and generalizing with rate  $\epsilon_{\text{gen}}(n)$ , then it is universally an AERM with rate*

$$\epsilon_{\text{erm}}(n) \leq \epsilon_{\text{gen}}(n) + 3\epsilon_{\text{cons}}(n^{1/4}) + \frac{4B}{\sqrt{n}}$$

Now, recall that learnability is defined as the existence of some universally consistent learning rule. Such a rule might not be generalizing, stable or even an AERM (see example 5 below). However, it turns out that if a universally consistent learning rule exist, then there is *another* learning rule for the same problem, which is generalizing (Lemma 22). Thus, by Theorems 7-8, this rule must also be average-RO stable AERM. In fact, by another application of Lemma 22, such an AERM must also be uniform-RO stable, leading to Theorem 5.

## 3.6 Randomization, Convexification, and a Generic Learning Rule

In this section we show that when one considers randomized learning rules, it is possible to get stronger results and we even propose a generic randomized algorithm for statistical learning problem that guarantees non-trivial learning rate whenever the problem is learnable.

### 3.6.1 Stronger Results with Randomized Learning Rules

The strongest result we were able to obtain for characterizing learnability so far is Theorem 5, which stated that a problem is learnable if and only if there exists a universally uniform-RO stable AERM. In fact, this result was obtained under the assumption that the learning rule  $\mathbf{A}$  is deterministic: given a fixed sample  $S$ ,  $\mathbf{A}$  returns a single specific hypothesis  $\mathbf{h}$ . However, we might relax this assumption and also consider *randomized* learning rules: given any fixed  $S$ ,  $\mathbf{A}(S)$  returns a distribution over the hypothesis class  $\mathcal{H}$ .

With this relaxation, we will see that we can obtain a stronger version of Theorem 5, and even provide a generic learning algorithm (at least for computationally unbounded learners) which successfully learns any learnable problem.

To simplify notation, we will override the notations  $\ell(\mathbf{A}(S), z)$ ,  $L(\mathbf{A}(S))$  and  $L_S(\mathbf{A}(S))$  to mean  $\mathbb{E}_{\mathbf{h} \sim \mathbf{A}(S)} [\ell(\mathbf{h}, z)]$ ,  $\mathbb{E}_{\mathbf{h} \sim \mathbf{A}(S)} [L(\mathbf{h})]$  and  $\mathbb{E}_{\mathbf{h} \sim \mathbf{A}(S)} [L_S(\mathbf{h})]$ . In other words,  $\mathbf{A}$  returns a distribution over  $\mathcal{H}$  and  $\ell(\mathbf{A}(S), z)$  for some fixed  $S, z$  is the expected loss of a random hypothesis picked according to that distribution, with



respect to  $z$ . Similarly,  $L(\mathbf{A}(S))$  for some fixed  $S$  is the expected generalization error, and  $L_S(\mathbf{A}(S))$  is the expected empirical risk on the fixed sample  $S$ . With this slight abuse of notation, all our previous definitions hold. For instance, we still define a learning rule  $\mathbf{A}$  to be consistent with rate  $\epsilon_{\text{cons}}(n)$  if  $\mathbb{E}_{S \sim \mathcal{D}^n} [L(\mathbf{A}(S)) - L^*] \leq \epsilon_{\text{cons}}(n)$ , only now we actually mean

$$\mathbb{E}_{S \sim \mathcal{D}^n} [\mathbb{E}_{\mathbf{h} \sim \mathbf{A}(S)} [L(\mathbf{h}) - L^*]] \leq \epsilon_{\text{cons}}(n).$$

The definitions for AERM, generalization etc. also hold with this subtle change in meaning.

An alternative way to view randomization is as a method to *linearize* the learning problem. In other words, randomization implicitly replaces the arbitrary hypothesis class  $\mathcal{H}$  by the space of probability distributions over  $\mathcal{H}$ ,

$$\mathcal{M} = \left\{ \alpha : \mathcal{H} \rightarrow [0, 1] \text{ s.t. } \int \alpha[\mathbf{h}] = 1 \right\},$$

and replaces the arbitrary function  $\ell(\mathbf{h}; z)$  by a *linear* function in its first argument

$$\ell(\alpha; z) = \mathbb{E}_{\mathbf{h} \sim \alpha} [\ell(\mathbf{h}, z)] = \int \ell(\mathbf{h}; z) \alpha[\mathbf{h}].$$

Linearity of the loss and convexity of  $\mathcal{M}$  are the key mechanism which allows us to obtain our stronger results. Moreover, if the learning problem is already convex (i.e.,  $f$  is convex and  $\mathcal{H}$  is convex), we can achieve the same results using a deterministic learning rule, as the following claim demonstrates:

**Claim 9.** *Assume that the hypothesis class  $\mathcal{H}$  is convex subset of a vector space, such that  $\mathbb{E}_{\mathbf{h} \sim \mathbf{A}(S)} [\mathbf{h}]$  is a well-defined element of  $\mathcal{H}$  for any  $S$ . Moreover, assume that  $\ell(\mathbf{h}; z)$  is convex in  $\mathbf{h}$ . Then from any (possibly randomized) learning rule  $\mathbf{A}$ , it is possible to construct a deterministic learning rule  $\mathbf{A}'$ , such that  $\ell(\mathbf{A}'(S), z) \leq \ell(\mathbf{A}(S), z)$  for any  $S, z$ . As a result, it also holds that  $L_S(\mathbf{A}'(S)) \leq L_S(\mathbf{A}(S))$  and  $L(\mathbf{A}'(S)) \leq L(\mathbf{A}(S))$ .*

*Proof.* Given a sample  $S$ , define  $\mathbf{A}'(S; z)$  as the single hypothesis  $\mathbb{E}_{\mathbf{h} \sim \mathbf{A}(S)} [\mathbf{h}]$ . The proof of the theorem is immediate by Jensen's inequality: since  $\ell(\cdot)$  is convex in its first argument,

$$\ell(\mathbf{A}'(S); z) = \ell(\mathbb{E}_{\mathbf{h} \sim \mathbf{A}(S)} [\mathbf{h}], z) \leq \mathbb{E}_{\mathbf{h} \sim \mathbf{A}(S)} [\ell(\mathbf{h}, z)],$$

where the r.h.s. is in fact  $\ell(\mathbf{A}(S), z)$  by the abuse of notation we have defined previously.  $\square$

Although linearization is the real mechanism at play here, we find it more convenient to display our results and proofs in the language of randomized learning rules.

Allowing randomization allows us to obtain results with respect to the following very strong notion of stability<sup>4</sup>:

<sup>4</sup>This definition of stability is very similar to the so-called "uniform stability", discussed in [25], although [25] consider deterministic learning rules.

**Definition 20.** A rule  $\mathbf{A}$  is **strongly-uniform-RO stable** with rate  $\epsilon_{\text{stable}}(n)$  if for all samples  $S$  of  $n$  points, for all  $i$ , and any  $z', z'_i \in \mathcal{Z}$ , it holds that

$$\left| \ell(\mathbf{A}(S^{(i)}); z') - \ell(\mathbf{A}(S); z') \right| \leq \epsilon_{\text{stable}}(n).$$

The strengthening of Theorem 5 that we will prove here is the following:

**Theorem 10.** *A learning problem is learnable if and only if there exists a (possibly randomized) learning rule which is an always AERM and strongly-uniform-RO stable.*

Compared to Theorem 5, we have replaced universal AERM by the stronger notion of an always AERM, and uniform-RO stability by strongly-uniform-RO stability. This makes the result strong enough to formulate a generic learning algorithm, as we will see later on.

The theorem is an immediate consequence of Theorem 5 and the following lemma:

**Lemma 11.** *For any deterministic learning rule  $\mathbf{A}$ , there exists a randomized learning rule  $\mathbf{A}'$  such that:*

- *For any  $\mathcal{D}$ , if  $\mathbf{A}$  is  $\epsilon_{\text{cons}}$ -consistent under  $\mathcal{D}$  then  $\mathbf{A}'$  is  $\epsilon_{\text{cons}}(\lfloor \sqrt{n} \rfloor)$  consistent under  $\mathcal{D}$ .*
- *$\mathbf{A}'$  universally generalizes with rate  $4B/\sqrt{n}$ .*
- *If  $\mathbf{A}$  is uniform-RO stable with rate  $\epsilon_{\text{stable}}(n)$ , then  $\mathbf{A}'$  is strongly-uniform-RO stable with rate  $\epsilon_{\text{stable}}(\lfloor \sqrt{n} \rfloor)$ .*
- *If  $\mathbf{A}$  is universally  $\epsilon_{\text{cons}}$ -consistent, then  $\mathbf{A}'$  is an always AERM with rate  $2\epsilon_{\text{cons}}(\lfloor \sqrt{n} \rfloor)$ .*

*Moreover,  $\mathbf{A}'$  is a symmetric learning rule (it does not depend on the order of elements in the sample on which it is applied).*

### 3.6.2 A Generic Learning Rule

Recall that a symmetric learning rule  $\mathbf{A}$  is such that  $\mathbf{A}(S) = \mathbf{A}(S')$  whenever  $S, S'$  are identical samples up to permutation. When we deal with randomized learning rules, we assume that the distribution of  $\mathbf{A}(S)$  is identical to the distribution of  $\mathbf{A}(S')$ . Also, let  $\bar{\mathcal{H}}$  denote the set of all distributions on  $\mathcal{H}$ . An element  $\bar{\mathbf{h}} \in \bar{\mathcal{H}}$  will be thought of as a possible outcome of a randomized learning rule.

Consider the following learning rule: given a sample size  $n$ , find a minimizer over all symmetric<sup>5</sup> functions  $\mathbf{A} : \mathcal{Z}^n \rightarrow \bar{\mathcal{H}}$  of

$$\sup_{S \in \mathcal{Z}^n} \left( L_S(\mathbf{A}(S)) - L_S(\hat{\mathbf{h}}_S) \right) + \sup_{S \in \mathcal{Z}^n, z'} \left| \ell(\mathbf{A}(S); z') - \ell(\mathbf{A}(S^{(i)}); z') \right|, \quad (3.6.1)$$

with  $i$  being an arbitrary fixed element in  $\{1, \dots, n\}$ . Once such a function  $\mathbf{A}$  is found, return  $\mathbf{A}_n(S)$ .

<sup>5</sup>The algorithm would still work, with slight modifications, if we minimize over all functions - symmetric or not. However, the search space would be larger.

**Theorem 12.** *If a learning problem is learnable (namely, there exist a universally consistent learning rule with rate  $\epsilon_{\text{cons}}(n)$ ), the learning algorithm described above is universally consistent with rate*

$$4\epsilon_{\text{cons}}(\lfloor \sqrt{n} \rfloor) + \frac{8B}{\sqrt{n}}.$$

The main drawback of the algorithm we described is that it is completely infeasible: in practice, we cannot hope to efficiently perform minimization of Eq. (3.6.1) over all functions from  $\mathcal{Z}^n$  to  $\mathcal{H}$ . Nevertheless, we believe it is conceptually important for three reasons: First, it hints that generic methods to develop learning algorithms might be possible in the General Learning Setting (similar to the more specific supervised classification setting); Second, it shows that stability might play a crucial role in the way such methods will work; And third, that stability might act in a similar manner to regularization. Indeed, Eq. (3.6.1) can be seen as a “regularized ERM” in the space of learning rules (i.e., functions from samples to hypotheses): if we take just the first term in Eq. (3.6.1),  $\sup_{S \in \mathcal{Z}^n} (L_S(\mathbf{A}(S)) - L_S(\hat{\mathbf{h}}_S))$ , then its minimizer is trivially the ERM learning rule. If we take just the second term in Eq. (3.6.1),  $\sup_{S \in \mathcal{Z}^n, z} |\ell(\mathbf{A}(S); z') - \ell(\mathbf{A}(S^{(i)}); z')|$ , then its minimizers are trivial learning rules which return the same hypothesis irrespective of the training sample. Minimizing a sum of both terms forces us to choose a learning rule which is an “almost”-ERM but also stable - a learning rule which must exist if the problem is learnable at all, as Theorem 10 proves.

## 3.7 Detailed Results and Proofs

### 3.7.1 Detailed Proof of Main Result (Section 3.5)

In this subsection we provide proofs for the main results contained in Section 3.5. We first establish that for AERMs, average-RO stability and generalization are equivalent.

#### Equivalence of Stability and Generalization

It will be convenient to work with a weaker version of generalization as an intermediate step: We say a rule  $\mathbf{A}$  **on-average generalizes** with rate  $\epsilon_{\text{oag}}(n)$  under distribution  $\mathcal{D}$  if for all  $n$ ,

$$|\mathbb{E}_{S \sim \mathcal{D}^n} [L(\mathbf{A}(S)) - L_S(\mathbf{A}(S))]| \leq \epsilon_{\text{oag}}(n). \quad (3.7.1)$$

It is straightforward to see that generalization implies on-average generalization with the same rate. We show that for AERMs, the converse is also true, and also that on-average generalization is equivalent to average-RO stability. This establishes the equivalence between generalization and average-RO stability (for AERMs).

**Lemma 13 (on-average generalization  $\Leftrightarrow$  average-RO stability).** *If  $\mathbf{A}$  is on-average generalizing with rate  $\epsilon_{\text{oag}}(n)$  then it is average-RO stable with rate  $\epsilon_{\text{oag}}(n)$ . If  $\mathbf{A}$  is average-RO stable with rate  $\epsilon_{\text{stable}}(n)$  then it is on-average generalizing with rate  $\epsilon_{\text{stable}}(n)$ .*

*Proof.* For any  $i$ ,  $z_i$  and  $z'_i$  are both drawn i.i.d. from  $\mathcal{D}$ , we have that

$$\mathbb{E}_{S \sim \mathcal{D}^n} [\ell(\mathbf{A}(S); z_i)] = \mathbb{E}_{S \sim \mathcal{D}^n, z'_i \sim \mathcal{D}} [\ell(\mathbf{A}(S^{(i)}); z'_i)].$$

Hence,

$$\begin{aligned} \mathbb{E}_{S \sim \mathcal{D}^n} [L_S(\mathbf{A}(S))] &= \mathbb{E}_{S \sim \mathcal{D}^n} \left[ \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{A}(S); z_i) \right] \\ &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{S \sim \mathcal{D}^n} [\ell(\mathbf{A}(S); z_i)] \\ &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{S \sim \mathcal{D}^n, z'_i \sim \mathcal{D}} [\ell(\mathbf{A}(S^{(i)}); z'_i)] \end{aligned}$$

Also note that  $L(\mathbf{A}(S)) = \mathbb{E}_{z'_i \sim \mathcal{D}} [\ell(\mathbf{A}(S); z'_i)] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{z'_i \sim \mathcal{D}} [\ell(\mathbf{A}(S); z'_i)]$ . Hence we can conclude that

$$\mathbb{E}_{S \sim \mathcal{D}^n} [L(\mathbf{A}(S)) - L_S(\mathbf{A}(S))] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{S \sim \mathcal{D}^n, (z'_1, \dots, z'_n) \sim \mathcal{D}^n} [\ell(\mathbf{A}(S); z'_i) - \ell(\mathbf{A}(S^{(i)}); z'_i)]$$

Hence we have the required result.  $\square$

For the next result, we will need the following two short utility lemmas.

**Utility Lemma 14.** For i.i.d.  $X_i$ ,  $|X_i| \leq B$  and  $X = \frac{1}{n} \sum_{i=1}^n X_i$  we have  $\mathbb{E} [|X - \mathbb{E}[X]|] \leq B/\sqrt{n}$ .

*Proof.*  $\mathbb{E} [|X - \mathbb{E}[X]|] \leq \sqrt{\mathbb{E} [|X - \mathbb{E}[X]|^2]} \leq \sqrt{\text{Var}[X]} = \sqrt{\text{Var}[X_i]/n} \leq B/\sqrt{n}$ .  $\square$

**Utility Lemma 15.** Let  $X, Y$  be random variables s.t.  $X \stackrel{a.s.}{\leq} Y$ . Then  $\mathbb{E} [|X|] \leq |\mathbb{E}[X]| + 2\mathbb{E} [|Y|]$ .

*Proof.*  $\mathbb{E} [|X|] = \mathbb{E} [(Y - X) - Y] \leq \mathbb{E}[Y - X] + \mathbb{E} [|Y|] \leq |\mathbb{E}[X]| + 2|\mathbb{E}[Y]|$ .  $\square$

**Lemma 16 (AERM + on-average generalization  $\Rightarrow$  generalization).** If  $\mathbf{A}$  is an AERM with rate  $\epsilon_{\text{erm}}(n)$  and on-average generalizes with rate  $\epsilon_{\text{oag}}(n)$  under  $\mathcal{D}$ , then  $\mathbf{A}$  generalizes with rate  $\epsilon_{\text{oag}}(n) + 2\epsilon_{\text{erm}}(n) + \frac{2B}{\sqrt{n}}$  under  $\mathcal{D}$ .

*Proof.* Recall that  $L^* = \inf_{\mathbf{h} \in \mathcal{H}} L(\mathbf{h})$ . For an arbitrarily small  $\nu > 0$ , let  $\mathbf{h}_\nu$  be a fixed hypothesis such that  $L(\mathbf{h}_\nu) \leq L^* + \nu$ . Using respective optimalities of  $\hat{\mathbf{h}}_S$  and  $L^*$  we can bound:

$$\begin{aligned} &L_S(\mathbf{A}(S)) - L(\mathbf{A}(S)) \\ &= L_S(\mathbf{A}(S)) - L_S(\hat{\mathbf{h}}_S) + L_S(\hat{\mathbf{h}}_S) - L_S(\mathbf{h}_\nu) + L_S(\mathbf{h}_\nu) - L(\mathbf{h}_\nu) + L(\mathbf{h}_\nu) - L(\mathbf{A}(S)) \\ &\leq L_S(\mathbf{A}(S)) - L_S(\hat{\mathbf{h}}_S) + L_S(\mathbf{h}_\nu) - L(\mathbf{h}_\nu) + \nu = Y_\nu \end{aligned}$$

Where the final equality defines a new random variable  $Y_\nu$ . By Lemma 14 and the AERM guarantee we have  $\mathbb{E}[|Y_\nu|] \leq \epsilon_{\text{erm}}(n) + B/\sqrt{n} + \nu$ . From Lemma 15 we can conclude that

$$\mathbb{E}[|L_S(\mathbf{A}(S)) - L(\mathbf{A}(S))|] \leq |\mathbb{E}[L_S(\mathbf{A}(S)) - L(\mathbf{A}(S))]| + 2\mathbb{E}[|Y_\nu|] \leq \epsilon_{\text{oag}}(n) + 2\epsilon_{\text{erm}}(n) + \frac{2B}{\sqrt{n}} + \nu.$$

Notice that the l.h.s. is a fixed quantity which does not depend on  $\nu$ . Therefore, we can take  $\nu$  in the r.h.s. to zero, and the result follows.  $\square$

Combining Lemma 13 and Lemma 16, we have now **established the stability $\leftrightarrow$ generalization parts of Theorem 6 and Theorem 7** (in fact, even a slightly stronger converse than in Theorem 7, as it does not require universality).

### A Sufficient Condition for Consistency

It is fairly straightforward to see that generalization (or even on-average generalization) of an AERM implies its consistency:

**Lemma 17 (AERM+generalization $\Rightarrow$ consistency).** *If  $\mathbf{A}$  is AERM with rate  $\epsilon_{\text{erm}}(n)$  and it on-average generalizes with rate  $\epsilon_{\text{oag}}(n)$  under  $\mathcal{D}$  then it is consistent with rate  $\epsilon_{\text{oag}}(n) + \epsilon_{\text{erm}}(n)$  under  $\mathcal{D}$ .*

*Proof.* For any  $\nu > 0$ , let  $\mathbf{h}_\nu$  be a hypothesis such that  $L(\mathbf{h}_\nu) \leq L^* + \nu$ . We have

$$\begin{aligned} \mathbb{E}[L(\mathbf{A}(S)) - L^*] &= \mathbb{E}[L(\mathbf{A}(S)) - L_S(\mathbf{h}_\nu) + \nu] \\ &= \mathbb{E}[L(\mathbf{A}(S)) - L_S(\mathbf{A}(S))] + \mathbb{E}[L_S(\mathbf{A}(S)) - L_S(\mathbf{h}_\nu)] + \nu \\ &\leq \mathbb{E}[L(\mathbf{A}(S)) - L_S(\mathbf{A}(S))] + \mathbb{E}[L_S(\mathbf{A}(S)) - L_S(\hat{\mathbf{h}}_S)] + \nu \\ &\leq \epsilon_{\text{oag}}(n) + \epsilon_{\text{erm}}(n) + \nu. \end{aligned}$$

Since this upper bound holds for any  $\nu$ , we can take  $\nu$  to zero, and the result follows.  $\square$

Combined with the results of Lemma 13, this completes the **proof of Theorem 6** and the **stability  $\rightarrow$  consistency and generalization  $\rightarrow$  consistency parts of Theorem 7**.

### Converse Direction

Lemma 13 already provides a converse result, establishing that stability is necessary for generalization. However, as it will turn out, in order to establish that stability is also necessary for *universal consistency*, we must prove that universal consistency of an AERM implies *universal generalization*. The assumption of *universal consistency* for the AERM is crucial here: mere consistency of an AERM with respect to a specific distribution does *not* imply generalization nor stability with respect to that distribution. The following example briefly illustrates this point.

**Example 4.** There exists a learning problem and a distribution on the instance space, such that the ERM (or any AERM) is consistent with rate  $\epsilon_{\text{cons}}(n) = 0$ , but does not generalize and is not average-RO stable (namely,  $\epsilon_{\text{gen}}(n), \epsilon_{\text{stable}}(n) = \Omega(1)$ ).

*Proof.* Let the instance space be  $[0, 1]$ , the hypothesis space consist of all finite subsets of  $[0, 1]$ , and define the objective function as  $\ell(\mathbf{h}, z) = \mathbf{1}_{\{z \notin \mathbf{h}\}}$ . Consider any continuous distribution on the instance space. Since the underlying distribution  $\mathcal{D}$  is continuous, we have  $L(\mathbf{h}) = 1$  for any hypothesis  $h$ . Therefore, any learning rule (including any AERM) will be consistent with  $L(\mathbf{A}(S)) = 1$ . On the other hand, the ERM here always achieves  $L_S(\hat{\mathbf{h}}_S) = 0$ , so any AERM cannot generalize, or even on-average-generalize (by Lemma 16), hence cannot be average-RO stable (by Lemma 13).  $\square$

The main tool we use to prove our desired converse result is the following lemma. It is here that we crucially use the universal consistency assumption (i.e., consistency with respect to *any* distribution). Intuitively, it states that if a problem is learnable at all, then although the ERM rule might fail, its empirical risk is a consistent estimator of the minimal achievable risk.

**Lemma 18 (Main Converse Lemma).** *If a problem is learnable, namely there exists a universally consistent rule  $\mathbf{A}$  with rate  $\epsilon_{\text{cons}}(n)$ , then under any distribution,*

$$\mathbb{E} \left[ \left| L_S(\hat{\mathbf{h}}_S) - L^* \right| \right] \leq \epsilon_{\text{emp}}(n) \quad \text{where} \quad (3.7.2)$$

$$\epsilon_{\text{emp}}(n) = 2\epsilon_{\text{cons}}(n') + \frac{2B}{\sqrt{n}} + \frac{2Bn'^2}{n}$$

for any  $n'$  such that  $2 \leq n' \leq n/2$ .

*Proof.* Let  $I = \{I_1, \dots, I_{n'}\}$  be a random sample of  $n'$  indexes in the range  $1..n$  where each  $I_i$  is independently uniformly distributed, and  $I$  is independent of  $S$ . Let  $S' = \{z_{I_i}\}_{i=1}^{n'}$ , i.e. a sample of size  $n'$  drawn from the uniform distribution over samples in  $S$  (with replacements). We first bound the probability that  $I$  has no repeated indexes (“duplicates”):

$$\Pr I \text{ has duplicates} \leq \frac{\sum_{i=1}^{n'} (i-1)}{n} \leq \frac{n'^2}{2n} \quad (3.7.3)$$

Conditioned on not having duplicates in  $I$ , the sample  $S'$  is actually distributed according to  $\mathcal{D}^{n'}$ , i.e. can be viewed as a sample from the original distribution. We therefore have by universal consistency:

$$\mathbb{E} [|L(\mathbf{A}(S')) - L^*| \mid \text{no dups}] \leq \epsilon_{\text{cons}}(n') \quad (3.7.4)$$

But viewed as a sample drawn from the uniform distribution over instances in  $S$ , we also have:

$$\mathbb{E}_{S'} \left[ \left| L_S(\mathbf{A}(S')) - L_S(\hat{\mathbf{h}}_S) \right| \right] \leq \epsilon_{\text{cons}}(n') \quad (3.7.5)$$

Conditioned on having no duplications in  $I$ , the set of those samples in  $S$  not chosen by  $I$  (i.e.  $S \setminus S'$ ) is independent of  $S'$ , and  $|S \setminus S'| = n - n'$ , and so by Lemma 14:

$$\mathbb{E}_S \left[ \left| L(\mathbf{A}(S')) - L_{S \setminus S'}(\mathbf{A}(S')) \right| \right] \leq \frac{B}{\sqrt{n - n'}} \quad (3.7.6)$$

Finally, if there are no duplicates, then for any hypothesis, and in particular for  $\mathbf{A}(S')$  we have:

$$\left| L_S(\mathbf{A}(S')) - L_{S \setminus S'}(\mathbf{A}(S')) \right| \leq \frac{2Bn'}{n} \quad (3.7.7)$$

Combining Eq. (3.7.4), Eq. (3.7.5), Eq. (3.7.6) and Eq. (3.7.7), accounting for a maximal discrepancy of  $B$  when we do have duplicates, and assuming  $2 \leq n' \leq n/2$ , we get the desired bound.  $\square$

Equipped with Lemma 18, we are now ready to show that universal consistency of an AERM implies universal generalization and that any universally consistent and generalizing rule must be an AERM. What we show is actually a bit stronger: that if a problem is learnable, and so Lemma 18 holds, then for any distribution  $\mathcal{D}$  separately, consistency of an AERM under  $\mathcal{D}$  implies generalization under  $\mathcal{D}$  and also any consistent and generalizing rule under  $\mathcal{D}$  must be an AERM.

**Lemma 19 (learnable+AERM+consistent $\Rightarrow$ generalizing).** *If Eq. (3.7.2) in Lemma 18 holds with rate  $\epsilon_{\text{emp}}(n)$ , and  $\mathbf{A}$  is an  $\epsilon_{\text{erm}}$ -AERM and  $\epsilon_{\text{cons}}$ -consistent under  $\mathcal{D}$ , then it is generalizing under  $\mathcal{D}$  with rate  $\epsilon_{\text{emp}}(n) + \epsilon_{\text{erm}}(n) + \epsilon_{\text{cons}}(n)$ .*

*Proof.*

$$\begin{aligned} \mathbb{E} \left[ \left| L_S(\mathbf{A}(S)) - L(\mathbf{A}(S)) \right| \right] &\leq \mathbb{E} \left[ \left| L_S(\mathbf{A}(S)) - L_S(\hat{\mathbf{h}}_S) \right| \right] + \mathbb{E} \left[ \left| L^* - L(\mathbf{A}(S)) \right| \right] + \mathbb{E} \left[ \left| L_S(\hat{\mathbf{h}}_S) - L^* \right| \right] \\ &\leq \epsilon_{\text{erm}}(n) + \epsilon_{\text{cons}}(n) + \epsilon_{\text{emp}}(n). \end{aligned}$$

$\square$

**Lemma 20 (learnable+consistent+generalizing $\Rightarrow$ AERM).** *If Eq. (3.7.2) in Lemma 18 holds with rate  $\epsilon_{\text{emp}}(n)$ , and  $\mathbf{A}$  is  $\epsilon_{\text{cons}}$ -consistent and  $\epsilon_{\text{gen}}$ -generalizing under  $\mathcal{D}$ , then it is AERM under  $\mathcal{D}$  with rate  $\epsilon_{\text{emp}}(n) + \epsilon_{\text{gen}}(n) + \epsilon_{\text{cons}}(n)$ .*

*Proof.*

$$\begin{aligned} \mathbb{E} \left[ \left| L_S(\mathbf{A}(S)) - L_S(\hat{\mathbf{h}}_S) \right| \right] &\leq \mathbb{E} \left[ \left| L_S(\mathbf{A}(S)) - L(\mathbf{A}(S)) \right| \right] + \mathbb{E} \left[ \left| L(\mathbf{A}(S)) - L^* \right| \right] + \mathbb{E} \left[ \left| L^* - L_S(\hat{\mathbf{h}}_S) \right| \right] \\ &\leq \epsilon_{\text{gen}}(n) + \epsilon_{\text{cons}}(n) + \epsilon_{\text{emp}}(n). \end{aligned}$$

$\square$

Lemma 19 establishes that universal consistency of an AERM implies universal generalization, and thus **completes the proof of Theorem 7**. Lemma 20 **establishes Theorem 8**. To get the rates in Subsection ??, we use  $n' = n^{1/4}$  in Lemma 18.

Lemma 17, Lemma 19 and Lemma 20 together establish an interesting relationship:

**Corollary 21.** *For a (universally) learnable problem, for any distribution  $\mathcal{D}$  and learning rule  $\mathbf{A}$ , any two of the following imply the third :*

- $\mathbf{A}$  is an AERM under  $\mathcal{D}$ .
- $\mathbf{A}$  is consistent under  $\mathcal{D}$ .
- $\mathbf{A}$  generalizes under  $\mathcal{D}$ .

Note, however, that any one property by itself is possible, even universally:

- In Subsection 3.3.1, we have discussed an example where the ERM learning rule is neither consistent nor generalizing, despite the problem being learnable.
- In the next subsection (Example 5) we demonstrate a universally consistent learning rule which is neither generalizing nor an AERM.
- A rule returning a fixed hypothesis always generalizes, but of course need not be consistent nor an AERM.

In contrast, for learnable supervised classification problems, it is not possible for a learning rule to be just universally consistent, without being an AERM and without generalization. Nor is it possible for a learning rule to be a universal AERM for a learnable problem, without being generalizing and consistent.

Corollary 21 can also provide a *certificate* of non-learnability. In other words, for the problem in Example 4 we show a specific distribution for which there is a consistent AERM that does not generalize. We can conclude that there is *no* universally consistent learning rule for the problem, otherwise the corollary is violated.

### **Existence of a Stable Rule**

Theorem 7 and Theorem 8, which we just completed proving, already establish that for AERMs, universal consistency is equivalent to universal average-RO stability. Existence of a universally average-RO stable AERM is thus sufficient for learnability. In order to prove that it is also necessary, it is enough to show that existence of a universally consistent learning rule implies existence of a universally consistent AERM. This AERM must then be average-RO stable by Theorem 7.

We actually show how to transform a consistent rule to a consistent and generalizing rule (Lemma 22 below). If this rule is universally consistent, then by Lemma 20 we can then conclude it must be an AERM, and by Lemma 13 it must be average-RO stable.



**Lemma 22.** For any rule  $\mathbf{A}$  there exists a rule  $\mathbf{A}'$ , such that:

- $\mathbf{A}'$  universally generalizes with rate  $\frac{3B}{\sqrt{n}}$ .
- For any  $\mathcal{D}$ , if  $\mathbf{A}$  is  $\epsilon_{\text{cons}}$ -consistent under  $\mathcal{D}$  then  $\mathbf{A}'$  is  $\epsilon_{\text{cons}}(\lfloor \sqrt{n} \rfloor)$  consistent under  $\mathcal{D}$ .
- $\mathbf{A}'$  is uniformly-RO-stable with rate  $\frac{2B}{\sqrt{n}}$ .

*Proof.* For a sample  $S$  of size  $n$ , let  $S'$  be a sub-sample consisting of some  $\lfloor \sqrt{n} \rfloor$  observation in  $S$ . To simplify the presentation, assume that  $\lfloor \sqrt{n} \rfloor$  is an integer. Define  $\mathbf{A}'(S) = \mathbf{A}(S')$ . That is,  $\mathbf{A}'$  applies  $A$  to only  $\sqrt{n}$  of the observation in  $S$ .

**$\mathbf{A}'$  generalizes:** We can decompose:

$$L_S(\mathbf{A}(S')) - L(\mathbf{A}(S')) = \frac{1}{\sqrt{n}}(L_{S'}(\mathbf{A}(S')) - L(\mathbf{A}(S'))) + (1 - \frac{1}{\sqrt{n}})(L_{S \setminus S'}(\mathbf{A}(S')) - L(\mathbf{A}(S')))$$

The first term can be bounded by  $2B/\sqrt{n}$ . As for the second term,  $S \setminus S'$  is statistically independent of  $S'$  and so we can use Lemma 14 to bound its expected magnitude to obtain:

$$\mathbb{E} [|L_S(\mathbf{A}(S')) - L(\mathbf{A}(S'))|] \leq \frac{2B}{\sqrt{n}} + (1 - \frac{1}{\sqrt{n}}) \frac{B}{\sqrt{n - \sqrt{n}}} \leq \frac{3B}{\sqrt{n}} \quad (3.7.8)$$

**$\mathbf{A}'$  is consistent:** If  $\mathbf{A}$  is consistent, then:

$$\mathbb{E} \left[ L(\mathbf{A}'(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L(\mathbf{h}) \right] = \mathbb{E} \left[ L(\mathbf{A}(S')) - \inf_{\mathbf{h} \in \mathcal{H}} L(\mathbf{h}) \right] \leq \epsilon_{\text{cons}}(\sqrt{n})$$

**$\mathbf{A}'$  is uniformly-RO-stable:** Since  $\mathbf{A}'$  only uses the first  $\sqrt{n}$  samples of  $S$ , for any  $i > \sqrt{n}$  we have  $\mathbf{A}'(S^{(i)}) = \mathbf{A}'(S)$  and so:

$$\frac{1}{n} \sum_{i=1}^n \left| \ell(\mathbf{A}'(S^{(i)}); z') - \ell(\mathbf{A}'(S); z') \right| = \frac{1}{n} \sum_{i=1}^{\sqrt{n}} \left| \ell(\mathbf{A}'(S^{(i)}); z') - \ell(\mathbf{A}'(S); z') \right| \leq \frac{2B}{\sqrt{n}}$$

□

**Proof of Converse in Theorem 5** If there exists a universally consistent rule with rate  $\epsilon_{\text{cons}}(n)$ , by Lemma 22 there exists  $\mathbf{A}'$  which is  $\epsilon_{\text{cons}}(\sqrt{n})$ - universally consistent,  $\frac{2B}{\sqrt{n}}$ -generalizing and  $\frac{2B}{\sqrt{n}}$ -uniformly-RO-stable. Further by Lemma 20 and Lemma 18 (with  $n' = n^{1/4}$ ), we can conclude that  $A'$  is  $\epsilon_{\text{erm}}$ -universally AERM where,

$$\epsilon_{\text{erm}}(n) \leq 3\epsilon_{\text{cons}}(n^{1/4}) + \frac{8B}{\sqrt{n}}.$$

Hence we get the specified rate for the converse direction. To see that if there exists a rule that is a universal AERM and stable it is consistent, we simply use Lemma 17.

As a final note, the following example shows that while learnability is equivalent to the existence of stable and consistent AERM's (Theorem 5 and Theorem 7), there might still exist other learning rules, which are neither stable, nor generalize, nor AERM's. In this sense, our results characterize learnability, but do not characterize all learning rules which “work”.

**Example 5.** There exists a learning problem with a universally consistent learning rule, which is not average-RO stable, generalizing nor an AERM.

*Proof.* Let the instance space be  $[0, 1]$ . Let the hypothesis space consist of all finite subsets of  $[0, 1]$ , and the objective function be the indicator function  $\ell(\mathbf{h}, z) = \mathbf{1}_{\{z \in \mathbf{h}\}}$ . Consider the following learning rule: given a sample  $S \subseteq [0, 1]$ , the learning rule checks if there are any two identical instances in the sample. If so, the learning rule returns the empty set  $\emptyset$ . Otherwise, it returns the sample.

Consider any continuous distribution on  $[0, 1]$ . In that case, the probability of having two identical instances is 0. Therefore, the learning rule always returns a countable non-empty set  $\mathbf{A}(S)$ , with  $L_S(\mathbf{A}(S)) = 1$ , while  $L_S(\emptyset) = 0$  (so it is not an AERM) and  $L(\mathbf{A}(S)) = 0$  (so it does not generalize). Also,  $\ell(\mathbf{A}(S), z_i) = 1$  while  $\ell(\mathbf{A}(S^{(i)}), z_i) = 0$  with probability 1, so it is not average-RO stable either.

However, the learning rule is universally consistent. If the underlying distribution is continuous on  $[0, 1]$ , then the returned hypothesis is  $S$ , which is countable hence  $L(S) = 0 = \inf_{\mathbf{h}} L(\mathbf{h})$ . For discrete distributions, let  $M_1$  denote the proportion of instances in the sample which appear exactly once, and let  $M_0$  be the probability mass of instances which did not appear in the sample. Using [32, Theorem 3], we have that for any  $\delta$ , it holds with probability at least  $1 - \delta$  over a sample of size  $n$  that

$$|M_0 - M_1| \leq O\left(\frac{\log(n/\delta)}{\sqrt{n}}\right),$$

uniformly for any discrete distribution. If this occurs, then either  $M_1 < 1$ , or  $M_0 \geq 1 - O(\log(n/\delta)/\sqrt{n})$ . But in the first event, we get duplicate instances in the sample, so the returned hypothesis is the optimal  $\emptyset$ , and in the second case, the returned hypothesis is the sample, which has a total probability mass of at most  $O(\log(n/\delta)/\sqrt{n})$ , and therefore  $L(\mathbf{A}(S)) \leq O(\log(n/\delta)/\sqrt{n})$ . As a result, regardless of the underlying distribution, with probability of at least  $1 - \delta$  over the sample,

$$L(\mathbf{A}(S)) \leq O\left(\frac{\log(n/\delta)}{\sqrt{n}}\right).$$

Since the r.h.s. converges to 0 with  $n$  for any  $\delta$ , it is easy to see that the learning rule is universally consistent.  $\square$

### 3.7.2 Other Proofs

*Proof of Theorem 1.* To prove the theorem, we use a stability argument. Denote

$$\hat{F}^{(i)}(\mathbf{h}) = \frac{1}{n} \left( \ell(\mathbf{h}, z'_i) + \sum_{j \neq i} \ell(\mathbf{h}, z_j) \right).$$

the empirical average with  $z_i$  replaced by an independently and identically drawn  $z'_i$ , and consider its minimizer:

$$\hat{\mathbf{h}}_S^{(i)} = \arg \min_{\mathbf{h} \in \mathcal{H}} \hat{F}^{(i)}(\mathbf{h}).$$

We first use strong convexity and Lipschitz-continuity to establish that empirical minimization is stable in the following sense:

$$\forall z \in \mathcal{Z}, \quad \left| \ell(\hat{\mathbf{h}}_S, z) - \ell(\hat{\mathbf{h}}_S^{(i)}, z) \right| \leq \frac{4L^2}{\lambda n}. \quad (3.7.9)$$

We have that

$$\begin{aligned} & \hat{L}(\hat{\mathbf{h}}_S^{(i)}) - \hat{L}(\hat{\mathbf{h}}_S) \\ &= \frac{\ell(\hat{\mathbf{h}}_S^{(i)}, z_i) - \ell(\hat{\mathbf{h}}_S, z_i)}{n} + \frac{\sum_{j \neq i} (\ell(\hat{\mathbf{h}}_S^{(i)}, z_j) - \ell(\hat{\mathbf{h}}_S, z_j))}{n} \\ &= \frac{\ell(\hat{\mathbf{h}}_S^{(i)}, z_i) - \ell(\hat{\mathbf{h}}_S, z_i)}{n} + \frac{\ell(\hat{\mathbf{h}}_S, z'_i) - \ell(\hat{\mathbf{h}}_S^{(i)}, z'_i)}{n} \\ &\quad + \left( \hat{F}^{(i)}(\hat{\mathbf{h}}_S^{(i)}) - \hat{F}^{(i)}(\hat{\mathbf{h}}_S) \right) \\ &\leq \frac{|\ell(\hat{\mathbf{h}}_S^{(i)}, z_i) - \ell(\hat{\mathbf{h}}_S, z_i)|}{n} + \frac{|\ell(\hat{\mathbf{h}}_S, z'_i) - \ell(\hat{\mathbf{h}}_S^{(i)}, z'_i)|}{n} \\ &\leq \frac{2L}{n} \left\| \hat{\mathbf{h}}_S^{(i)} - \hat{\mathbf{h}}_S \right\| \end{aligned} \quad (3.7.10)$$

where the first inequality follows from the fact that  $\hat{\mathbf{h}}_S^{(i)}$  is the minimizer of  $\hat{F}^{(i)}(\mathbf{h})$  and for the second inequality we use Lipschitz continuity. But from strong convexity of  $\hat{L}(\mathbf{h})$  and the fact that  $\hat{\mathbf{h}}_S$  minimizes  $\hat{L}(\mathbf{h})$  we also have that

$$\hat{L}(\hat{\mathbf{h}}_S^{(i)}) \geq \hat{L}(\hat{\mathbf{h}}_S) + \frac{\lambda}{2} \left\| \hat{\mathbf{h}}_S^{(i)} - \hat{\mathbf{h}}_S \right\|^2. \quad (3.7.11)$$

Combining Eq. (3.7.11) with Eq. (3.7.10) we get  $\left\| \hat{\mathbf{h}}_S^{(i)} - \hat{\mathbf{h}}_S \right\| \leq 4L/(\lambda n)$  and combining this with Lipschitz continuity of  $f$  we obtain that Eq. (3.7.9) holds. Later in this chapter we show that a stable ERM is sufficient for learnability. More formally, Eq. (3.7.9) implies that the ERM is uniform-RO stability (Definition 14) with rate  $\epsilon_{\text{stable}}(n) = 4L^2/(\lambda n)$  and therefore Theorem 6 implies that the ERM is consistent with rate  $\leq \epsilon_{\text{stable}}(n)$ , namely

$$\mathbb{E}_{S \sim \mathcal{D}^n} \left[ L(\hat{\mathbf{h}}_S) - L^* \right] \leq \frac{4L^2}{\lambda n}.$$

This concludes the proof. □

We now turn to the proof of Theorem 2.

*Proof of Theorem 2.* Let  $r(\mathbf{h}; z) = \frac{\lambda}{2}\|\mathbf{h}\|^2 + \ell(\mathbf{h}; z)$  and let  $R(\mathbf{h}) = \mathbb{E}_z[r(\mathbf{h}, z)]$ . Note that  $\hat{\mathbf{h}}_\lambda$  is the empirical minimizer for the stochastic optimization problem defined by  $r(\mathbf{h}; z)$ .

We apply Theorem 1 to  $r(\mathbf{h}; z)$ , to this end note that since  $f$  is  $L$ -Lipschitz and  $\forall \mathbf{h} \in \mathcal{H}$ ,  $\|\mathbf{h}\| \leq B$  we see that  $r$  is in fact  $L + \lambda B$ -Lipschitz. Applying Theorem 1, we see that

$$\frac{\lambda}{2} \left\| \hat{\mathbf{h}}_\lambda \right\|^2 + \mathbb{E}_S \left[ L(\hat{\mathbf{h}}_\lambda) \right] = \mathbb{E}_S \left[ R(\hat{\mathbf{h}}_\lambda) \right] \leq \inf_{\mathbf{h}} R(\mathbf{h}) + \frac{4(L + \lambda B)^2}{\lambda n}$$

Now note that  $\inf_{\mathbf{h}} R(\mathbf{h}) \leq \inf_{\mathbf{h} \in \mathcal{H}} L(\mathbf{h}) + \frac{\lambda}{2} B^2 = L^* + \frac{\lambda}{2} B^2$ , and so we get that

$$\begin{aligned} \mathbb{E}_S \left[ L(\hat{\mathbf{h}}_\lambda) \right] &\leq \inf_{\mathbf{h} \in \mathcal{H}} L(\mathbf{h}) + \frac{\lambda}{2} B^2 + \frac{4(L + \lambda B)^2}{\lambda n} \\ &\leq \inf_{\mathbf{h} \in \mathcal{H}} L(\mathbf{h}) + \frac{\lambda}{2} B^2 + \frac{8L^2}{\lambda n} + \frac{8\lambda B^2}{n} \end{aligned}$$

Plugging in the value of  $\lambda$  given in the theorem statement we see that

$$L(\hat{\mathbf{h}}_\lambda) \leq \inf_{\mathbf{h} \in \mathcal{H}} L(\mathbf{h}) + 4\sqrt{\frac{L^2 B^2}{n}} + \frac{32}{n}\sqrt{\frac{L^2 B^2}{n}}$$

This gives us the required bound. □

*Proof of Theorem 4.* Lemma 17 and Lemma 19 from subsection 3.7.1 already tell us that for ERM's, universal consistency is equivalent to universal generalization. Moreover, Lemma 16 implies that for ERM's, generalization is equivalent to on-average generalization (see Eq. (3.7.1) for the exact definition). Thus, is left to prove that for ERM's, generalization implies LOO stability, and LOO stability implies on-average generalization. stability.

First, suppose the ERM learning rule is generalizing with rate  $\epsilon_{\text{gen}}(n)$ . Note that  $\ell(\hat{\mathbf{h}}_{S \setminus i}; z_i) - \ell(\hat{\mathbf{h}}_S; z_i)$  is

always nonnegative. Therefore the LOO stability of the ERM can be upper bounded as follows:

$$\begin{aligned}
& \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[ |\ell(\hat{\mathbf{h}}_{S \setminus i}; z_i) - \ell(\hat{\mathbf{h}}_S; z_i)| \right] \\
&= \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[ \ell(\hat{\mathbf{h}}_{S \setminus i}; z_i) - \ell(\hat{\mathbf{h}}_S; z_i) \right] \\
&= \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[ L(\hat{\mathbf{h}}_{S \setminus i}) \right] - \mathbb{E} \left[ \frac{1}{n} \sum_{i=1}^n \ell(\hat{\mathbf{h}}_S; z_i) \right] \\
&\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[ L_{S \setminus i}(\hat{\mathbf{h}}_{S \setminus i}) + \epsilon_{\text{gen}}(n-1) \right] - \mathbb{E} \left[ L_S(\hat{\mathbf{h}}_S) \right] \\
&= \epsilon_{\text{gen}}(n-1) + \mathbb{E} \left[ \frac{1}{n} \sum_{i=1}^n L_{S \setminus i}(\hat{\mathbf{h}}_{S \setminus i}) - L_S(\hat{\mathbf{h}}_S) \right] \\
&\leq \epsilon_{\text{gen}}(n-1).
\end{aligned}$$

For the opposite direction, suppose the ERM learning rule is LOO stable with rate  $\epsilon_{\text{stable}}(n)$ . Notice that we can get any sample of size  $n-1$  by picking a sample  $S$  of size  $n$  and discarding any instance  $i$ . Therefore, the on-average generalization rate of the ERM for samples of size  $n-1$  is equal to the following:

$$\begin{aligned}
& \left| \mathbb{E} \left[ L(\hat{\mathbf{h}}_{S \setminus i}) - L_{S \setminus i}(\hat{\mathbf{h}}_{S \setminus i}) \right] \right| \\
&= \left| \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[ L(\hat{\mathbf{h}}_{S \setminus i}) - L_{S \setminus i}(\hat{\mathbf{h}}_{S \setminus i}) \right] \right| \\
&= \left| \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[ \ell(\hat{\mathbf{h}}_{S \setminus i}; z_i) \right] - \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[ L_{S \setminus i}(\hat{\mathbf{h}}_{S \setminus i}) \right] \right|
\end{aligned}$$

Now, note that for the ERM's of  $S$  and  $S \setminus i$  we have  $\left| L_{S \setminus i}(\hat{\mathbf{h}}_{S \setminus i}) - L_S(\hat{\mathbf{h}}_S) \right| \leq \frac{2B}{n}$ . Therefore, we can upper bound the above by

$$\begin{aligned}
& \left| \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[ \ell(\hat{\mathbf{h}}_{S \setminus i}; z_i) \right] - \mathbb{E} \left[ L_S(\hat{\mathbf{h}}_S) \right] \right| + \frac{2B}{n} \\
&= \left| \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[ \ell(\hat{\mathbf{h}}_{S \setminus i}; z_i) - \ell(\hat{\mathbf{h}}_S; z_i) \right] \right| \\
&\leq \epsilon_{\text{stable}}(n)
\end{aligned}$$

using the assumption that the learning rule is  $\epsilon_{\text{stable}}(n)$ -stable.  $\square$

*Proof of Lemma 11.* Consider the learning rule  $A'$  which given a sample  $S$ , returns a uniform distribution

over  $A(S')$ , where  $S'$  ranges over all subsets of  $S$  of size  $\lfloor \sqrt{n} \rfloor$ .

The fact that  $\mathbf{A}'$  is symmetric is trivial. We now prove the other assertions in the lemma.

**$\mathbf{A}'$  is consistent:** First note that  $L(\mathbf{A}'(S)) = \mathbb{E}_{S'} [L(\mathbf{A}(S'))]$ , and so:

$$\mathbb{E}_S [|L(\mathbf{A}'(S)) - L^*|] \leq \mathbb{E}_{S,S'} [|L(\mathbf{A}(S')) - L^*|] = \mathbb{E}_{[S']} [\mathbb{E}_{S|[S']} [|L(\mathbf{A}(S')) - L^*|]]$$

where  $[S']$  designates a choice of indices for  $S'$ . This decomposition of the random choice of  $S'$  (e.g., first deciding on the indices and only then sampling  $S$ ) allows us think of  $[S']$  and  $S$  as statistically independent. Given a fixed choice of indices  $[S']$ ,  $S'$  is simply an i.i.d. sample of size  $\lfloor \sqrt{n} \rfloor$ . Therefore, if  $\mathbf{A}$  is consistent,  $|L(\mathbf{A}(S')) - L^*| \leq \epsilon_{\text{cons}}(\lfloor \sqrt{n} \rfloor)$ , this holds for any possible fixed  $[S']$ , and therefore

$$\mathbb{E}_{[S']} [\mathbb{E}_{S|[S']} [|L(\mathbf{A}(S')) - L^*|]] = \mathbb{E}_{[S']} [\epsilon_{\text{cons}}(\lfloor \sqrt{n} \rfloor)] \leq \epsilon_{\text{cons}}(\lfloor \sqrt{n} \rfloor).$$

**$\mathbf{A}'$  generalizes:** For convenience, let  $b(S, S') = |L_S(\mathbf{A}(S')) - L(\mathbf{A}(S'))|$ . Using similar arguments and notation as above:

$$\begin{aligned} & \mathbb{E}_S [|L_S(\mathbf{A}'(S)) - L(\mathbf{A}'(S))|] \\ & \leq \mathbb{E}_{[S']} [\mathbb{E}_{S|[S']} [b(S, S')]] \\ & \leq \mathbb{E}_{[S']} \left[ \mathbb{E}_{S|[S']} \left[ \frac{\lfloor \sqrt{n} \rfloor}{n} b(S', S') \right] + \mathbb{E}_{S|[S']} \left[ \left( 1 - \frac{\lfloor \sqrt{n} \rfloor}{n} \right) b(S \setminus S', S') \right] \right] \\ & \leq \mathbb{E}_{[S']} \left[ \frac{\lfloor \sqrt{n} \rfloor}{n} 2B + \left( 1 - \frac{\lfloor \sqrt{n} \rfloor}{n} \right) \frac{B}{\sqrt{n} - \lfloor \sqrt{n} \rfloor + 1} \right], \end{aligned}$$

where the last line follows from Lemma 14 and the fact that  $b(S, S') \leq 2B$  for any  $S, S'$ . It is not hard to show that the expression above is at most  $4B/\sqrt{n}$ , assuming  $n \geq 1$ .

**$\mathbf{A}'$  is strongly-uniform-RO stable:** For any sample  $S$ , any  $i$  and replacement instance  $z_i$ , and any instance  $z'$ , we have that

$$\left| \ell(\mathbf{A}'(S^{(i)}); z') - \ell(\mathbf{A}'(S); z') \right| \leq \mathbb{E}_{S'} \left[ \left| \ell(\mathbf{A}(S'^{(i)}); z') - \ell(\mathbf{A}(S'); z') \right| \right],$$

where we take  $S'^{(i)}$  in the expectation to mean  $S'$  if  $i \notin [S']$ . Notice that if  $i \notin [S']$ , then  $\ell(\mathbf{A}(S'^{(i)}); z_i) - \ell(\mathbf{A}(S'); z_i)$  is trivially 0. Thus, we can upper bound the expression above by

$$\mathbb{E}_{S'} \left[ \left| \ell(\mathbf{A}(S'^i); z') - \ell(\mathbf{A}(S'); z') \right| \mid i \in [S'] \right].$$

Since  $S'$  is chosen uniformly over all  $\lfloor \sqrt{n} \rfloor$ -subsets of  $S$ , all permutations of  $[S']$  are equally happen to occur, and therefore the above is equal to

$$\mathbb{E}_{S'} \left[ \frac{1}{\lfloor \sqrt{n} \rfloor} \sum_{j \in S'} \left| \ell(\mathbf{A}(S'^{(j)}); z') - \ell(\mathbf{A}(S'); z') \right| \right] \leq \mathbb{E}_{S'} [\epsilon_{\text{stable}}(\lfloor \sqrt{n} \rfloor)] = \epsilon_{\text{stable}}(\lfloor \sqrt{n} \rfloor).$$

**A' is an always AERM:** For any fixed sample  $S$ , we note that

$$\begin{aligned} |L_S(\mathbf{A}'(S)) - L_S(\hat{\mathbf{h}}_S)| &= \mathbb{E}_{S'} [L_S(\mathbf{A}(S')) - L_S(\hat{\mathbf{h}}_S)] \\ &= \mathbb{E}_{S' \sim \mathcal{U}(S)^{\lfloor \sqrt{n} \rfloor}} [L_S(\mathbf{A}(S')) - L_S(\hat{\mathbf{h}}_S) \mid \text{no dups}], \end{aligned}$$

where  $\mathcal{U}(S)^{\lfloor \sqrt{n} \rfloor}$  signifies the distribution of i.i.d. samples of size  $\lfloor \sqrt{n} \rfloor$ , picked uniformly at random (with replacement) from  $\lfloor \sqrt{n} \rfloor$ , and 'no dups' signifies the event that no element in  $S$  was picked twice. By the law of total expectation, this is at most

$$\frac{\mathbb{E}_{S' \sim \mathcal{U}(S)^{\lfloor \sqrt{n} \rfloor}} [L_S(\mathbf{A}(S')) - L_S(\hat{\mathbf{h}}_S)]}{\Pr \text{ no dups}}.$$

Since the learning rule  $\mathbf{A}$  is universally consistent, it is in particular consistent with respect to the distribution  $\mathcal{U}(S)$ , and therefore the expectation in the expression above is at most  $\epsilon_{\text{cons}}(\lfloor \sqrt{n} \rfloor)$ . As to  $\Pr \text{ no dups}$ , an analysis identical to the one performed in the proof of Lemma 18 (see Eq. (3.7.3)) implies that it is at least  $1 - (\lfloor \sqrt{n} \rfloor)^2/n \geq 1/2$ . Overall, we get that  $L_S(\mathbf{A}'(S)) - L_S(\hat{\mathbf{h}}_S) \leq 2\epsilon_{\text{cons}}(\lfloor \sqrt{n} \rfloor)$ , so in particular

$$\frac{\mathbb{E}_{S' \sim \mathcal{U}(S)^{\lfloor \sqrt{n} \rfloor}} [L_S(\mathbf{A}(S')) - L_S(\hat{\mathbf{h}}_S)]}{\Pr \text{ no dups}} \leq 2\epsilon_{\text{cons}}(\lfloor \sqrt{n} \rfloor),$$

from which the claim follows. □

*Proof of Theorem 12.* By Lemma 11, if a learning problem is learnable, there exists a (possibly randomized) symmetric learning rule  $\mathbf{A}'$ , which is an always AERM and strongly-uniform-RO stable. More specifically, we have that

$$\sup_{S \in \mathcal{Z}^n} \left( L_S(\mathbf{A}'(S)) - L_S(\hat{\mathbf{h}}_S) \right) \leq 2\epsilon_{\text{cons}}(\lfloor \sqrt{n} \rfloor),$$

as well as

$$\sup_{S \in \mathcal{Z}^n, z'} \left| \ell(\mathbf{A}'(S); z') - \ell(\mathbf{A}'(S^{(i)}); z') \right| \leq \frac{4B}{\sqrt{n}}.$$

In particular, there exists some symmetric  $\mathbf{A} : \mathcal{Z}^n \rightarrow \bar{\mathcal{H}}$ , for which the expression in Eq. (3.6.1) is at most

$$2\epsilon_{\text{cons}}(\lfloor \sqrt{n} \rfloor) + \frac{4B}{\sqrt{n}}.$$

Therefore, by definition, the  $\mathbf{A}$  found satisfies

$$\sup_{S \in \mathcal{Z}^n} \left( L_S(\mathbf{A}_n(S)) - L_S(\hat{\mathbf{h}}_S) \right) \leq 2\epsilon_{\text{cons}}(\lfloor \sqrt{n} \rfloor) + \frac{4B}{\sqrt{n}}, \quad (3.7.12)$$

as well as

$$\sup_{S \in \mathcal{Z}^n} \left| \ell(\mathbf{A}_n(S); z') - \ell(\mathbf{A}_n(S^{(i)}); z') \right| \leq 2\epsilon_{\text{cons}}(\lfloor \sqrt{n} \rfloor) + \frac{4B}{\sqrt{n}}. \quad (3.7.13)$$

In Theorem 7, we have seen that a universally average-RO stable AERM learning rule has to be universally consistent. The inequalities above essentially say that  $\mathbf{A}$  is in fact both strongly-uniform-RO stable (and in particular, universally average-RO stable) and an AERM, and thus is a universally consistent learning rule. Formally speaking, this is not entirely accurate, because  $\mathbf{A}$  is defined only with respect to samples of size  $n$ , and hence is not formally a learning rule which can be applied to samples of any size. However, the analysis we have done earlier in fact carries through also for learning rules  $\mathbf{A}$  which are defined just on a specific sample size  $n$ . In particular, the analysis of Lemma 13 and Lemma 17 hold verbatim for  $\mathbf{A}$  (with trivial modifications due to the fact that  $\mathbf{A}$  is randomized), and together imply that since Eq. (3.7.12) and Eq. (3.7.13) hold, then

$$\mathbb{E} [L(\mathbf{A}(S)) - L^*] \leq 4\epsilon_{\text{cons}}(\lfloor \sqrt{n} \rfloor) + \frac{8B}{\sqrt{n}}.$$

Therefore, our learning algorithm is consistent with rate  $4\epsilon_{\text{cons}}(\lfloor \sqrt{n} \rfloor) + \frac{8B}{\sqrt{n}}$ .  $\square$

## 3.8 Discussion

In this chapter we began exploring the issue of statistical learnability in the General Setting, and uncovered important relationships between learnability and stability. However problems are left open and avenues left to explore, some of which are listed below.

First, a natural question that might arise is whether it is possible to come up with well-known machine learning applications, where learnability is achievable despite uniform convergence failing to hold. Subsequently in a very recent work [33] it was shown that for multi-class learning problems with large number of classes, problems could still be learnable while uniform convergence fails and ERM approach may not be successful (at least not all ERM's are good).

In Section 3.6.2, we have managed to obtain a completely generic learning algorithm: an algorithm which in principle allows us to learn any learnable problem. However, the algorithm suffers from the severe drawback that in general, it requires unbounded computational power and is not in any succinct form. Can we derive an algorithm in a simple form, or characterize classes of learning problems where our algorithm, or some other generic learning algorithm utilizing the notion of stability, can be written for instance, as a regularized ERM learning rule?

On a related vein, it would be interesting to develop learning algorithms (perhaps for specific settings rather than generic learning problems) which directly use stability in order to learn. Convex regularization is one



such mechanism, as discussed earlier to induce stability. Are there other mechanisms, which use the notion of stability in a different way?

Another issue is that even the existence of uniform-RO stable AERM (or strongly-uniform-RO stable, always-AERM allowing for convexity/randomization) is not as elegant and simple as having finite VC dimension or fat-shattering dimension. It would be very interesting to derive equivalent but more “combinatorial” conditions for learnability.

## Chapter 4

# Online Learning/Optimization

In the online learning framework, the learner is faced with a sequence of data appearing at discrete time intervals. In contrast to statistical learning scenario where the learner is being evaluated after the sequence is completely revealed, in the online framework the learner is evaluated at every round. Furthermore, in the statistical learning scenario the data source is typically assumed to be drawn *i.i.d.* while in the online framework we relax or eliminate any stochastic assumptions on the data source. As such, the online learning problem can be phrased as a repeated two-player game between the learner (player) and the adversary (Nature).

Unlike the statistical learning framework, there has been surprisingly little work on characterizing learnability and developing generic tools to obtain rates for the online learning framework. Littlestone [34] has shown that, in the setting of prediction of binary outcomes, a certain combinatorial property of the binary-valued function class characterizes learnability in the realizable case (that is, when the outcomes presented by the adversary are given according to some function in the class  $\mathcal{F}$ ). The result has been extended to the non-realizable case by Shai Ben-David, Dávid Pál and Shai Shalev-Shwartz [35] who named this combinatorial quantity the *Littlestone's dimension*. Coincident with [35], *minimax* analysis of online convex optimization yielded new insights into the value of the game, its minimax dual representation, as well as algorithm-independent upper and lower bounds [36, 7]. In this chapter we will build tools analogous to the ones we have to analyze statistical learning problems like Rademacher complexity, covering numbers etc that work for online learning framework.

Section 4.1 introduces the problem at hand formally and provides various key definitions like that of online learning algorithm. Section 4.2 formally defines value of the online learning game and uses it to define learnability of an online learning problem and provides the main minimax theorem which is key in getting many results. In Section 4.3 we formally define the sequential Rademacher complexity and shows that it can be used to bound the value and hence get bounds on optimal rates for online learning problems. Structural properties of this complexity measure is also provided. This is perhaps the most important tool we introduce in this chapter. Section 4.4 introduces sequential covering numbers and sequential fat-shattering dimension and shows relation between them and how they can be used to bound the sequential Rademacher complexity.

Section 4.5 shows how these complexity tools provide a martingale uniform convergence story for online learning. Section 4.6 shows how these complexity measures can be used to characterize online learnability of supervised learning problems and goes on to provide a generic algorithm for online supervised learning problems. Section 4.7 provides various examples illustrating how the results can be used to obtain bounds for various online learning problems.

## 4.1 The Online Learning Problem

The online learning problem is a continual learning process that proceeds in rounds where in each round adversary picks an instance, learner in turn picks a hypothesis. At the end of the round, the learner pays loss for picking the particular hypothesis against the instance chosen by the adversary for that round. Specifically the online learning protocol can be written as :

**Online Learning Protocol :**

**for**  $t = 1$  **to**  $n$

Learner picks hypothesis  $\mathbf{h}_t \in \tilde{\mathcal{H}}$

Adversary simultaneously picks instance  $z_t \in \mathcal{Z}$

Learner pays loss  $\ell(\mathbf{h}_t, z_t)$

**end for**

Notice that unlike the statistical learning framework, the instances need not be selected statistically according to some fixed distribution. The adversary at round  $t$  can select the instance  $z_t$  in an adversarial worst case fashion based on previous instances  $z_1, \dots, z_{t-1}$  and based on previous hypotheses  $\mathbf{h}_1, \dots, \mathbf{h}_{t-1}$  selected by the learner. The automatic question that will rise in the reader's mind would be what is the goal of the learner in this online learning framework. The goal we consider for this online learning framework at the end of  $n$  rounds is to have low "regret" w.r.t. the best single hypothesis from target class  $\mathcal{H}$  that the learner could have picked at hind-sight after knowing  $z_1, \dots, z_n$ . That is, the regret after  $n$  rounds is defined as :

$$\mathbf{R}_n((\mathbf{h}_1, z_1), \dots, (\mathbf{h}_n, z_n)) = \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \quad (4.1.1)$$

**On the Importance of Randomization :** Unlike the statistical learning framework, due to the adversarial nature of the online learning framework, randomization of learning rules become necessary even for very simple (non-convex) problems. To illustrate this let us consider the simple problem of binary classification w.r.t. target hypothesis class  $\mathcal{H}$  consisting of exactly two functions, the constant mapping to 1 and the constant mapping to  $-1$ . Now if we consider only deterministic rule, the adversary (knowing the learning rule of the learner), at each round can pick label  $y_t$  to be opposite of what the learner predicts. At the end of  $n$  rounds, the average loss suffered by the learner is  $\frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, (x_t, y_t)) = 1$ . However  $\inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, (x_t, y_t)) \leq \frac{1}{2}$  Hence regret is lower bounded by  $1/2$ . This example can easily be extended to other non-convex (in  $\mathcal{H}$ ) supervised learning problems. Thus we see that in the online learning

framework considering randomized learning rules is imperative to get useful results. However whenever  $\bar{\mathcal{H}}$  is a vector space and loss  $\ell$  is convex in its first argument, simple application of Jensen's inequality shows that it is enough to consider only deterministic learning rules.

Just like we introduced the notion of "Statistical Learning Rule" to refer to the way the learner picks hypothesis for statistical learning problems, we now formally define online learning rules to refer to learner's strategy for an online learning problem. However as discussed in the previous paragraph, we shall right from the start define learning rule to be a randomized one.

**Definition 21.** A "Randomized Online Learning Rule"  $\mathbf{A} : \bigcup_{n \in \mathbb{N} \cup \{\}} (\bar{\mathcal{H}}^n \times \mathcal{Z}^n) \mapsto \Delta(\bar{\mathcal{H}})$  is a mapping from sequences of hypothesis, instance pair in  $\bar{\mathcal{H}} \times \mathcal{Z}$  to the set of all Borel distributions over hypothesis set  $\bar{\mathcal{H}}$ .

Given a "Randomized Online Learning Rule"  $\mathbf{A}$ , the way the learner uses this for picking the hypothesis for each round is as follows. On any round 1, learner simply sampling  $\mathbf{h}_1 \sim \mathbf{A}(\{\})$  which is some fixed distribution over the set  $\bar{\mathcal{H}}$ . Further recursively at each round  $t$ , the learner uses the learning rule on instances  $z_1, \dots, z_{t-1}$  seen so far and hypotheses  $\mathbf{h}_1, \dots, \mathbf{h}_{t-1}$  sampled so far to pick the hypothesis for current round as  $\mathbf{h}_t \sim \mathbf{A}(\mathbf{h}_{1:t-1}, z_{1:t-1})$ . For a randomized learning algorithm our aim will be to ensure that under expectation over randomization, the regret of the learner is small.

We shall often refer to the the "Randomized Online Learning Rule" as player/learner's strategy as well. Further, whenever the output of the learning rule is deterministic, ie. whenever for each input sequence, the learning rule picks a particular hypothesis with probability one, we will refer to the rule as a "Deterministic Online Learning Rule". Notice that if the learning rule is deterministic then at any round  $t$ , giving hypotheses  $\mathbf{h}_1, \dots, \mathbf{h}_{t-1}$  as argument to the learning is redundant as they can be calculated using just the instances. Hence we see that when we talk of "Deterministic Learning Rule", it is no different in form from "Statistical Learning Rules" which are mapping from sequence of instances to hypothesis set  $\bar{\mathcal{H}}$ .

## 4.2 Online Learnability and the Value of the Game

How can one define learnability in the online learning framework? Of course, we will refer to a problem as online learnable if there exists a randomized online learning algorithm which can guarantee diminishing expected regret against any strategy of the adversary. We can use the concept of value of a game to succinctly write down that it means for a problem to be online learnable. We will assume that  $\Delta(\bar{\mathcal{H}})$ , the set of all Borel probability measures on  $\bar{\mathcal{H}}$  is weakly compact. Note that if  $\bar{\mathcal{H}}$  is a compact set or is the unit ball of a reflexive Banach space then one can guarantee that  $\Delta(\bar{\mathcal{H}})$  is automatically weakly compact. Hence this restriction is automatically true in most practical cases. We consider randomized learners who predicts a distribution  $q_t \in \Delta(\bar{\mathcal{H}})$  on every round. We can define the value of the game as

$$\mathcal{V}_n(\mathcal{H}, \mathcal{Z}) = \inf_{q_1 \in \Delta(\bar{\mathcal{H}})} \sup_{z_1 \in \mathcal{Z}} \mathbb{E}_{\mathbf{h}_1 \sim q_1} \cdots \inf_{q_n \in \Delta(\bar{\mathcal{H}})} \sup_{z_n \in \mathcal{Z}} \mathbb{E}_{\mathbf{h}_n \sim q_n} \left[ \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \inf_{\mathbf{h} \in \bar{\mathcal{H}}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \right] \quad (4.2.1)$$

where  $\mathbf{h}_t$  has distribution  $q_t$ . We consider here the *adaptive* adversary who gets to choose instance  $z_t$  at round  $t$  based on the history of moves  $\mathbf{h}_{1:t-1}$  and  $z_{1:t-1}$ .

The above definition is stated in the extensive form, but can be equivalently written in a strategic form. Just like we used  $\mathbf{A}$  to describe the learner's strategy, analogously we can define the adversarial strategy as a sequence of  $\tau : \bigcup_{n \in \mathbb{N}} (\mathcal{H}^n \times \mathcal{Z}^n) \mapsto \Delta(\mathcal{Z})$  where  $\Delta(\mathcal{Z})$  refers to the set of all Borel probability distribution on  $\mathcal{Z}$ . The value can then be written as

$$\mathcal{V}_n(\mathcal{H}, \mathcal{Z}) = \inf_{\mathbf{A}} \sup_{\tau} \mathbb{E} \left\{ \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \right\}. \quad (4.2.2)$$

where it is understood that each  $\mathbf{h}_t$  and  $z_t$  are successively drawn according to law  $\mathbf{A}(\mathbf{h}_{1:t-1}, z_{1:t-1})$  and  $\tau(\mathbf{h}_{1:t-1}, z_{1:t-1})$  respectively. While the strategic notation is more succinct, it hides the important sequential structure of the problem. This is the reason why we opt for the more explicit, yet more cumbersome, extensive form. We are now ready to formally define online learnability of a problem.

**Definition 22.** We say that a problem is *online learnable* with respect to the given instance space  $\mathcal{Z}$  against target hypothesis set  $\mathcal{H}$  if

$$\limsup_{n \rightarrow \infty} \mathcal{V}_n(\mathcal{H}, \mathcal{Z}) = 0.$$

The first key step is to appeal to the minimax theorem and exchange the pairs of infima and suprema in Eq. (4.2.1). This dual formulation is easier to analyze because the choice of the player comes *after* the choice of the mixed strategy of the adversary. We remark that the minimax theorem holds under a very general assumption of weak compactness of  $\Delta(\bar{\mathcal{H}})$ . The assumptions on  $\bar{\mathcal{H}}$  that translate into weak compactness of  $\Delta(\bar{\mathcal{H}})$  are discussed in Appendix ???. Compactness under weak topology allows us to appeal to Theorem 23 stated below.

**Theorem 23.** *Let  $\bar{\mathcal{H}}$  and  $\mathcal{Z}$  be the sets satisfying the necessary conditions for the minimax theorem to hold, then*

$$\begin{aligned} \mathcal{V}_n(\mathcal{H}, \mathcal{Z}) &= \inf_{q_1 \in \Delta(\bar{\mathcal{H}})} \sup_{z_1 \in \mathcal{Z}} \mathbb{E}_{\mathbf{h}_1 \sim q_1} \cdots \inf_{q_n \in \Delta(\bar{\mathcal{H}})} \sup_{z_n \in \mathcal{Z}} \mathbb{E}_{\mathbf{h}_n \sim q_n} \left[ \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \right] \\ &= \sup_{p_1 \in \Delta(\mathcal{Z})} \mathbb{E}_{z_1 \sim p_1} \cdots \sup_{p_n \in \Delta(\mathcal{Z})} \mathbb{E}_{z_n \sim p_n} \left[ \frac{1}{n} \sum_{t=1}^n \inf_{\mathbf{h}_t \in \bar{\mathcal{H}}} \mathbb{E} [\ell(\mathbf{h}_t, z_t)] - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \right] \end{aligned} \quad (4.2.3)$$

The question of learnability in the online learning model is now reduced to the study of  $\mathcal{V}_n(\mathcal{H}, \mathcal{Z})$ , taking Eq. (4.2.3) as the starting point. In particular, under our definition, showing that the value grows sublinearly with  $n$  is equivalent to showing learnability.

One of the key notions introduced in this chapter is the complexity which we term *Sequential Rademacher complexity*. A natural generalization of Rademacher complexity [37, 38, 39], the sequential analogue possesses many of the nice properties of its classical cousin. The properties are proved in Section 4.3.1 and then used to show learnability for many of the examples in Section ??. The first step, however, is to show

that Sequential Rademacher complexity upper bounds the value of the game. This is the subject of the next section.

### 4.3 Sequential Rademacher Complexity

We propose the following definition of sequential Rademacher complexity of any function class  $\mathcal{F} \subseteq \mathbb{R}^{\mathcal{Z}}$ . The key difference from the classical notion is the dependence of the sequence of data on the sequence of signs (Rademacher random variables). As shown in the sequel, this dependence captures the sequential nature of the problem.

**Definition 23.** The *Sequential Rademacher Complexity* of a function class  $\mathcal{F} \subseteq \mathbb{R}^{\mathcal{Z}}$  is defined as

$$\mathfrak{R}_n^{\text{seq}}(\mathcal{F}) = \sup_{\mathbf{z}} \mathbb{E}_{\epsilon} \left[ \sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \epsilon_t f(z_t(\epsilon)) \right]$$

where the outer supremum is taken over all  $\mathcal{Z}$ -valued trees ( $\mathbf{z}$ ) of depth  $n$  and  $\epsilon = (\epsilon_1, \dots, \epsilon_n)$  is a sequence of i.i.d. Rademacher random variables.

In statistical learning, Rademacher complexity is shown to control uniform deviations of means and expectations, and this control is key for learnability in the “batch” setting. We now show that Sequential Rademacher complexity upper-bounds the value of the game, suggesting its importance for online learning (see Section ?? for a lower bound).

**Theorem 24.** *The minimax value of a randomized game is bounded as*

$$\mathcal{V}_n(\mathcal{H}, \mathcal{Z}) \leq 2\mathfrak{R}_n^{\text{seq}}(\mathcal{F})$$

where the function class  $\mathcal{F}$  is given by,  $\mathcal{F} = \{z \mapsto \ell(\mathbf{h}, z) : \mathbf{h} \in \mathcal{H}\}$ .

*Proof.* From Eq. Eq. (4.2.3),

$$\begin{aligned} \mathcal{V}_n(\mathcal{H}, \mathcal{Z}) &= \sup_{p_1} \mathbb{E}_{z_1 \sim p_1} \dots \sup_{p_n} \mathbb{E}_{z_n \sim p_n} \left[ \frac{1}{n} \sum_{t=1}^n \inf_{\mathbf{h}_t \in \mathcal{H}} \mathbb{E}_{z_t \sim p_t} [\ell(\mathbf{h}_t, z_t)] - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \right] \\ &= \sup_{p_1} \mathbb{E}_{z_1 \sim p_1} \dots \sup_{p_n} \mathbb{E}_{z_n \sim p_n} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \left\{ \frac{1}{n} \sum_{t=1}^n \inf_{\mathbf{h}_t \in \mathcal{H}} \mathbb{E}_{z_t \sim p_t} [\ell(\mathbf{h}_t, z_t)] - \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \right\} \right] \\ &\leq \sup_{p_1} \mathbb{E}_{z_1 \sim p_1} \dots \sup_{p_n} \mathbb{E}_{z_n \sim p_n} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \left\{ \frac{1}{n} \sum_{t=1}^n \mathbb{E}_{z_t \sim p_t} [\ell(\mathbf{h}, z_t)] - \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \right\} \right] \end{aligned} \quad (4.3.1)$$

The last step, in fact, is the first time we deviated from keeping equalities. The upper bound is obtained by

replacing each infimum by a particular choice  $\mathbf{h}$ . Now renaming variables we have,

$$\begin{aligned} \mathcal{V}_n(\mathcal{H}, \mathcal{Z}) &= \sup_{p_1} \mathbb{E}_{z_1 \sim p_1} \dots \sup_{p_n} \mathbb{E}_{z_n \sim p_n} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \left\{ \frac{1}{n} \sum_{t=1}^n \mathbb{E}_{z'_t \sim p_t} [\ell(\mathbf{h}, z'_t)] - \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \right\} \right] \\ &\leq \sup_{p_1} \mathbb{E}_{z_1 \sim p_1} \dots \sup_{p_n} \mathbb{E}_{z_n \sim p_n} \left[ \mathbb{E}_{z'_1 \sim p_1} \dots \mathbb{E}_{z'_n \sim p_n} \sup_{\mathbf{h} \in \mathcal{H}} \left\{ \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z'_t) - \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \right\} \right] \\ &\leq \sup_{p_1} \mathbb{E}_{z_1, z'_1 \sim p_1} \dots \sup_{p_n} \mathbb{E}_{z_n, z'_n \sim p_n} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \left\{ \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z'_t) - \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \right\} \right]. \end{aligned}$$

where the last two steps are using Jensen inequality for the supremum.

By the Key Technical Lemma (see Lemma 25 below) with  $\phi(u) = u$ ,

$$\begin{aligned} \sup_{p_1} \mathbb{E}_{z_1, z'_1 \sim p_1} \dots \sup_{p_n} \mathbb{E}_{z_n, z'_n \sim p_n} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \left\{ \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z'_t) - \ell(\mathbf{h}, z_t) \right\} \right] \\ \leq \sup_{z_1, z'_1} \mathbb{E}_{\epsilon_1} \dots \sup_{z_n, z'_n} \mathbb{E}_{\epsilon_n} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \epsilon_t (\ell(\mathbf{h}, z'_t) - \ell(\mathbf{h}, z_t)) \right] \end{aligned}$$

Thus,

$$\begin{aligned} \mathcal{V}_n(\mathcal{F}) &\leq \sup_{z_1, z'_1} \mathbb{E}_{\epsilon_1} \dots \sup_{z_n, z'_n} \mathbb{E}_{\epsilon_n} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \epsilon_t (\ell(\mathbf{h}, z'_t) - \ell(\mathbf{h}, z_t)) \right] \\ &\leq \sup_{z_1, z'_1} \mathbb{E}_{\epsilon_1} \dots \sup_{z_n, z'_n} \mathbb{E}_{\epsilon_n} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \left\{ \frac{1}{n} \sum_{t=1}^n \epsilon_t \ell(\mathbf{h}, z'_t) \right\} + \sup_{\mathbf{h} \in \mathcal{H}} \left\{ \frac{1}{n} \sum_{t=1}^n -\epsilon_t \ell(\mathbf{h}, z_t) \right\} \right] \\ &= 2 \sup_{z_1} \mathbb{E}_{\epsilon_1} \dots \sup_{z_n} \mathbb{E}_{\epsilon_n} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \epsilon_t \ell(\mathbf{h}, z_t) \right] \end{aligned}$$

Now, we need to move the suprema over  $z_t$ 's outside. This is achieved via an idea similar to skolemization in logic. We basically exploit the identity

$$\mathbb{E}_{\epsilon_{1:t-1}} \left[ \sup_{z_t} G(\epsilon_{1:t-1}, z_t) \right] = \sup_{z_t} \mathbb{E}_{\epsilon_{1:t-1}} [G(\epsilon_{1:t-1}, z_t(\epsilon_{1:t-1}))]$$

that holds for any  $G : \{\pm 1\}^{t-1} \times \mathcal{Z} \mapsto \mathbb{R}$ . On the right the supremum is over functions  $z_t : \{\pm 1\}^{t-1} \rightarrow \mathcal{Z}$ .

Using this identity once, we get,

$$\mathcal{V}_n(\mathcal{F}) \leq \frac{2}{n} \sup_{z_1, z_2} \left\{ \mathbb{E}_{\epsilon_1, \epsilon_2} \left[ \sup_{z_3} \dots \sup_{z_n} \left\{ \mathbb{E}_{\epsilon_n} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \left\{ \epsilon_1 \ell(\mathbf{h}, z_1) + \epsilon_2 \ell(\mathbf{h}, z_2(\epsilon_1)) + \sum_{t=3}^n \epsilon_t \ell(\mathbf{h}, z_t) \right\} \right] \dots \right] \right] \right\}$$

Now, use the identity  $n - 2$  more times to successively move the supremums over  $z_3, \dots, z_n$  outside, to get

$$\begin{aligned} \mathcal{V}_n(\mathcal{F}) &\leq \frac{2}{n} \sup_{z_1, z_2, \dots, z_n} \mathbb{E}_{\epsilon_1, \dots, \epsilon_n} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \left\{ \epsilon_1 \ell(\mathbf{h}, z_1) + \sum_{t=2}^n \epsilon_t \ell(\mathbf{h}, z_t(\epsilon_{1:t-1})) \right\} \right] \\ &= 2 \sup_z \mathbb{E}_{\epsilon_1, \dots, \epsilon_n} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \left\{ \frac{1}{n} \sum_{t=1}^n \epsilon_t \ell(\mathbf{h}, z_t(\epsilon)) \right\} \right] \end{aligned}$$

where the last supremum is over  $\mathcal{Z}$ -valued trees of depth  $n$ . Thus we have proved the required statement.  $\square$

Theorem 24 relies on the following technical lemma, which will be used again in Section ???. Its proof requires considerably more work than the classical symmetrization proof [40, 39] due to the non-i.i.d. nature of the sequences.

**Lemma 25** (Key Technical Lemma). *Let  $(z_1, \dots, z_n) \in \mathcal{Z}^n$  be a sequence distributed according to  $\mathcal{D}$  and let  $(z'_1, \dots, z'_n) \in \mathcal{Z}^n$  be a tangent sequence. Let  $\phi : \mathbb{R} \mapsto \mathbb{R}$  be a measurable function. Then*

$$\begin{aligned} \sup_{p_1} \mathbb{E}_{z_1, z'_1 \sim p_1} \dots \sup_{p_n} \mathbb{E}_{z_n, z'_n \sim p_n} \left[ \phi \left( \sup_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \Delta_{\mathbf{h}}(z_t, z'_t) \right) \right] \\ \leq \sup_{z_1, z'_1} \mathbb{E}_{\epsilon_1} \dots \sup_{z_n, z'_n} \mathbb{E}_{\epsilon_n} \left[ \phi \left( \sup_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \epsilon_t \Delta_{\mathbf{h}}(z_t, z'_t) \right) \right] \end{aligned}$$

where  $\epsilon_1, \dots, \epsilon_n$  are independent (of each other and everything else) Rademacher random variables and  $\Delta_{\mathbf{h}}(z_t, z'_t) = \frac{1}{n} (\ell(\mathbf{h}, z'_t) - \ell(\mathbf{h}, z_t))$ . The inequality also holds when an absolute value of the sum is introduced on both sides.

Before proceeding, let us give some intuition behind the attained bounds. Theorem 23 establishes an upper bound on the value of the game in terms of a stochastic process on  $\mathcal{Z}$ . In general, it is difficult to get a handle on the behavior of this process. The key idea is to relate this process to a symmetrized version, and then pass to a new process obtained by fixing a binary tree  $z$  and then following a path in  $z$  using i.i.d. coin flips. In some sense, we are replacing the  $\sigma$ -algebra generated by the random process of Theorem 23 by a simpler process generated by Rademacher random variables and a tree  $z$ . It can be shown that the original process and the simpler process are in fact close in a certain sense, yet the process generated by the Rademacher random variables is much easier to work with. It is precisely due to symmetrization that the trees we consider in this work are binary trees and not full game trees. Passing to binary trees allows us to define covering numbers, combinatorial parameters, and other analogues of the classical notions from statistical learning theory.

### 4.3.1 Structural Results

Being able to bound complexity of a function class by a complexity of a simpler class is of great utility for proving bounds. In statistical learning theory, such structural results are obtained through properties of Rademacher averages [39, 38]. In particular, the contraction inequality due to Ledoux and Talagrand [41,



Corollary 3.17], allows one to pass from a composition of a Lipschitz function with a class to the function class itself. This wonderful property permits easy convergence proofs for a vast array of problems.

We show that the notion of Sequential Rademacher complexity also enjoys many of the same properties. In Section ??, the effectiveness of the results is illustrated on a number of examples.

The next lemma bounds the Sequential Rademacher complexity for the product of function classes.

**Lemma 26.** *Let  $\mathcal{F} = \mathcal{F}_1 \times \dots \times \mathcal{F}_k$  where each  $\mathcal{F}_j \subset [-1, 1]^{\mathcal{Z}}$ . Also let  $\phi : \mathbb{R}^k \times \mathcal{Z} \mapsto \mathbb{R}$  be such that  $\phi(\cdot, z)$  is  $L$ -Lipschitz w.r.t.  $\|\cdot\|_\infty$  norm for any  $z \in \mathcal{Z}$ . Then we have that*

$$\mathfrak{R}_n^{\text{seq}}(\phi \circ \mathcal{F}) \leq 8L \left(1 + 4\sqrt{2} \log^{3/2}(en^2)\right) \sum_{j=1}^k \mathfrak{R}_n^{\text{seq}}(\mathcal{F}_j)$$

as long as  $\mathfrak{R}_n^{\text{seq}}(\mathcal{F}_j) \geq 1$  for each  $j$ .

As a special case of the result, we get a sequential counterpart of the Ledoux-Talagrand [41] contraction inequality.

**Lemma 27.** *Fix a class  $\mathcal{F} \subseteq [-1, 1]^{\mathcal{Z}}$  with  $\mathfrak{R}_n^{\text{seq}}(\mathcal{F}) \geq 1$  and a function  $\phi : \mathbb{R} \times \mathcal{Z} \mapsto \mathbb{R}$ . Assume, for all  $z \in \mathcal{Z}$ ,  $\phi(\cdot, z)$  is a Lipschitz function with a constant  $L$ .*

$$\mathfrak{R}_n^{\text{seq}}(\phi(\mathcal{F})) \leq 8L \left(1 + 4\sqrt{2} \log^{3/2}(en^2)\right) \cdot \mathfrak{R}_n^{\text{seq}}(\mathcal{F})$$

where  $\phi(\mathcal{F}) = \{z \mapsto \phi(f(z), z) : f \in \mathcal{F}\}$ .

We remark that the lemma above encompasses the case of a Lipschitz  $\phi : \mathbb{R} \mapsto \mathbb{R}$ , as stated in [41, 38]. However, here we get an extra logarithmic factor in  $n$  which is absent in the classical case. Whether the same can be proved in the sequential case remains an open question.

We state another useful corollary of Lemma 26.

**Corollary 28.** *For a fixed binary function  $b : \{\pm 1\}^k \mapsto \{\pm 1\}$  and classes  $\mathcal{F}_1, \dots, \mathcal{F}_k$  of  $\{\pm 1\}$ -valued functions,*

$$\mathfrak{R}_n^{\text{seq}}(g(\mathcal{F}_1, \dots, \mathcal{F}_k)) \leq \mathcal{O}\left(\log^{3/2}(n)\right) \sum_{j=1}^k \mathfrak{R}_n^{\text{seq}}(\mathcal{F}_j)$$

In the next proposition, we summarize some useful properties of Sequential Rademacher complexity (see [39, 38] for the results in the i.i.d. setting)

**Proposition 29.** *Sequential Rademacher complexity satisfies the following properties.*

1. If  $\mathcal{F} \subset \mathcal{G}$ , then  $\mathfrak{R}_n^{\text{seq}}(\mathcal{F}) \leq \mathfrak{R}_n^{\text{seq}}(\mathcal{G})$ .
2.  $\mathfrak{R}_n^{\text{seq}}(\mathcal{F}) = \mathfrak{R}_n^{\text{seq}}(\text{conv}(\mathcal{F}))$ .

3.  $\mathfrak{N}_n^{\text{seq}}(c\mathcal{F}) = |c|\mathfrak{N}_n^{\text{seq}}(\mathcal{F})$  for all  $c \in \mathbb{R}$ .
4. For any  $h$ ,  $\mathfrak{N}_n^{\text{seq}}(\mathcal{F} + h) = \mathfrak{N}_n^{\text{seq}}(\mathcal{F})$  where  $\mathcal{F} + h = \{f + h : f \in \mathcal{F}\}$ .

## 4.4 Sequential Covering Number and Combinatorial Parameters

In statistical learning theory, learnability for binary classes of functions is characterized by the Vapnik-Chervonenkis combinatorial dimension [42]. For real-valued function classes, the corresponding notions are the scale-sensitive dimensions, such as  $P_\gamma$  [43, 44]. For online learning, the notion characterizing learnability for binary prediction in the realizable case has been introduced by Littlestone [34] and extended to the non-realizable case of binary prediction by Shai Ben-David, Dávid Pál and Shai Shalev-Shwartz [35]. Next, we define the Littlestone’s dimension [34, 35] and propose its scale-sensitive versions for real-valued function classes. In the sequel, these combinatorial parameters are shown to control the growth of covering numbers on trees. In the setting of prediction, the combinatorial parameters are shown to exactly characterize learnability (see Section ??).

**Definition 24.** A  $\mathcal{Z}$ -valued tree  $z$  of depth  $d$  is *shattered* by a function class  $\mathcal{F} \subseteq \{\pm 1\}^{\mathcal{Z}}$  if for all  $\epsilon \in \{\pm 1\}^d$ , there exists  $f \in \mathcal{F}$  such that  $f(z_t(\epsilon)) = \epsilon_t$  for all  $t \in [d]$ . The *Littlestone dimension*  $\text{Ldim}(\mathcal{F}, \mathcal{Z})$  is the largest  $d$  such that  $\mathcal{F}$  shatters an  $\mathcal{Z}$ -valued tree of depth  $d$ .

**Definition 25.** An  $\mathcal{Z}$ -valued tree  $z$  of depth  $d$  is  $\alpha$ -*shattered* by a function class  $\mathcal{F} \subseteq \mathbb{R}^{\mathcal{Z}}$ , if there exists an  $\mathbb{R}$ -valued tree  $s$  of depth  $d$  such that

$$\forall \epsilon \in \{\pm 1\}^d, \exists f \in \mathcal{F} \text{ s.t. } \forall t \in [d], \epsilon_t(f(z_t(\epsilon)) - s_t(\epsilon)) \geq \alpha/2$$

The tree  $s$  is called the *witness to shattering*. The *fat-shattering dimension*  $\text{fat}_\alpha^{\text{seq}}(\mathcal{F}, \mathcal{Z})$  at scale  $\alpha$  is the largest  $d$  such that  $\mathcal{F}$   $\alpha$ -shatters an  $\mathcal{Z}$ -valued tree of depth  $d$ .

With these definitions it is easy to see that  $\text{fat}_\alpha^{\text{seq}}(\mathcal{F}, \mathcal{Z}) = \text{Ldim}(\mathcal{F}, \mathcal{Z})$  for a binary-valued function class  $\mathcal{F} \subseteq \{\pm 1\}^{\mathcal{Z}}$  for any  $0 < \alpha \leq 2$ .

When  $\mathcal{Z}$  and/or  $\mathcal{F}$  is understood from the context, we will simply write  $\text{fat}_\alpha^{\text{seq}}$  or  $\text{fat}_\alpha^{\text{seq}}(\mathcal{F})$  instead of  $\text{fat}_\alpha^{\text{seq}}(\mathcal{F}, \mathcal{Z})$ . Furthermore, we will write  $\text{fat}_\alpha^{\text{seq}}(\mathcal{F}, z)$  for  $\text{fat}_\alpha^{\text{seq}}(\mathcal{F}, \text{Img}(z))$ . In other words,  $\text{fat}_\alpha^{\text{seq}}(\mathcal{F}, z)$  is the largest  $d$  such that  $\mathcal{F}$   $\alpha$ -shatters a tree  $y$  of depth  $d$  with  $\text{Img}(y) \subseteq \text{Img}(z)$ .

Let us mention that if trees  $z$  are defined by constant mappings  $z_t(\epsilon) = z_t$ , the combinatorial parameters coincide with the Vapnik-Chervonenkis dimension and with the scale-sensitive dimension  $P_\gamma$ . Therefore, the notions we are studying are strict “temporal” generalizations of the VC theory.

As in statistical learning theory, the combinatorial parameters are only useful if they can be shown to capture that aspect of  $\mathcal{F}$  which is important for learnability. In particular, a “size” of a function class is known to be related to complexity of learning from i.i.d. data., and the classical way to measure “size” is through a cover or of a packing set. We propose the following definitions for online learning.

**Definition 26.** A set  $V$  of  $\mathbb{R}$ -valued trees of depth  $n$  is an  $\alpha$ -cover (with respect to  $\ell_p$ -norm) of  $\mathcal{F} \subseteq \mathbb{R}^{\mathcal{Z}}$  on a tree  $z$  of depth  $n$  if

$$\forall f \in \mathcal{F}, \forall \epsilon \in \{\pm 1\}^n \exists \mathbf{v} \in V \text{ s.t. } \left( \frac{1}{n} \sum_{t=1}^n |\mathbf{v}_t(\epsilon) - f(\mathbf{z}_t(\epsilon))|^p \right)^{1/p} \leq \alpha$$

The *covering number* of a function class  $\mathcal{F}$  on a given tree  $z$  is defined as

$$\mathfrak{N}_p^{\text{seq}}(\alpha, \mathcal{F}, z) = \min\{|V| : V \text{ is an } \alpha\text{-cover w.r.t. } \ell_p\text{-norm of } \mathcal{F} \text{ on } z\}.$$

Further define  $\mathfrak{N}_p^{\text{seq}}(\alpha, \mathcal{F}, n) = \sup_z \mathfrak{N}_p^{\text{seq}}(\alpha, \mathcal{F}, z)$ , the maximal  $\ell_p$  covering number of  $\mathcal{F}$  over depth  $n$  trees.

In particular, a set  $V$  of  $\mathbb{R}$ -valued trees of depth  $n$  is a *0-cover* of  $\mathcal{F} \subseteq \mathbb{R}^{\mathcal{Z}}$  on a tree  $z$  of depth  $n$  if

$$\forall f \in \mathcal{F}, \forall \epsilon \in \{\pm 1\}^n \exists \mathbf{v} \in V \text{ s.t. } \mathbf{v}_t(\epsilon) = f(\mathbf{z}_t(\epsilon))$$

We denote by  $\mathfrak{N}^{\text{seq}}(0, \mathcal{F}, z)$  the size of a smallest 0-cover on  $z$  and  $\mathfrak{N}^{\text{seq}}(0, \mathcal{F}, n) = \sup_z \mathfrak{N}^{\text{seq}}(0, \mathcal{F}, z)$ .

Let us discuss a subtle point. The 0-cover should not be mistaken for the size  $|\mathcal{F}(z)|$  of the projection of  $\mathcal{F}$  onto the tree  $z$ , and the same care should be taken when dealing with  $\alpha$ -covers. Let us illustrate this with an example. Consider a tree  $z$  of depth  $n$  and suppose for simplicity that  $|\text{Img}(z)| = 2^n - 1$ , i.e. the values of  $z$  are all distinct. Suppose  $\mathcal{F}$  consists of  $2^{n-1}$  binary-valued functions defined as zero on all of  $\text{Img}(z)$  except for a single value of  $\text{Img}(z_n)$ . In plain words, each function is zero everywhere on the tree except for a single leaf. While the projection  $\mathcal{F}(z)$  has  $2^{n-1}$  distinct trees, the size of a 0-cover is only 2. It is enough to take an all-zero function  $g_0$  along with a function  $g_1$  which is zero on all of  $\text{Img}(z)$  except  $\text{Img}(z_n)$  (i.e. on the leaves). It is easy to verify that  $g_0(z)$  and  $g_1(z)$  provide a 0-cover for  $\mathcal{F}$  on  $z$ , and therefore, unlike  $|\mathcal{F}(z)|$ , the size of the cover does not grow with  $n$ . The example is encouraging: our definition of a cover captures the fact that the function class is “simple” for any given path.

Next, we naturally propose a definition of a packing.

**Definition 27.** A set  $V$  of  $\mathbb{R}$ -valued trees of depth  $n$  is said to be  $\alpha$ -separated if

$$\forall \mathbf{v} \in V, \exists \epsilon \in \{\pm 1\}^n \text{ s.t. } \forall \mathbf{w} \in V \setminus \{\mathbf{v}\} \left( \frac{1}{n} \sum_{t=1}^n |\mathbf{v}_t(\epsilon) - \mathbf{w}_t(\epsilon)|^p \right)^{1/p} > \alpha$$

The *weak packing number*  $\mathcal{D}_p(\alpha, \mathcal{F}, z)$  of a function class  $\mathcal{F}$  on a given tree  $z$  is the size of the largest  $\alpha$ -separated subset of  $\{f(z) : f \in \mathcal{F}\}$ .

**Definition 28.** A set  $V$  of  $\mathbb{R}$ -valued trees of depth  $n$  is said to be *strongly*  $\alpha$ -separated if

$$\exists \epsilon \in \{\pm 1\}^n \text{ s.t. } \forall \mathbf{v}, \mathbf{w} \in V, \mathbf{v} \neq \mathbf{w} \left( \frac{1}{n} \sum_{t=1}^n |\mathbf{v}_t(\epsilon) - \mathbf{w}_t(\epsilon)|^p \right)^{1/p} > \alpha$$

The *strong packing number*  $\mathcal{M}_p(\alpha, \mathcal{F}, \mathbf{z})$  of a function class  $\mathcal{F}$  on a given tree  $\mathbf{z}$  is the size of the largest strongly  $\alpha$ -separated subset of  $\{f(\mathbf{z}) : f \in \mathcal{F}\}$ .

Note the distinction between the packing number and the strong packing number. For the former, it must be that every member of the packing is  $\alpha$ -separated from every other member on *some* path. For the latter, there must be a path on which every member of the packing is  $\alpha$ -separated from every other member. This distinction does not arise in the classical scenario of “batch” learning. We observe that if a tree  $\mathbf{z}$  is defined by constant mappings  $z_t = x_t$ , the two notions of packing and strong packing coincide, i.e.  $\mathcal{D}_p(\alpha, \mathcal{F}, \mathbf{z}) = \mathcal{M}_p(\alpha, \mathcal{F}, \mathbf{z})$ . The following lemma gives a relationship between covering numbers and the two notions of packing numbers. The form of this should be familiar, except for the distinction between the two types of packing numbers.

**Lemma 30.** *For any  $\mathcal{F} \subseteq \mathbb{R}^{\mathcal{Z}}$ , any  $\mathcal{Z}$ -valued tree  $\mathbf{z}$  of depth  $n$ , and any  $\alpha > 0$*

$$\mathcal{M}_p(2\alpha, \mathcal{F}, \mathbf{z}) \leq \mathfrak{N}_p^{\text{seq}}(\alpha, \mathcal{F}, \mathbf{z}) \leq \mathcal{D}_p(\alpha, \mathcal{F}, \mathbf{z}).$$

It is important to note that the gap between the two types of packing can be as much as  $2^n$ .

#### 4.4.1 A Combinatorial Upper Bound

We now relate the combinatorial parameters introduced in the previous section to the size of a cover. In the binary case ( $k = 1$  below), a reader might notice a similarity of Theorems 31 and 33 to the classical results due to Sauer [45], Shelah [46] (also, Perles and Shelah), and Vapnik and Chervonenkis [42]. There are several approaches to proving what is often called the Sauer-Shelah lemma. We opt for the inductive-style proof (e.g. Alon and Spencer [47]). Dealing with trees, however, requires more work than in the VC case .

**Theorem 31.** *Let  $\mathcal{F} \subseteq \{0, \dots, k\}^{\mathcal{Z}}$  be a class of functions with  $\text{fat}_2^{\text{seq}}(\mathcal{F}) = d$ . Then*

$$\mathfrak{N}_{\infty}^{\text{seq}}(1/2, \mathcal{F}, n) \leq \sum_{i=0}^d \binom{n}{i} k^i \leq (ekn)^d .$$

Furthermore, for  $n \geq d$

$$\sum_{i=0}^d \binom{n}{i} k^i \leq \left( \frac{ekn}{d} \right)^d .$$

Armed with Theorem 31, we can approach the problem of bounding the size of a cover at an  $\alpha$  scale by a discretization trick. For the classical case of a cover based on a set points, the discretization idea appears in [43, 48]. When passing from the combinatorial result to the cover at scale  $\alpha$  in Corollary 32, it is crucial that Theorem 31 is in terms of  $\text{fat}_2^{\text{seq}}(\mathcal{F})$  and not  $\text{fat}_1^{\text{seq}}(\mathcal{F})$ . This point can be seen in the proof of Corollary 32 (also see [48]): the discretization process can assign almost identical function values to discrete values which differ by 1. This explains why the combinatorial result of Theorem 31 is proved for the 2-shattering dimension.

We now show that the covering numbers are bounded in terms of the fat-shattering dimension.

**Corollary 32.** *Suppose  $\mathcal{F}$  is a class of  $[-1, 1]$ -valued functions on  $\mathcal{Z}$ . Then for any  $\alpha > 0$ , any  $n > 0$ , and any  $\mathcal{Z}$ -valued tree  $z$  of depth  $n$ ,*

$$\mathfrak{N}_1^{\text{seq}}(\alpha, \mathcal{F}, z) \leq \mathfrak{N}_2^{\text{seq}}(\alpha, \mathcal{F}, z) \leq \mathfrak{N}_\infty^{\text{seq}}(\alpha, \mathcal{F}, z) \leq \left(\frac{2en}{\alpha}\right)^{\text{fat}_\alpha^{\text{seq}}(\mathcal{F})}$$

With a proof similar to Theorem 31, a bound on the 0-cover can be proved in terms of the  $\text{fat}_1^{\text{seq}}(\mathcal{F})$  combinatorial parameter. Of particular interest is the case  $k = 1$ , when  $\text{fat}_1^{\text{seq}}(\mathcal{F}) = \text{Ldim}(\mathcal{F})$ .

**Theorem 33.** *Let  $\mathcal{F} \subseteq \{0, \dots, k\}^{\mathcal{Z}}$  be a class of functions with  $\text{fat}_1^{\text{seq}}(\mathcal{F}) = d$ . Then*

$$\mathfrak{N}^{\text{seq}}(0, \mathcal{F}, n) \leq \sum_{i=0}^d \binom{n}{i} k^i \leq (ekn)^d.$$

Furthermore, for  $n \geq d$

$$\sum_{i=0}^d \binom{n}{i} k^i \leq \left(\frac{ekn}{d}\right)^d.$$

In particular, the result holds for binary-valued function classes ( $k = 1$ ), in which case  $\text{fat}_1^{\text{seq}}(\mathcal{F}) = \text{Ldim}(\mathcal{F})$ .

When bounding deviations of means from expectations uniformly over the function class, the usual approach proceeds by a symmetrization argument [49] followed by passing to a cover of the function class and a union bound (e.g. [39]). Alternatively, a more refined *chaining* analysis integrates over covering at different scales (e.g. [50]). By following the same path, we are able to prove a number of similar results for our setting. In the next section we present a bound similar to Massart’s finite class lemma [51, Lemma 5.2], and in the following section this result will be used when integrating over different scales for the cover.

#### 4.4.2 Finite Class Lemma and the Chaining Method

**Lemma 34.** *For any finite set  $V$  of  $\mathbb{R}$ -valued trees of depth  $n$  we have that*

$$\mathbb{E}_\epsilon \left[ \max_{\mathbf{v} \in V} \sum_{t=1}^n \epsilon_t \mathbf{v}_t(\epsilon) \right] \leq \sqrt{2 \log(|V|) \max_{\mathbf{v} \in V} \max_{\epsilon \in \{\pm 1\}^n} \sum_{t=1}^n \mathbf{v}_t(\epsilon)^2}$$

A simple consequence of the above lemma is that if  $\mathcal{F} \subseteq [0, 1]^{\mathcal{Z}}$  is a finite class, then for any given tree  $z$  we have that

$$\mathbb{E}_\epsilon \left[ \max_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \epsilon_t f(z_t(\epsilon)) \right] \leq \frac{1}{n} \mathbb{E}_\epsilon \left[ \max_{\mathbf{v} \in \mathcal{F}(z)} \sum_{t=1}^n \epsilon_t \mathbf{v}_t(\epsilon) \right] \leq \sqrt{\frac{2 \log |\mathcal{F}|}{n}}.$$

Note that if  $f \in \mathcal{F}$  is associated with an “expert”, this result combined with Theorem 24 yields a bound given

by the exponential weighted average forecaster algorithm (see [52]). In Section ?? we discuss this case in more detail. However, as we show next, Lemma 34 goes well beyond just finite classes and can be used to get an analog of Dudley entropy bound [17] for the online setting through a chaining argument.

**Definition 29.** The *Integrated complexity* of a function class  $\mathcal{F} \subseteq [-1, 1]^{\mathcal{Z}}$  is defined as

$$\mathfrak{D}_n^{\text{seq}}(\mathcal{F}) = \inf_{\alpha} \left\{ 4\alpha + 12 \int_{\alpha}^1 \sqrt{\frac{\log \mathcal{N}_2(\delta, \mathcal{F}, n)}{n}} d\delta \right\}.$$

To prove the next theorem, we consider covers of the class  $\mathcal{F}$  at different scales that form a geometric progression. We zoom into a given function  $f \in \mathcal{F}$  using covering elements at successive scales. This zooming in procedure is visualized as forming a chain that consists of links connecting elements of covers at successive scales. The Rademacher complexity of  $\mathcal{F}$  can then be bounded by controlling the Rademacher complexity of the link classes, i.e. the class consisting of differences of functions from covers at neighbouring scales. This last part of the argument is the place where our proof becomes a bit more involved than the classical case.

**Theorem 35.** For any function class  $\mathcal{F} \subseteq [-1, 1]^{\mathcal{Z}}$ ,

$$\mathfrak{R}_n^{\text{seq}}(\mathcal{F}) \leq \mathfrak{D}_n^{\text{seq}}(\mathcal{F}).$$

If a fat-shattering dimension of the class can be controlled, Corollary 32 together with Theorem 35 yield an upper bound on the value.

We can now show that, in fact, the two complexity measures  $\mathfrak{R}_n^{\text{seq}}(\mathcal{F})$  and  $\mathfrak{D}_n^{\text{seq}}(\mathcal{F})$  are equivalent, up to a logarithmic factor. Before stating this result formally, we prove the following lemma which asserts that the fat-shattering dimensions at “large enough” scales cannot be too large.

**Lemma 36.** For any  $\beta > \frac{2}{n} \mathfrak{R}_n^{\text{seq}}(\mathcal{F})$ , we have that  $\text{fat}_{\beta}^{\text{seq}}(\mathcal{F}) < n$ .

The following lemma complements Theorem 35.

**Lemma 37.** For any function class  $\mathcal{F} \subseteq [-1, 1]^{\mathcal{Z}}$ , we have that

$$\mathfrak{D}_n^{\text{seq}}(\mathcal{F}) \leq 8 \mathfrak{R}_n^{\text{seq}}(\mathcal{F}) \left( 1 + 4\sqrt{2} \log^{3/2}(en^2) \right)$$

as long as  $\mathfrak{R}_n^{\text{seq}}(\mathcal{F}) \geq \frac{1}{n}$ .

## 4.5 Martingale Uniform Convergence

As we discussed in the previous chapter, in the statistical learning setting learnability of supervised learning problem is equivalent to the so called uniform Glivenko-Cantelli property (uniform convergence) of the class. The property refers to the empirical averages converging to expected value of the function for any fixed

distribution (samples drawn i.i.d.) and uniformly over the function class almost surely. Tools like classical Rademacher complexity, statistical covering numbers and statistical fat-shattering dimension can be used in analyzing rate of uniform convergence of empirical average to expected value of the function for samples drawn i.i.d from any fixed distribution. Analogously sequential counterparts of these complexity measures can be used to bound rate of uniform convergence of average value of the function to average conditional expectation of the function values for arbitrary distributions over sequence of random variables. In fact in the proof of Theorem 24 we already encountered this general martingale uniform convergence in expectation. Specifically Equation 4.3.1 which we showed is bounded by the sequential Rademacher complexity. We now formally define Universal Uniform Convergence which is analogous to the usual definition of uniform Glivenko-Cantelli property for general dependent processes and in this section we will show how the sequential complexity measures provide tools for bounding this martingale version of uniform convergence.

**Definition 30.** A function class  $\mathcal{F}$  satisfies a *Universal Uniform Convergence* if for all  $\alpha > 0$ ,

$$\lim_{N \rightarrow \infty} \sup_{\mathcal{D}} \mathbb{P}_{\mathcal{D}} \left( \sup_{n \geq N} \sup_{f \in \mathcal{F}} \frac{1}{n} \left| \sum_{t=1}^n (f(z_t) - \mathbb{E}_{t-1}[f(z_t)]) \right| > \alpha \right) = 0$$

where the supremum is over distributions  $\mathcal{D}$  over infinite sequences  $(x_1, \dots, x_n, \dots)$

We remark that the notion of uniform Glivenko-Cantelli classes is recovered if the supremum is taken over i.i.d. distributions. The theorem below shows that finite fat shattering dimension at all scales is a sufficient condition for *Universal Uniform Convergence*.

**Theorem 38.** Let  $\mathcal{F}$  be a class of  $[-1, 1]$ -valued functions. If  $\text{fat}_{\alpha}^{\text{seq}}(\mathcal{F})$  is finite for all  $\alpha > 0$ , then  $\mathcal{F}$  satisfies *Universal Uniform Convergence*.

The proof follows from the Lemma 39 and Lemma 40 below, while Lemma 41 is an even stronger version of Lemma 40. We remark that Lemma 39 is the “in-probability” version of sequential symmetrization technique of Theorem 24 and Lemma 41 is the “in-probability” version of Theorem 35.

**Lemma 39.** Let  $\mathcal{F}$  be a class of  $[-1, 1]$ -valued functions. Then for any  $\alpha > 0$

$$\mathbb{P}_{\mathcal{D}} \left( \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n (f(z_t) - \mathbb{E}_{t-1}[f(z_t)]) \right| > \alpha \right) \leq 4 \sup_{\mathbf{z}} \mathbb{P}_{\epsilon} \left( \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n \epsilon_t f(\mathbf{z}_t(\epsilon)) \right| > \alpha/4 \right)$$

**Lemma 40.** Let  $\mathcal{F}$  be a class of  $[-1, 1]$ -valued functions. For any  $\mathcal{Z}$ -valued tree  $\mathbf{z}$  of depth  $n$  and  $\alpha > 0$

$$\mathbb{P}_{\epsilon} \left( \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n \epsilon_t f(\mathbf{z}_t(\epsilon)) \right| > \alpha/4 \right) \leq 2\mathcal{N}_1(\alpha/8, \mathcal{F}, \mathbf{z}) e^{-n\alpha^2/128} \leq 2 \left( \frac{16en}{\alpha} \right)^{\text{fat}_{\alpha/8}^{\text{seq}}} e^{-n\alpha^2/128}$$

Next, we show that the sequential Rademacher complexity is, in some sense, the “right” complexity measure even when one considers high probability statements.

**Lemma 41.** *Let  $\mathcal{F}$  be a class of  $[-1, 1]$ -valued functions and suppose  $\text{fat}_\alpha^{\text{seq}}(\mathcal{F})$  is finite for all  $\alpha > 0$ . Then for any  $\theta > \sqrt{8/n}$ , for any  $\mathcal{Z}$ -valued tree  $\mathbf{z}$  of depth  $n$ ,*

$$\begin{aligned} & \mathbb{P}_\epsilon \left( \sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{t=1}^n \epsilon_t f(\mathbf{z}_t(\epsilon)) \right| > 128 \left( 1 + \theta \sqrt{n} \log^{3/2}(2n) \right) \cdot \mathfrak{N}_n^{\text{seq}}(\mathcal{F}) \right) \\ & \leq \mathbb{P}_\epsilon \left( \sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{t=1}^n \epsilon_t f(\mathbf{z}_t(\epsilon)) \right| > \inf_{\alpha > 0} \left\{ 4\alpha + 12\theta \int_\alpha^1 \sqrt{\frac{\log \mathfrak{N}_\infty^{\text{seq}}(\delta, \mathcal{F}, n)}{n}} d\delta \right\} \right) \\ & \leq L e^{-\frac{n\theta^2}{4}} \end{aligned}$$

where  $L$  is a constant such  $L > \sum_{j=1}^{\infty} \mathfrak{N}_\infty^{\text{seq}}(2^{-j}, \mathcal{F}, n)^{-1}$ .

While throughout this chapter we are mostly concerned with expected versions of minimax regret and the corresponding complexities, the above lemmas can be employed to give an analogous in-probability treatment. To obtain such in-probability statements, the value is defined as the minimax probability of regret exceeding a threshold.

## 4.6 Characterizing Learnability of Supervised Learning Problem

In this section we study the specific case of online supervised learning problem. In this setting, the instance space  $\mathcal{Z}$  is of form  $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$  where  $\mathcal{X}$  is some arbitrary input space and  $\mathcal{Y} \subset \mathbb{R}$ . We shall assume that  $\mathcal{Y} \subseteq [-1, 1]$  (of course the bound of 1 can be changed to an arbitrary value). The target hypothesis set  $\mathcal{H}$  for the supervised learning problem corresponds to a set of functions that map input instances from  $\mathcal{X}$  to some predicted label in  $[-1, 1]$ , that is  $\mathcal{H} \subset [-1, 1]^\mathcal{X}$ .  $\bar{\mathcal{H}} \subset [-1, 1]^\mathcal{X}$  can be an arbitrary superset of  $\mathcal{H}$  and the results in fact hold even if  $\bar{\mathcal{H}} = \mathcal{H}$  (ie. proper learning setting). The loss function we consider is for the form

$$\ell(\mathbf{h}, (x, y)) = |\mathbf{h}(x) - y|$$

In the online supervised learning problem at each round  $t$ , the player picks hypothesis  $\mathbf{h}_t \in \bar{\mathcal{H}}$  and the adversary provides input target pair  $(x_t, y_t)$  and the player suffers loss  $|\mathbf{h}_t(x_t) - y_t|$ . Note that if  $\bar{\mathcal{H}} \subseteq \{\pm 1\}^\mathcal{X}$  and each  $y_t \in \{\pm 1\}$  then the problem boils down to binary classification problem.

Though we use the absolute loss in this section, it is easy to see that all the results hold (with modified rates) for any loss  $\ell(\mathbf{h}(x), y)$  which is such that for all  $\mathbf{h}, x$  and  $y$ ,

$$\phi(\ell(\hat{y}, y)) \leq |\hat{y} - y| \leq \Phi(\ell(\hat{y}, y))$$

where  $\Phi$  and  $\phi$  are monotonically increasing functions. For instance the squared loss is a classic example.

To formally define the value of the online supervised learning game, fix a set of labels  $\mathcal{Y} \subseteq [-1, 1]$ . For the sake of brevity, we shall use the notation  $\mathcal{V}_n^S(\mathcal{H}) = \mathcal{V}_n(\mathcal{H}, \mathcal{X} \times \mathcal{Y})$ . Binary classification is, of course, a special case when  $\mathcal{Y} = \{\pm 1\}$  and  $\bar{\mathcal{H}} \subseteq \{\pm 1\}^\mathcal{X}$ . In that case, we simply use  $\mathcal{V}_n^{\text{Binary}}(\mathcal{H})$  for  $\mathcal{V}_n^S(\mathcal{H})$ .



**Proposition 42.** For the supervised learning game played with a target hypothesis class  $\mathcal{H} \subseteq [-1, 1]^{\mathcal{X}}$ , for any  $n \geq 2$

$$\begin{aligned} \frac{1}{4\sqrt{2}} \sup_{\alpha} \left\{ \alpha \sqrt{n \min \{ \text{fat}_{\alpha}^{\text{seq}}(\mathcal{H}), n \}} \right\} &\leq \frac{1}{2} \mathcal{V}_n^S(\mathcal{H}) \leq \mathfrak{D}_n^{\text{seq}}(\mathcal{H}) \leq \inf_{\alpha} \left\{ 4\alpha + \frac{12}{\sqrt{n}} \int_{\alpha}^1 \sqrt{\text{fat}_{\beta}^{\text{seq}}(\mathcal{H}) \log \left( \frac{2en}{\beta} \right)} d\beta \right\} \\ &\leq 58 \log^{\frac{3}{2}} n \mathfrak{R}_n^{\text{seq}}(\mathcal{H}). \end{aligned} \quad (4.6.1)$$

Moreover, the lower bound  $\mathfrak{R}_n^{\text{seq}}(\mathcal{H}) \leq \mathcal{V}_n^S(\mathcal{H})$  on the value of the supervised game also holds.

The proposition above implies that finiteness of the fat-shattering dimension is necessary and sufficient for learnability of a supervised game. Further, all the complexity notions introduced so far are within a logarithmic factor from each other whenever the problem is learnable. These results are summarized in the next theorem.

**Theorem 43.** For any target hypothesis class  $\mathcal{H} \subseteq [-1, 1]^{\mathcal{X}}$ , the following statements are equivalent

1. Target hypothesis class  $\mathcal{H}$  is online learnable in the supervised setting.
2. For any  $\alpha > 0$ ,  $\text{fat}_{\alpha}^{\text{seq}}(\mathcal{H})$  is finite.

Moreover, if the function class is online learnable, then the value of the supervised game  $\mathcal{V}_n^S(\mathcal{H})$ , the Sequential Rademacher complexity  $\mathfrak{R}_n^{\text{seq}}(\mathcal{H})$ , and the Integrated complexity  $\mathfrak{D}_n^{\text{seq}}(\mathcal{H})$  are within a multiplicative factor of  $\mathcal{O}(\log^{3/2} n)$  of each other.

*Proof.* The equivalence of 1 and 2 follows directly from Proposition 42. As for relating the various complexity measures, note that again by Proposition 42,  $\mathfrak{R}_n^{\text{seq}}(\mathcal{H}) \leq \mathcal{V}_n^S(\mathcal{H})$ ,  $\mathcal{V}_n^S(\mathcal{H}) \leq 2\mathfrak{D}_n^{\text{seq}}(\mathcal{H})$  and  $\mathfrak{D}_n^{\text{seq}}(\mathcal{H}) \leq 58 \log^{\frac{3}{2}} n \mathfrak{R}_n^{\text{seq}}(\mathcal{H})$ . Hence  $\mathcal{V}_n^S(\mathcal{H})$  and  $\mathfrak{D}_n^{\text{seq}}(\mathcal{H})$  are sandwiched between  $\mathfrak{R}_n^{\text{seq}}(\mathcal{H})$  and  $O(\log^{3/2} n) \mathfrak{R}_n^{\text{seq}}(\mathcal{H})$  which concludes the proof.  $\square$

**Corollary 44.** For the binary classification game played with function class  $\mathcal{F}$  we have that

$$K_1 \sqrt{n \min \{ \text{Ldim}(\mathcal{H}), n \}} \leq \mathcal{V}_n^{\text{Binary}}(\mathcal{H}) \leq K_2 \sqrt{n \text{Ldim}(\mathcal{H}) \log n}$$

for some universal constants  $K_1, K_2$ .

We wish to point out that the lower bound of Proposition 42 holds for arbitrary class  $\bar{\mathcal{H}}$  that is a superset of the target hypothesis class  $\mathcal{H}$ . Since a proper learning rule can always be seen also as an improper learning rule, we trivially have that if class is properly online learnable in the supervised setting then it is improperly online learnable too. However by the above mentioned fact that the lower bound of Proposition 42 holds for arbitrary class  $\bar{\mathcal{H}}$ , we also have the non-trivial reverse implication that : if a class is improperly online learnable in the supervised setting, it is online learnable.

It is natural to ask whether being able to learn in the online model is different from learning in a batch model (in the supervised setting). The standard example (e.g. [34, 35]) is the class of step functions on a bounded interval, which has a VC dimension 1, but is not learnable in the online setting. Indeed, it is possible to verify that the Littlestone’s dimension is not bounded. Interestingly, the closely-related class of “ramp” functions (modified step functions with a Lipschitz transition between 0’s and 1’s) is learnable in the online setting (and in the batch case). We extend this example as follows. By taking a convex hull of step-up and step-down functions on a unit interval, we arrive at a class of functions of bounded variation, which is learnable in the batch model, but not in the online learning model. However, the class of *Lipschitz* functions of bounded variation is learnable in both models. Online learnability of the latter class is shown with techniques analogous to Section 4.7.5.

#### 4.6.1 Generic Algorithm for Supervised Learning Problem

We shall now present a generic improper learning algorithm for the online supervised setting that achieves a low regret bound whenever the function class is online learnable. For any  $\alpha > 0$  define an  $\alpha$ -discretization of the  $[-1, 1]$  interval as  $B_\alpha = \{-1 + \alpha/2, -1 + 3\alpha/2, \dots, -1 + (2k+1)\alpha/2, \dots\}$  for  $0 \leq k$  and  $(2k+1)\alpha \leq 4$ . Also for any  $a \in [-1, 1]$  define  $\lfloor a \rfloor_\alpha = \operatorname{argmin}_{r \in B_\alpha} |r - a|$ . For a set of functions  $V \subseteq \mathcal{H}$ , any  $r \in B_\alpha$  and  $x \in \mathcal{X}$  define

$$V(r, x) = \{\mathbf{h} \in V \mid \mathbf{h}(x) \in (r - \alpha/2, r + \alpha/2]\}$$

---

##### Algorithm 1 Fat-SOA Algorithm ( $\mathcal{F}, \alpha$ )

---

```

 $V_1 \leftarrow \mathcal{F}$ 
for  $t = 1$  to  $n$  do
   $R_t(x) = \{r \in B_\alpha : \mathbf{fat}_\alpha^{\text{seq}}(V_t(r, x)) = \max_{r' \in B_\alpha} \mathbf{fat}_\alpha^{\text{seq}}(V_t(r', x))\}$ 
  For each  $x \in \mathcal{X}$ , let  $\mathbf{h}_t(x) = \frac{1}{|R_t(x)|} \sum_{r \in R_t(x)} r$ 
  Play  $\mathbf{h}_t$  and receive  $(x_t, y_t)$ 
  if  $|\mathbf{h}_t(x_t) - y_t| \leq \alpha$  then
     $V_{t+1} = V_t$ 
  else
     $V_{t+1} = V_t(\lfloor y_t \rfloor_\alpha, x_t)$ 
  end if
end for

```

---

**Lemma 45.** *Let  $\mathcal{H} \subseteq [-1, 1]^\mathcal{X}$  be a function class with finite  $\mathbf{fat}_\alpha^{\text{seq}}(\mathcal{H})$ . Suppose the learner is presented with a sequence  $(x_1, y_1), \dots, (x_n, y_n)$ , where  $y_t = \mathbf{h}(x_t)$  for some fixed  $\mathbf{h} \in \mathcal{H}$  unknown to the player. Then for  $\mathbf{h}_t$ ’s computed by the Algorithm 1 it must hold that*

$$\sum_{t=1}^T \mathbf{1}_{\{|\mathbf{h}_t(x_t) - y_t| > \alpha\}} \leq \mathbf{fat}_\alpha^{\text{seq}}(\mathcal{H}).$$

Lemma 45 proves a bound on the performance of Algorithm 1 in the realizable setting. We now provide an

algorithm for the agnostic setting. We achieve this by generating “experts” in a way similar to [35]. Using these experts along with the exponentially weighted average (EWA) algorithm we shall provide the generic algorithm for online supervised learning. The EWA (Algorithm 3) and its regret bound are provided in the appendix for completeness (p. 101).

---

**Algorithm 2** Expert  $(\mathcal{F}, \alpha, 1 \leq i_1 < \dots < i_L \leq n, Y_1, \dots, Y_L)$

---

$V_1 \leftarrow \mathcal{F}$   
**for**  $t = 1$  to  $n$  **do**  
 $R_t(x) = \{r \in B_\alpha : \text{fat}_\alpha^{\text{seq}}(V_t(r, x)) = \max_{r' \in B_\alpha} \text{fat}_\alpha^{\text{seq}}(V_t(r', x))\}$   
For each  $x \in \mathcal{X}$ , let  $f'_t(x) = \frac{1}{|R_t(x)|} \sum_{r \in R_t(x)} r$   
**if**  $t \in \{i_1, \dots, i_L\}$  **then**  
 $\forall x \in \mathcal{X}, \mathbf{h}_t(x) = Y_j$  where  $j$  is s.t.  $t = i_j$   
Play  $\mathbf{h}_t$  and receive  $x_t$   
 $V_{t+1} = V_t(\mathbf{h}_t(x_t), x_t)$   
**else**  
Play  $\mathbf{h}_t = \mathbf{h}'_t$  and receive  $x_t$   
 $V_{t+1} = V_t$   
**end if**  
**end for**

---

For each  $L \leq \text{fat}_\alpha^{\text{seq}}(\mathcal{H})$  and every possible choice of  $1 \leq i_1 < \dots < i_L \leq n$  and  $Y_1, \dots, Y_L \in B_\alpha$  we generate an expert. Denote this set of experts as  $E_n$ . Each expert outputs a function  $\mathbf{h}_t \in \mathcal{H}$  at every round  $n$ . Hence each expert  $e \in E_n$  can be seen as a sequence  $(e_1, \dots, e_n)$  of mappings  $e_t : \mathcal{X}^{t-1} \mapsto \mathcal{H}$ . The total number of unique experts is clearly

$$|E_n| = \sum_{L=0}^{\text{fat}_\alpha^{\text{seq}}} \binom{n}{L} (|B_\alpha| - 1)^L \leq \left(\frac{2n}{\alpha}\right)^{\text{fat}_\alpha^{\text{seq}}}$$

**Lemma 46.** For any  $\mathbf{h} \in \mathcal{H}$  there exists an expert  $e \in E_n$  such that for any  $t \in [n]$ ,

$$|\mathbf{h}(x_t) - e(x_{1:t-1})(x_t)| \leq \alpha$$

*Proof.* By Lemma 45, for any function  $\mathbf{h} \in \mathcal{H}$ , the number of rounds on which  $|\mathbf{h}_t(x_t) - \mathbf{h}(x_t)| > \alpha$  for the output of the fat-SOA algorithm  $\mathbf{h}_t$  is bounded by  $\text{fat}_\alpha^{\text{seq}}(\mathcal{H})$ . Further on each such round there are  $|B_\alpha| - 1$  other possibilities. For any possible such sequence of “mistakes”, there is an expert that predicts the right label on those time steps and on the remaining time agrees with the fat-SOA algorithm for that target function. Hence we see that there is always an expert  $e \in E_n$  such that

$$|\mathbf{h}(x_t) - e(x_{1:t-1})(x_t)| \leq \alpha$$

□

**Theorem 47.** For any  $\alpha > 0$  if we run the exponentially weighted average (EWA) algorithm with the set  $E_n$

of experts then the expected regret of the algorithm is bounded as

$$\mathbb{E} [\mathbf{R}_n] \leq \alpha + \sqrt{\frac{\text{fat}_\alpha^{\text{seq}} \log\left(\frac{2n}{\alpha}\right)}{n}}$$

*Proof.* For any  $\alpha \geq 0$  if we run EWA with corresponding set of experts  $E_n$  then we can guarantee that regret w.r.t. best expert in the set  $E_n$  is bounded by  $\sqrt{n \text{fat}_\alpha^{\text{seq}} \log\left(\frac{2n}{\alpha}\right)}$ . However by Lemma 46 we have that the regret of the best expert in  $E_n$  w.r.t. best function in function class  $\mathcal{H}$  is at most  $\alpha n$ . Combining we get the required result.  $\square$

The above theorem holds for a fixed  $\alpha$ . To provide a regret statement that optimizes over  $\alpha$  we consider  $\alpha_i$ 's of form  $2^{-i}$  and assign weights  $p_i = \frac{6}{\pi^2} i^{-2}$  to experts generated in above theorem for each  $\alpha_i$  and run EWA on the entire set of experts with these initial weights. Hence we get the following corollary.

**Corollary 48.** *Let  $\mathcal{H} \subseteq [-1, 1]^{\mathcal{X}}$ . The expected regret of the algorithm described above is bounded as*

$$\mathbb{E} [\mathbf{R}_n] \leq \inf_{\alpha} \left\{ \alpha + \sqrt{\frac{\text{fat}_\alpha^{\text{seq}} \log\left(\frac{2n}{\alpha}\right)}{n}} + \frac{(3 + 2 \log \log\left(\frac{1}{\alpha}\right))}{\sqrt{n}} \right\}$$

## 4.7 Examples

### 4.7.1 Example: Margin Based Regret

In the classical statistical setting, margin bounds provide guarantees on expected zero-one loss of a classifier based on the empirical margin zero-one error. These results form the basis of the theory of large margin classifiers (see [53, 54]). Recently, in the online setting, margin bounds have been shown through the concept of margin via the Littlestone dimension [35]. We show that our machinery can easily lead to margin bounds for the binary classification games for general function classes  $\mathcal{F}$  based on their sequential Rademacher Complexity. We use ideas from [54] to do this.

**Proposition 49.** *For any target hypothesis class  $\mathcal{H} \subset [-1, 1]^{\mathcal{Z}}$ , there exists a randomized online algorithm  $\mathbf{A}$  such that for any sequence  $z_1, \dots, z_n$  where each  $z_t = (x_t, y_t) \in \mathcal{Z} \times \{\pm 1\}$ , played by the adversary,*

$$\mathbb{E} \left[ \frac{1}{n} \sum_{t=1}^n \mathbb{E}_{\mathbf{h}_t \sim \mathbf{A}(z_{1:t-1})} [\mathbf{1}_{\{\mathbf{h}_t(x_t) y_t < 0\}}] \right] \leq \inf_{\gamma > 0} \left\{ \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \mathbf{1}_{\{\mathbf{h}(x_t) y_t < \gamma\}} + \frac{4}{\gamma} \mathfrak{R}_n^{\text{seq}}(\mathcal{H}) + \frac{(3 + \log \log\left(\frac{1}{\gamma}\right))}{\sqrt{n}} \right\}$$

### 4.7.2 Example : Neural Networks

We provide below a bound on sequential Rademacher complexity for classic multi-layer neural networks thus showing they are learnable in the online setting. The model of neural network we consider below and the

bounds we provide are analogous to the ones considered in the batch setting in [38]. We now consider a  $k$ -layer 1-norm neural network. To this end let function class  $\mathcal{F}_1$  be given by

$$\mathcal{F}_1 = \left\{ x \mapsto \sum_j w_j^1 x_j \mid \|w\|_1 \leq B_1 \right\}$$

and further for each  $2 \leq i \leq k$  define

$$\mathcal{F}_i = \left\{ x \mapsto \sum_j w_j^i \sigma(f_j(x)) \mid \forall j f_j \in \mathcal{F}_{i-1}, \|w^i\|_1 \leq B_i \right\}$$

**Proposition 50.** *Say  $\sigma : \mathbb{R} \mapsto [-1, 1]$  is  $L$ -Lipschitz, then*

$$\mathfrak{N}_n^{\text{seq}}(\mathcal{F}_k) \leq \left( \prod_{i=1}^k 2B_i \right) L^{k-1} X_\infty \sqrt{\frac{2 \log d}{n}}$$

where  $X_\infty$  is such that  $\forall x \in \mathcal{Z}, \|x\|_\infty \leq X_\infty$  and  $\mathcal{Z} \subset \mathbb{R}^d$

### 4.7.3 Example: Decision Trees

We consider here the supervised learning game where adversary provides instances from input space  $\mathcal{X}$  and binary labels  $\pm 1$  corresponding to the input and the player plays decision trees of depth no more than  $d$  with decision functions from set  $\mathcal{C} \subset \{\pm 1\}^{\mathcal{X}}$  of binary valued functions. The following proposition shows that there exists a randomized learning algorithm which under certain circumstances could have low regret for the supervised learning (binary) game played with class of decision trees of depth at most  $d$  with decision functions from  $\mathcal{C}$ . The proposition is analogous to the one in [38] considered in the batch (classical) setting.

**Proposition 51.** *Denote by  $\mathcal{T}$  the class of decision trees of depth at most  $d$  with decision functions in  $\mathcal{C}$ . There exists a randomized online learning algorithm  $\mathbf{A}$  such that for any sequence of instances  $z_1 = (x_1, y_1), \dots, z_n = (x_n, y_n) \in (\mathcal{X} \times \{\pm 1\})^n$  played by the adversary,*

$$\mathbb{E} \left[ \frac{1}{n} \sum_{t=1}^n \mathbb{E}_{\tau_t \sim \mathbf{A}(z_{1:t-1})} [\mathbf{1}_{\{\tau_t(x_t) \neq y_t\}}] \right] \leq \inf_{\tau \in \mathcal{T}} \frac{1}{n} \sum_{t=1}^n \mathbf{1}_{\{\tau(x_t) \neq y_t\}} + \mathcal{O} \left( \sum_l \min \left( \tilde{C}_n(l), d \log^{3/2}(n) \mathfrak{N}_n^{\text{seq}}(\mathcal{H}) \right) + \frac{(3 + 2 \log(N_{\text{leaf}}))}{\sqrt{n}} \right)$$

where  $\tilde{C}_n(l)$  denotes the number of instances which reach the leaf  $l$  and are correctly classified in the decision tree  $t$  that minimizes  $\sum_{t=1}^n \mathbf{1}_{\{\tau(x_t) \neq y_t\}}$  and let  $N_{\text{leaf}}$  be the number of leaves in this tree.

#### 4.7.4 Example: Online Transductive Learning

Let  $\mathcal{F}$  be a class of functions from  $\mathcal{Z}$  to  $\mathbb{R}$ . Let

$$\widehat{\mathcal{N}}_\infty(\alpha, \mathcal{F}) = \min \{ |G| : G \subseteq \mathbb{R}^{\mathcal{Z}} \text{ s.t. } \forall f \in \mathcal{F} \exists g \in G \text{ satisfying } \|f - g\|_\infty \leq \alpha \}. \quad (4.7.1)$$

be the  $\ell_\infty$  covering number at scale  $\alpha$ , where the cover is pointwise on all of  $\mathcal{Z}$ . It is easy to verify that

$$\forall n, \quad N_\infty(\alpha, \mathcal{F}, n) \leq \widehat{\mathcal{N}}_\infty(\alpha, \mathcal{F}) \quad (4.7.2)$$

Indeed, let  $G$  be a minimal cover of  $\mathcal{F}$  at scale  $\alpha$ . We claim that the set  $V = \{\mathbf{v}^g = g(\mathbf{z}) : g \in G\}$  of  $\mathbb{R}$ -valued trees is an  $\ell_\infty$  cover of  $\mathcal{F}$  on  $\mathbf{z}$ . Fix any  $\epsilon \in \{\pm 1\}^n$  and  $f \in \mathcal{F}$ , and let  $g \in G$  be such that  $\|f - g\|_\infty \leq \alpha$ . Then clearly  $|\mathbf{v}_t^g(\epsilon) - f(\mathbf{z}_t(\epsilon))|$  for any  $1 \leq t \leq n$ , which concludes the proof.

This simple observation can be applied in several situations. First, consider the problem of *transductive learning*, where the set  $\mathcal{Z} = \{z_1, \dots, z_m\}$  is a finite set. To ensure online learnability, it is sufficient to consider an assumption on the dependence of  $\widehat{\mathcal{N}}_\infty(\alpha, \mathcal{F})$  on  $\alpha$ . An obvious example of such a class is a VC-type class with  $\widehat{\mathcal{N}}_\infty(\alpha, \mathcal{F}) \leq (c/\alpha)^d$  for some  $c$  which can depend on  $m$ . Assume that  $\mathcal{F} \subset [-1, 1]^{\mathcal{Z}}$ . Substituting this bound on the covering number into

$$\mathfrak{D}_n^{\text{seq}}(\mathcal{F}) = \inf_\alpha \left\{ 4\alpha + 12 \int_\alpha^1 \sqrt{\frac{\log \mathcal{N}_2(\delta, \mathcal{F}, n)}{n}} d\delta \right\}$$

and choosing  $\alpha = 0$ , we observe that the value of the supervised game is upper bounded by  $2\mathfrak{D}_n^{\text{seq}}(\mathcal{F}) \leq 4\sqrt{\frac{d \log c}{n}}$  by Proposition 42. It is easy to see that if  $m$  is fixed and the problem is learnable in the batch (e.g. PAC) setting, then the problem is learnable in the online transductive model.

In the transductive setting considered by Kakade and Kalai [55], it is assumed that  $m \leq n$  and  $\mathcal{F}$  are binary-valued. If  $\mathcal{F}$  is a class with VC dimension  $d$ , the Sauer-Shelah lemma ensures that the  $\ell_\infty$  cover is smaller than  $(em/d)^d \leq (en/d)^d$ . Using the previous argument with  $c = en$ , we obtain a bound of  $4\sqrt{dn \log(en)}$  for the value of the game, matching [55] up to a constant 2.

We also consider the problem of prediction of individual sequences, which has been studied both in information theory and in learning theory. In particular, in the case of binary prediction, Cesa-Bianchi and Lugosi [56] proved upper bounds on the value of the game in terms of the (classical) Rademacher complexity and the (classical) Dudley integral. The particular assumption made in [56] is that experts are *static*. That is, their prediction only depends on the current round, not on the past information. Formally, we define static experts as mappings  $f : \{1, \dots, n\} \mapsto \mathcal{Y} = [-1, 1]$ , and let  $\mathcal{F}$  denote a class of such experts. Defining  $\mathcal{Z} = \{1, \dots, n\}$  puts us in the setting considered earlier with  $m = n$ . We immediately obtain  $4\sqrt{dn \log(en)}$ , matching the results on [56, p. 1873]. We mention that the upper bound in Theorem 4 in [56] is tighter by a  $\log n$  factor if a sharper bound on the  $\ell_2$  cover is considered. Finally, for the case of a finite number of experts, clearly  $\widehat{\mathcal{N}}_\infty \leq N$  which gives the classical  $O(\sqrt{n \log N})$  bound on the value of the game [52].

### 4.7.5 Example: Isotron

Recently, Kalai and Sastry [57] introduced a method called *Isotron* for learning Single Index Models (SIM). These models generalize linear and logistic regression, generalized linear models, and classification by linear threshold functions. For brevity, we only describe the Idealized SIM problem from [57]. In its “batch” version, we assume that the data is revealed at once as a set  $\{(\mathbf{x}_i, y_i)\}_{i=1}^n \in \mathbb{R}^d \times \mathbb{R}$  where  $y_i = u(\langle \mathbf{w}, \mathbf{x}_i \rangle)$  for some unknown  $\mathbf{w} \in \mathbb{R}^d$  of bounded norm and an unknown non-decreasing  $u : \mathbb{R} \mapsto \mathbb{R}$  with a bounded Lipschitz constant. Given this data, the goal is to iteratively find the function  $u$  and the direction  $\mathbf{w}$ , making as few mistakes as possible. The error is measured as  $\frac{1}{n} \sum_{i=1}^n (\mathbf{h}_i(\mathbf{x}_i) - y_i)^2$ , where  $\mathbf{h}_i(\mathbf{x}) = u_i(\langle \mathbf{w}_i, \mathbf{x} \rangle)$  is the iterative approximation found by the algorithm on the  $i$ th round. The elegant computationally efficient method presented in [57] is motivated by Perceptron, and a natural open question posed by the authors is whether there is an online variant of Isotron. Before even attempting a quest for such an algorithm, we can ask a more basic question: is the (Idealized) SIM problem even learnable in the online framework? After all, most online methods deal with convex functions, but  $u$  is only assumed to be Lipschitz and non-decreasing. We answer the question easily with the tools we have developed.

We are interested in online learnability of the above described problem. Specifically, the setting of the problem is a supervised learning one where instance set is  $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$  where  $\mathcal{X} = B_2$  (the unit Euclidean ball in  $\mathbb{R}^d$ ) and  $\mathcal{Y} = [-1, 1]$ . The target hypothesis set is given by  $\mathcal{H} = \mathcal{U} \times B_2$  where  $\mathcal{U} = \{u : [-1, 1] \mapsto [-1, 1] : u \text{ is 1-Lipschitz}\}$ . Finally the loss function is the squared loss, that is given hypothesis  $\mathbf{h} = (u, \mathbf{x})$  and instance  $z = (\mathbf{x}, y)$ , the loss function is given by  $\ell(\mathbf{h}, z) = (u(\langle \mathbf{x}, \mathbf{w} \rangle) - y)^2$ . It is evident that the loss class  $\mathcal{F}$  associated with the problem is a composition with three levels: the squared loss, the Lipschitz non-decreasing function, and the linear function. That is

$$\mathcal{F} = \{(\mathbf{x}, y) \mapsto (y - u(\langle \mathbf{w}, \mathbf{x} \rangle))^2 \mid u \in \mathcal{U}, \|\mathbf{w}\|_2 \leq 1\} \quad (4.7.3)$$

The proof of the following Proposition boils down to showing that the covering number of the class does not increase much under these compositions.

**Proposition 52.** *The target class  $\mathcal{H} = \mathcal{U} \times B_2$  is online learnable in the supervised setting. Moreover,*

$$\mathcal{V}_n(\mathcal{H}, \mathcal{X} \times \mathcal{Y}) = O\left(\sqrt{\frac{\log^3 n}{n}}\right).$$

## 4.8 Detailed Proofs and More Results

### 4.8.1 Proofs

**Proof of Theorem 23.** For brevity, denote  $\psi(z_{1:n}) = \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t)$ . The first step in the proof is to appeal to the minimax theorem for every couple of inf and sup:

$$\begin{aligned} \mathcal{V}_n(\mathcal{H}, \mathcal{Z}) &= \inf_{q_1 \in \Delta(\mathcal{H})} \sup_{p_1 \in \Delta(\mathcal{Z})} \mathbb{E}_{\substack{\mathbf{h}_1 \sim q_1 \\ z_1 \sim p_1}} \dots \inf_{q_n \in \Delta(\mathcal{H})} \sup_{p_n \in \Delta(\mathcal{Z})} \mathbb{E}_{\substack{\mathbf{h}_n \sim q_n \\ z_n \sim p_n}} \left\{ \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \psi(z_{1:n}) \right\} \\ &= \sup_{p_1 \in \Delta(\mathcal{Z})} \inf_{q_1 \in \Delta(\mathcal{H})} \mathbb{E}_{\substack{\mathbf{h}_1 \sim q_1 \\ z_1 \sim p_1}} \dots \sup_{p_n \in \Delta(\mathcal{Z})} \inf_{q_n \in \Delta(\mathcal{H})} \mathbb{E}_{\substack{\mathbf{h}_n \sim q_n \\ z_n \sim p_n}} \left\{ \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \psi(z_{1:n}) \right\} \\ &= \sup_{p_1 \in \Delta(\mathcal{Z})} \inf_{\mathbf{h}_1 \in \mathcal{H}} \mathbb{E}_{x_1 \sim p_1} \dots \sup_{p_n \in \Delta(\mathcal{Z})} \inf_{\mathbf{h}_n \in \mathcal{H}} \mathbb{E}_{z_n \sim p_n} \left\{ \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \psi(z_{1:n}) \right\} \end{aligned}$$

From now on, it will be understood that  $z_t$  has distribution  $p_t$ . Moving the expectation with respect to  $z_n$  and then the infimum with respect to  $\mathbf{h}_n$  inside the expression, we arrive at,

$$\begin{aligned} &= \sup_{p_1 \in \Delta(\mathcal{Z})} \inf_{\mathbf{h}_1 \in \mathcal{H}} \mathbb{E}_{z_1} \dots \sup_{p_{n-1} \in \Delta(\mathcal{Z})} \inf_{\mathbf{h}_{n-1} \in \mathcal{H}} \mathbb{E}_{z_{n-1}} \sup_{p_n \in \Delta(\mathcal{Z})} \left\{ \frac{1}{n} \sum_{t=1}^{n-1} \ell(\mathbf{h}_t, z_t) + \frac{1}{n} \inf_{\mathbf{h}_n \in \mathcal{H}} \mathbb{E}_{z_n} \ell(\mathbf{h}_n, z_n) - \mathbb{E}_{z_n} \psi(z_{1:n}) \right\} \\ &= \sup_{p_1 \in \Delta(\mathcal{Z})} \inf_{\mathbf{h}_1 \in \mathcal{H}} \mathbb{E}_{z_1} \dots \sup_{p_{n-1} \in \Delta(\mathcal{Z})} \inf_{\mathbf{h}_{n-1} \in \mathcal{H}} \mathbb{E}_{z_{n-1}} \sup_{p_n \in \Delta(\mathcal{Z})} \mathbb{E}_{z_n} \left[ \frac{1}{n} \sum_{t=1}^{n-1} \ell(\mathbf{h}_t, z_t) + \frac{1}{n} \inf_{\mathbf{h}_n \in \mathcal{H}} \mathbb{E}_{z_n} \ell(\mathbf{h}_n, z_n) - \psi(z_{1:n}) \right] \end{aligned}$$

Repeating the procedure for step  $n-1$ ,

$$\begin{aligned} &= \sup_{p_1 \in \Delta(\mathcal{Z})} \inf_{\mathbf{h}_1 \in \mathcal{H}} \mathbb{E}_{z_1} \dots \sup_{p_{n-1} \in \Delta(\mathcal{Z})} \inf_{\mathbf{h}_{n-1} \in \mathcal{H}} \mathbb{E}_{z_{n-1}} \left[ \frac{1}{n} \sum_{t=1}^{n-1} \ell(\mathbf{h}_t, z_t) + \sup_{p_n \in \Delta(\mathcal{Z})} \mathbb{E}_{z_n} \left[ \frac{1}{n} \inf_{\mathbf{h}_n \in \mathcal{H}} \mathbb{E}_{z_n} \ell(\mathbf{h}_n, z_n) - \psi(z_{1:n}) \right] \right] \\ &= \sup_{p_1 \in \Delta(\mathcal{Z})} \inf_{\mathbf{h}_1 \in \mathcal{H}} \mathbb{E}_{z_1} \dots \sup_{p_{n-1} \in \Delta(\mathcal{Z})} \left\{ \frac{1}{n} \sum_{t=1}^{n-2} \ell(\mathbf{h}_t, z_t) + \frac{1}{n} \left[ \inf_{\mathbf{h}_{n-1} \in \mathcal{H}} \mathbb{E}_{z_{n-1}} \ell(\mathbf{h}_{n-1}, z_{n-1}) \right. \right. \\ &\quad \left. \left. + \mathbb{E}_{z_{n-1}} \sup_{p_n \in \Delta(\mathcal{Z})} \mathbb{E}_{z_n} \left[ \frac{1}{n} \inf_{\mathbf{h}_n \in \mathcal{H}} \mathbb{E}_{z_n} \ell(\mathbf{h}_n, z_n) - \psi(z_{1:n}) \right] \right] \right\} \\ &= \sup_{p_1 \in \Delta(\mathcal{Z})} \inf_{\mathbf{h}_1 \in \mathcal{H}} \mathbb{E}_{z_1} \dots \sup_{p_n \in \Delta(\mathcal{Z})} \mathbb{E}_{z_n} \left\{ \frac{1}{n} \sum_{t=1}^{n-2} \ell(\mathbf{h}_t, z_t) + \frac{1}{n} \sum_{t=n-1}^n \inf_{\mathbf{h}_t \in \mathcal{H}} \mathbb{E}_{z_t} \ell(\mathbf{h}_t, z_t) - \psi(z_{1:n}) \right\} \end{aligned}$$

Continuing in this fashion for  $n-2$  and all the way down to  $t=1$  proves the theorem.  $\square$

**Proof of the Key Technical Lemma (Lemma 25).** We start by noting that since  $z_n, z'_n$  are both drawn from



$p_n$ ,

$$\begin{aligned}
\mathbb{E}_{z_n, z'_n \sim p_n} \left[ \Phi \left( \sum_{t=1}^n \Delta_{\mathbf{h}}(z_t, z'_t) \right) \right] &= \mathbb{E}_{z_n, z'_n \sim p_n} \left[ \Phi \left( \sum_{t=1}^{n-1} \Delta_{\mathbf{h}}(z_t, z'_t) + \Delta_{\mathbf{h}}(z_n, z'_n) \right) \right] \\
&= \mathbb{E}_{z'_n, z_n \sim p_n} \left[ \Phi \left( \sum_{t=1}^{n-1} \Delta_{\mathbf{h}}(z_t, z'_t) + \Delta_{\mathbf{h}}(z_n, z'_n) \right) \right] \\
&= \mathbb{E}_{z_n, z'_n \sim p_n} \left[ \Phi \left( \sum_{t=1}^{n-1} \Delta_{\mathbf{h}}(z_t, z'_t) - \Delta_{\mathbf{h}}(z_n, z'_n) \right) \right],
\end{aligned}$$

where the last line is by antisymmetry of  $\Delta_{\mathbf{h}}$ . Since the first and last lines are equal, they are both equal to their average and hence

$$\mathbb{E}_{z_n, z'_n \sim p_n} \left[ \Phi \left( \sum_{t=1}^{n-1} \Delta_{\mathbf{h}}(z_t, z'_t) \right) \right] = \mathbb{E}_{z_n, z'_n \sim p_n} \left[ \mathbb{E}_{\epsilon_n} \left[ \Phi \left( \sum_{t=1}^{n-1} \Delta_{\mathbf{h}}(z_t, z'_t) + \epsilon_n \Delta_{\mathbf{h}}(z_n, z'_n) \right) \right] \right].$$

Hence we conclude that

$$\begin{aligned}
&\sup_{p_n} \mathbb{E}_{z_n, z'_n \sim p_n} \left[ \Phi \left( \sum_{t=1}^n \Delta_{\mathbf{h}}(z_t, z'_t) \right) \right] \\
&= \sup_{p_n} \mathbb{E}_{z_n, z'_n \sim p_n} \left[ \mathbb{E}_{\epsilon_n} \left[ \Phi \left( \sum_{t=1}^{n-1} \Delta_{\mathbf{h}}(z_t, z'_t) + \epsilon_n \Delta_{\mathbf{h}}(z_n, z'_n) \right) \right] \right] \\
&\leq \sup_{z_n, z'_n} \mathbb{E}_{\epsilon_n} \left[ \Phi \left( \sum_{t=1}^{n-1} \Delta_{\mathbf{h}}(z_t, z'_t) + \epsilon_n \Delta_{\mathbf{h}}(z_n, z'_n) \right) \right].
\end{aligned}$$

Using the above and noting that  $z_{n-1}, z'_{n-1}$  are both drawn from  $p_{n-1}$  and hence similar to previous step introducing Rademacher variable  $\epsilon_{n-1}$  we get that

$$\begin{aligned}
&\sup_{p_{n-1}} \mathbb{E}_{z_{n-1}, z'_{n-1} \sim p_{n-1}} \sup_{p_n} \mathbb{E}_{z_n, z'_n \sim p_n} \left[ \Phi \left( \sum_{t=1}^n \Delta_{\mathbf{h}}(z_t, z'_t) \right) \right] \\
&\leq \sup_{p_{n-1}} \mathbb{E}_{z_{n-1}, z'_{n-1} \sim p_{n-1}} \left[ \sup_{z_n, z'_n} \mathbb{E}_{\epsilon_n} \left[ \Phi \left( \sum_{t=1}^{n-1} \Delta_{\mathbf{h}}(z_t, z'_t) + \epsilon_n \Delta_{\mathbf{h}}(z_n, z'_n) \right) \right] \right] \\
&= \sup_{p_{n-1}} \mathbb{E}_{z_{n-1}, z'_{n-1} \sim p_{n-1}} \mathbb{E}_{\epsilon_{n-1}} \left[ \sup_{z_n, z'_n} \mathbb{E}_{\epsilon_n} \left[ \Phi \left( \sum_{t=1}^{n-2} \Delta_{\mathbf{h}}(z_t, z'_t) + \epsilon_{n-1} \Delta_{\mathbf{h}}(z_{n-1}, z'_{n-1}) + \epsilon_n \Delta_{\mathbf{h}}(z_n, z'_n) \right) \right] \right] \\
&\leq \sup_{z_{n-1}, z'_{n-1}} \mathbb{E}_{\epsilon_{n-1}} \left[ \sup_{z_n, z'_n} \mathbb{E}_{\epsilon_n} \left[ \Phi \left( \sum_{t=1}^{n-2} \Delta_{\mathbf{h}}(z_t, z'_t) + \epsilon_{n-1} \Delta_{\mathbf{h}}(z_{n-1}, z'_{n-1}) + \epsilon_n \Delta_{\mathbf{h}}(z_n, z'_n) \right) \right] \right].
\end{aligned}$$

Proceeding in similar fashion introducing Rademacher variables all the way upto  $\epsilon_1$  we finally get the required

statement that

$$\begin{aligned} & \sup_{p_1} \mathbb{E}_{z_1, z'_1 \sim p_1} \cdots \sup_{p_n} \mathbb{E}_{z_n, z'_n \sim p_n} \left[ \Phi \left( \sum_{t=1}^n \Delta_{\mathbf{h}}(z_t, z'_t) \right) \right] \\ & \leq \sup_{z_1, z'_1} \left\{ \mathbb{E}_{\epsilon_1} \left[ \cdots \sup_{z_n, z'_n} \left\{ \mathbb{E}_{\epsilon_n} \left[ \Phi \left( \sum_{t=1}^n \epsilon_t \Delta_{\mathbf{h}}(z_t, z'_t) \right) \right] \right\} \cdots \right] \right\} \end{aligned}$$

□

**Proof of Lemma 26.** Without loss of generality assume that the Lipschitz constant  $L = 1$  because the general case follows by scaling  $\phi$ . Now note that by Theorem 35 we have that

$$\mathfrak{R}_n^{\text{seq}}(\phi \circ \mathcal{F}) \leq \inf_{\alpha} \left\{ 4\alpha + 12 \int_{\alpha}^1 \sqrt{\frac{\log \mathcal{N}_2(\delta, \phi \circ \mathcal{F}, n)}{n}} d\delta \right\} \quad (4.8.1)$$

Now we claim that we can bound

$$\log \mathcal{N}_2(\delta, \phi \circ \mathcal{F}, n) \leq \sum_{j=1}^k \log \mathcal{N}_{\infty}(\delta, \mathcal{F}_j, n)$$

To see this we first start by noting that

$$\begin{aligned} & \sqrt{\frac{1}{n} \sum_{t=1}^n (\phi(f(z_t(\epsilon)), z_t(\epsilon)) - \phi(\mathbf{v}_t(\epsilon), z_t(\epsilon)))^2} \\ & \leq \sqrt{\frac{1}{n} \sum_{t=1}^n \max_j (f_j(z_t(\epsilon)) - \mathbf{v}_t^j(\epsilon))^2} \\ & \leq \sqrt{\max_{t \in [n]} \max_j (f_j(z_t(\epsilon)) - \mathbf{v}_t^j(\epsilon))^2} \\ & \leq \max_{j \in [k], t \in [n]} |f_j(z_t(\epsilon)) - \mathbf{v}_t^j(\epsilon)| \end{aligned}$$

Now suppose we have  $V_1, \dots, V_k$  that are minimal  $\ell_{\infty}$ -covers for  $\mathcal{F}_1, \dots, \mathcal{F}_k$  on the tree  $z$  at level  $\delta$ . Consider the set:

$$V_{\phi} = \{\mathbf{v}_{\phi} : \mathbf{v} \in V_1 \times \dots \times V_k\}$$

where  $\mathbf{v}_{\phi}$  is the tree such that  $(\mathbf{v}_{\phi})_t(\epsilon) = \phi(\mathbf{v}_t(\epsilon), z_t(\epsilon))$ . Then, for any  $f = (f_1, \dots, f_k) \in \mathcal{F}$  (with

representatives  $(\mathbf{v}^1, \dots, \mathbf{v}^k) \in V_1 \times \dots \times V_k$  and any  $\epsilon \in \{\pm 1\}^n$ , we have,

$$\begin{aligned} & \sqrt{\frac{1}{n} \sum_{t=1}^n (\phi(f(z_t(\epsilon)), z_t(\epsilon)) - (\mathbf{v}_\phi)_t(\epsilon))^2} \\ &= \sqrt{\frac{1}{n} \sum_{t=1}^n (\phi(f(z_t(\epsilon)), z_t(\epsilon)) - \phi(\mathbf{v}_t(\epsilon), z_t(\epsilon)))^2} \\ &\leq \max_{j \in [k]} \max_{t \in [n]} |f_j(z_t(\epsilon)) - \mathbf{v}_t^j(\epsilon)| \leq \delta \end{aligned}$$

Thus we see that  $V_\phi$  is an  $\ell_\infty$ -cover at scale  $\delta$  for  $\phi \circ \mathcal{F}$  on  $z$ . Hence

$$\begin{aligned} \log \mathcal{N}_2(\delta, \phi \circ \mathcal{F}, n) &\leq \log \mathcal{N}_\infty(\delta, \phi \circ \mathcal{F}, n) \\ &\leq \log(|V|) = \sum_{j=1}^k \log(|V_j|) \\ &= \sum_{j=1}^k \log \mathcal{N}_\infty(\delta, \mathcal{F}_j, n) \end{aligned}$$

as claimed. Now using this in Equation 4.8.1 we have that

$$\begin{aligned} \mathfrak{R}_n^{\text{seq}}(\phi \circ \mathcal{F}) &\leq \inf_\alpha \left\{ 4\alpha + 12 \int_\alpha^1 \sqrt{\frac{\sum_{j=1}^k \log \mathcal{N}_\infty(\delta, \mathcal{F}_j, n)}{n}} d\delta \right\} \\ &\leq \inf_\alpha \left\{ 4\alpha + 12 \sum_{j=1}^k \int_\alpha^1 \sqrt{\frac{\log \mathcal{N}_\infty(\delta, \mathcal{F}_j, n)}{n}} d\delta \right\} \\ &\leq \sum_{j=1}^k \inf_\alpha \left\{ 4\alpha + 12 \int_\alpha^1 \sqrt{\frac{\log \mathcal{N}_\infty(\delta, \mathcal{F}_j, n)}{n}} d\delta \right\} \end{aligned}$$

Now applying Lemma37 we conclude, as required, that

$$\mathfrak{R}_n^{\text{seq}}(\phi \circ \mathcal{F}) \leq 8 \left( 1 + 4\sqrt{2} \log^{3/2}(en^2) \right) \sum_{j=1}^k \mathfrak{R}_n^{\text{seq}}(\mathcal{F}_j)$$

as long as  $\mathfrak{R}_n^{\text{seq}}(\mathcal{F}_j) \geq 1$  for each  $j$ . □

**Proof of Corollary 28.** We first extend the binary function  $b$  to a function  $\bar{b}$  to any  $x \in \mathbb{R}^k$  as follows :

$$\bar{b}(x) = \begin{cases} (1 - \|x - a\|_\infty) b(a) & \text{if } \|x - a\|_\infty < 1 \text{ for some } a \in \{\pm 1\}^k \\ 0 & \text{otherwise} \end{cases}$$

First note that  $\bar{b}$  is well-defined since all points in the  $k$ -cube are separated by  $L_\infty$  distance 2. Further

note that  $\bar{b}$  is 1-Lipschitz w.r.t. the  $L_\infty$  norm and so applying Lemma 26 we conclude the statement of the corollary.  $\square$

**Proof of Proposition 29.** The most difficult of these is Part 4, which follows immediately by Lemma 27 by taking  $\phi(\cdot, z)$  there to be simply  $\phi(\cdot)$ . The other items follow similarly to Theorem 15 in [39] and we provide the proofs for completeness. Note that, unlike Rademacher complexity defined in [39], Sequential Rademacher complexity does not have the absolute value around the sum.

Part 1 is immediate because for any fixed tree  $z$  and fixed realization of  $\{\epsilon_i\}$ ,

$$\sup_{f \in \mathcal{F}} \sum_{t=1}^n \epsilon_t f(z_t(\epsilon)) \leq \sup_{f \in \mathcal{G}} \sum_{t=1}^n \epsilon_t f(z_t(\epsilon)),$$

Now taking expectation over  $\epsilon$  and supremum over  $z$  completes the argument.

To show Part 2, first observe that, according to Part 1,

$$\mathfrak{R}_n^{\text{seq}}(\mathcal{F}) \leq \mathfrak{R}_n^{\text{seq}}(\text{conv}(\mathcal{F})).$$

Now, any  $h \in \text{conv}(\mathcal{F})$  can be written as  $h = \sum_{j=1}^m \alpha_j f_j$  with  $\sum_{j=1}^m \alpha_j = 1$ ,  $\alpha_j \geq 0$ . Then, for fixed tree  $z$  and sequence  $\epsilon$ ,

$$\sum_{t=1}^n \epsilon_t h(z_t(\epsilon)) = \sum_{j=1}^m \alpha_j \sum_{t=1}^n \epsilon_t f_j(z_t(\epsilon)) \leq \sup_{f \in \mathcal{F}} \sum_{t=1}^n \epsilon_t f(z_t(\epsilon))$$

and thus

$$\sup_{h \in \text{conv}(\mathcal{F})} \sum_{t=1}^n \epsilon_t h(z_t(\epsilon)) \leq \sup_{f \in \mathcal{F}} \sum_{t=1}^n \epsilon_t f(z_t(\epsilon)).$$

Taking expectation over  $\epsilon$  and supremum over  $z$  completes the proof.

To prove Part 3, first observe that the statement is easily seen to hold for  $c \geq 0$ . That is,  $\mathfrak{R}_n^{\text{seq}}(c\mathcal{F}) = c\mathfrak{R}_n^{\text{seq}}(\mathcal{F})$  follows directly from the definition. Hence, it remains to convince ourselves of the statement for  $c = -1$ . That is,  $\mathfrak{R}_n^{\text{seq}}(-\mathcal{F}) = \mathfrak{R}_n^{\text{seq}}(\mathcal{F})$ . To prove this, consider a tree  $z^R$  that is a reflection of  $z$ . That is,  $z_t^R(\epsilon) = z_t(-\epsilon)$  for all  $t \in [n]$ . It is then enough to observe that

$$\begin{aligned} \mathbb{E}_\epsilon \left[ \sup_{f \in -\mathcal{F}} \sum_{t=1}^n \epsilon_t f(z_t(\epsilon)) \right] &= \mathbb{E}_\epsilon \left[ \sup_{f \in \mathcal{F}} \sum_{t=1}^n -\epsilon_t f(z_t(\epsilon)) \right] \\ &= \mathbb{E}_\epsilon \left[ \sup_{f \in \mathcal{F}} \sum_{t=1}^n \epsilon_t f(z_t(-\epsilon)) \right] = \mathbb{E}_\epsilon \left[ \sup_{f \in \mathcal{F}} \sum_{t=1}^n \epsilon_t f(z_t^R(\epsilon)) \right] \end{aligned}$$

where we used the fact that  $\epsilon$  and  $-\epsilon$  have the same distribution. As  $z$  varies over all trees,  $z^R$  also varies over all trees. Hence taking the supremum over  $z$  above finishes the argument.

Finally, for Part 5,

$$\sup_{f \in \mathcal{F}} \left\{ \sum_{t=1}^n \epsilon_t (f + h)(\mathbf{z}_t(\epsilon)) \right\} = \left\{ \sup_{f \in \mathcal{F}} \sum_{t=1}^n \epsilon_t f(\mathbf{z}_t(\epsilon)) \right\} + \left\{ \sum_{t=1}^n \epsilon_t h(\mathbf{z}_t(\epsilon)) \right\}$$

Note that, since  $h(\mathbf{z}_t(\epsilon))$  only depends on  $\epsilon_{1:t-1}$ , we have

$$\mathbb{E}_\epsilon [\epsilon_t h(\mathbf{z}_t(\epsilon))] = \mathbb{E}_{\epsilon_{1:t-1}} [\mathbb{E}[\epsilon_t | \epsilon_{1:t-1}] h(\mathbf{z}_t(\epsilon))] = 0.$$

Thus,

$$\mathfrak{R}_n^{\text{seq}}(\mathcal{F} + h) = \mathfrak{R}_n^{\text{seq}}(\mathcal{F}).$$

□

**Proof of Lemma 34.** For any  $\lambda > 0$ , we invoke Jensen's inequality to get

$$\begin{aligned} M(\lambda) &:= \exp \left\{ \lambda \mathbb{E}_\epsilon \left[ \max_{\mathbf{v} \in V} \sum_{t=1}^n \epsilon_t \mathbf{v}_t(\epsilon) \right] \right\} \leq \mathbb{E}_\epsilon \left[ \exp \left\{ \lambda \max_{\mathbf{v} \in V} \sum_{t=1}^n \epsilon_t \mathbf{v}_t(\epsilon) \right\} \right] \\ &\leq \mathbb{E}_\epsilon \left[ \max_{\mathbf{v} \in V} \exp \left\{ \lambda \sum_{t=1}^n \epsilon_t \mathbf{v}_t(\epsilon) \right\} \right] \leq \mathbb{E}_\epsilon \left[ \sum_{\mathbf{v} \in V} \exp \left\{ \lambda \sum_{t=1}^n \epsilon_t \mathbf{v}_t(\epsilon) \right\} \right] \end{aligned}$$

With the usual technique of peeling from the end,

$$\begin{aligned} M(\lambda) &\leq \sum_{\mathbf{v} \in V} \mathbb{E}_{\epsilon_1, \dots, \epsilon_n} \left[ \prod_{t=1}^n \exp \{ \lambda \epsilon_t \mathbf{v}_t(\epsilon_{1:t-1}) \} \right] \\ &= \sum_{\mathbf{v} \in V} \mathbb{E}_{\epsilon_1, \dots, \epsilon_{n-1}} \left[ \prod_{t=1}^{n-1} \exp \{ \lambda \epsilon_t \mathbf{v}_t(\epsilon_{1:t-1}) \} \times \left( \frac{\exp \{ \lambda \mathbf{v}_n(\epsilon_{1:n-1}) \} + \exp \{ -\lambda \mathbf{v}_n(\epsilon_{1:n-1}) \}}{2} \right) \right] \\ &\leq \sum_{\mathbf{v} \in V} \mathbb{E}_{\epsilon_1, \dots, \epsilon_{n-1}} \left[ \prod_{t=1}^{n-1} \exp \{ \lambda \epsilon_t \mathbf{v}_t(\epsilon_{1:t-1}) \} \times \exp \left\{ \frac{\lambda^2 \mathbf{v}_n(\epsilon_{1:n-1})^2}{2} \right\} \right] \end{aligned}$$

where we used the inequality  $\frac{1}{2} \{ \exp(a) + \exp(-a) \} \leq \exp(a^2/2)$ , valid for all  $a \in \mathbb{R}$ . Peeling off the second term is a bit more involved:

$$\begin{aligned} M(\lambda) &\leq \sum_{\mathbf{v} \in V} \mathbb{E}_{\epsilon_1, \dots, \epsilon_{n-2}} \left[ \prod_{t=1}^{n-2} \exp \{ \lambda \epsilon_t \mathbf{v}_t(\epsilon_{1:t-1}) \} \times \right. \\ &\quad \left. \frac{1}{2} \left( \exp \{ \lambda \mathbf{v}_{n-1}(\epsilon_{1:n-2}) \} \exp \left\{ \frac{\lambda^2 \mathbf{v}_n((\epsilon_{1:n-2}, 1))^2}{2} \right\} \right. \right. \\ &\quad \left. \left. + \exp \{ -\lambda \mathbf{v}_{n-1}(\epsilon_{1:n-2}) \} \exp \left\{ \frac{\lambda^2 \mathbf{v}_n((\epsilon_{1:n-2}, -1))^2}{2} \right\} \right) \right] \end{aligned}$$

Consider the term inside:

$$\begin{aligned}
& \frac{1}{2} \left( \exp \{ \lambda \mathbf{v}_{n-1}(\epsilon_{1:n-2}) \} \exp \left\{ \frac{\lambda^2 \mathbf{v}_n((\epsilon_{1:n-2}, 1))^2}{2} \right\} + \exp \{ -\lambda \mathbf{v}_{n-1}(\epsilon_{1:n-2}) \} \exp \left\{ \frac{\lambda^2 \mathbf{v}_n((\epsilon_{1:n-2}, -1))^2}{2} \right\} \right) \\
& \leq \max_{\epsilon_{n-1}} \left( \exp \left\{ \frac{\lambda^2 \mathbf{v}_n((\epsilon_{1:n-2}, \epsilon_{n-1}))^2}{2} \right\} \right) \frac{\exp \{ \lambda \mathbf{v}_{n-1}(\epsilon_{1:n-2}) \} + \exp \{ -\lambda \mathbf{v}_{n-1}(\epsilon_{1:n-2}) \}}{2} \\
& \leq \max_{\epsilon_{n-1}} \left( \exp \left\{ \frac{\lambda^2 \mathbf{v}_n((\epsilon_{1:n-2}, \epsilon_{n-1}))^2}{2} \right\} \right) \exp \left\{ \frac{\lambda^2 \mathbf{v}_{n-1}(\epsilon_{1:n-2})^2}{2} \right\} \\
& = \exp \left\{ \frac{\lambda^2 \max_{\epsilon_{n-1} \in \{\pm 1\}} (\mathbf{v}_{n-1}(\epsilon_{1:n-2})^2 + \mathbf{v}_n(\epsilon_{1:n-1})^2)}{2} \right\}
\end{aligned}$$

Repeating the last steps, we show that for any  $i$ ,

$$M(\lambda) \leq \sum_{\mathbf{v} \in V} \mathbb{E}_{\epsilon_1, \dots, \epsilon_{i-1}} \left[ \prod_{t=1}^{i-1} \exp \{ \lambda \epsilon_t \mathbf{v}_t(\epsilon_{1:t-1}) \} \times \exp \left\{ \frac{\lambda^2 \max_{\epsilon_1, \dots, \epsilon_{n-1} \in \{\pm 1\}} \sum_{t=i}^n \mathbf{v}_t(\epsilon_{1:t-1})^2}{2} \right\} \right]$$

We arrive at

$$\begin{aligned}
M(\lambda) & \leq \sum_{\mathbf{v} \in V} \exp \left\{ \frac{\lambda^2 \max_{\epsilon_1, \dots, \epsilon_{n-1} \in \{\pm 1\}} \sum_{t=1}^n \mathbf{v}_t(\epsilon_{1:t-1})^2}{2} \right\} \\
& \leq |V| \exp \left\{ \frac{\lambda^2 \max_{\mathbf{v} \in V} \max_{\epsilon \in \{\pm 1\}^n} \sum_{t=1}^n \mathbf{v}_t(\epsilon)^2}{2} \right\}
\end{aligned}$$

Taking logarithms on both sides, dividing by  $\lambda$  and setting  $\lambda = \sqrt{\frac{2 \log(|V|)}{\max_{\mathbf{v} \in V} \max_{\epsilon \in \{\pm 1\}^n} \sum_{t=1}^n \mathbf{v}_t(\epsilon)^2}}$  we conclude that

$$\mathbb{E}_{\epsilon_1, \dots, \epsilon_n} \left[ \max_{\mathbf{v} \in V} \sum_{t=1}^n \epsilon_t \mathbf{v}_t(\epsilon) \right] \leq \sqrt{2 \log(|V|) \max_{\mathbf{v} \in V} \max_{\epsilon \in \{\pm 1\}^n} \sum_{t=1}^n \mathbf{v}_t(\epsilon)^2}$$

□

**Proof of Lemma 30.** We prove the first inequality. Let  $\{\mathbf{w}^1, \dots, \mathbf{w}^M\}$  be a largest strongly  $2\alpha$ -separated set of  $\mathcal{F}(z)$  with  $M = \mathcal{M}_p(2\alpha, \mathcal{F}, z)$ . Let  $\{\mathbf{v}^1, \dots, \mathbf{v}^N\}$  be a smallest  $\alpha$ -cover of  $\mathcal{F}$  on  $z$  with  $N = \mathfrak{N}_p^{\text{seq}}(\alpha, \mathcal{F}, z)$ . For the sake of contradiction, assume  $M > N$ . Consider a path  $\epsilon \in \{\pm 1\}^n$  on which all the trees  $\{\mathbf{w}^1, \dots, \mathbf{w}^M\}$  are  $(2\alpha)$ -separated. By the definition of a cover, for any  $\mathbf{w}^i$  there exists a tree  $\mathbf{v}^j$  such that

$$\left( \frac{1}{n} \sum_{t=1}^n |\mathbf{v}_t^j(\epsilon) - \mathbf{w}_t^i(\epsilon)|^p \right)^{1/p} \leq \alpha.$$

Since  $M > N$ , there must exist distinct  $\mathbf{w}^i$  and  $\mathbf{w}^k$ , for which the covering tree  $\mathbf{v}^j$  is the same for the given path  $\epsilon$ . By triangle inequality,

$$\left( \frac{1}{n} \sum_{t=1}^n |\mathbf{w}_t^i(\epsilon) - \mathbf{w}_t^k(\epsilon)|^p \right)^{1/p} \leq 2\alpha,$$

which is a contradiction. We conclude that  $M \leq N$ .

Now, we prove the second inequality. Consider a maximal  $\alpha$ -packing  $V \subseteq \mathcal{F}(z)$  of size  $\mathcal{D}_p(\alpha, \mathcal{F}, z)$ . Since this is a *maximal*  $\alpha$ -packing, for any  $f \in \mathcal{F}$ , there is no path on which  $f(z)$  is  $\alpha$ -separated from every member of the packing. In other words, for every path  $\epsilon \in \{\pm 1\}^n$ , there is a member of the packing  $v \in V$  such that

$$\left( \frac{1}{n} \sum_{t=1}^n |v_t(\epsilon) - f(z_t(\epsilon))|^p \right)^{1/p} \leq \alpha$$

which means that the packing  $V$  is a cover.  $\square$

**Proof of Theorem 31.** For any  $d \geq 0$  and  $n \geq 0$ , define the function

$$g_k(d, n) = \sum_{i=0}^d \binom{n}{i} k^i.$$

It is not difficult to verify that this function satisfies the recurrence

$$g_k(d, n) = g_k(d, n-1) + k g_k(d-1, n-1)$$

for all  $d, n \geq 1$ . To visualize this recursion, consider a  $k \times n$  matrix and ask for ways to choose at most  $d$  columns followed by a choice among the  $k$  rows for each chosen column. The task can be decomposed into (a) making the  $d$  column choices out of the first  $n-1$  columns, followed by picking rows (there are  $g_k(d, n-1)$  ways to do it) or (b) choosing  $d-1$  columns (followed by row choices) out of the first  $n-1$  columns and choosing a row for the  $n$ th column (there are  $k g_k(d-1, n-1)$  ways to do it). This gives the recursive formula.

In what follows, we shall refer to an  $L_\infty$  cover at scale  $1/2$  simply as a  $1/2$ -cover. The theorem claims that the size of a minimal  $1/2$ -cover is at most  $g_k(d, n)$ . The proof proceeds by induction on  $n+d$ .

**Base:** For  $d=1$  and  $n=1$ , there is only one node in the tree, i.e. the tree is defined by the constant  $z_1 \in \mathcal{Z}$ . Functions in  $\mathcal{F}$  can take up to  $k+1$  values on  $z_1$ , i.e.  $\mathfrak{N}^{\text{seq}}(0, \mathcal{F}, 1) \leq k+1$  (and, thus, also for the  $1/2$ -cover). Using the convention  $\binom{n}{0} = 1$ , we indeed verify that  $g_k(1, 1) = 1+k = k+1$ . The same calculation gives the base case for  $n=1$  and any  $d \in \mathbb{N}$ . Furthermore, for any  $n \in \mathbb{N}$  if  $d=0$ , then there is no point which is 2-shattered by  $\mathcal{F}$ . This means that functions in  $\mathcal{F}$  differ by at most 1 on any point of  $\mathcal{Z}$ . Thus, there is a  $1/2$  cover of size  $1 = g_k(0, n)$ , verifying this base case.

**Induction step:** Suppose by the way of induction that the statement holds for  $(d, n-1)$  and  $(d-1, n-1)$ . Consider any tree  $z$  of depth  $n$  with  $\text{fat}_2^{\text{seq}}(\mathcal{F}, z) = d$ . Define the partition  $\mathcal{F} = \mathcal{F}_0 \cup \dots \cup \mathcal{F}_k$  with  $\mathcal{F}_i = \{f \in \mathcal{F} : f(z_1) = i\}$  for  $i \in \{0, \dots, k\}$ , where  $z_1$  is the root of  $z$ . Let  $n = |\{i : \text{fat}_2^{\text{seq}}(\mathcal{F}_i, z) = d\}|$ . Suppose first, for the sake of contradiction, that  $\text{fat}_2^{\text{seq}}(\mathcal{F}_i, z) = \text{fat}_2^{\text{seq}}(\mathcal{F}_j, z) = d$  for  $|i-j| \geq 2$ . Then there exist two trees  $z$  and  $v$  of depth  $d$  which are 2-shattered by  $\mathcal{F}_i$  and  $\mathcal{F}_j$ , respectively, and with

$\text{Img}(z), \text{Img}(\mathbf{v}) \subseteq \text{Img}(z)$ . Since functions within each subset  $\mathcal{F}_i$  take on the same values on  $z_1$ , we conclude that  $z_1 \notin \text{Img}(z), z_1 \notin \text{Img}(\mathbf{v})$ . This follows immediately from the definition of shattering. We now *join* the two shattered  $z$  and  $\mathbf{v}$  trees with  $z_1$  at the root and observe that  $\mathcal{F}_i \cup \mathcal{F}_j$  2-shatters this resulting tree of depth  $d+1$ , which is a contradiction. Indeed, the witness  $\mathbb{R}$ -valued tree  $\mathbf{s}$  is constructed by joining the two witnesses for the 2-shattered trees  $z$  and  $\mathbf{v}$  and by defining the root as  $s_1 = (i+j)/2$ . It is easy to see that  $\mathbf{s}$  is a witness to the shattering. Given any  $\epsilon \in \{\pm 1\}^{d+1}$ , there is a function  $f^i \in \mathcal{F}_i$  which realizes the desired separation under the signs  $(\epsilon_2, \dots, \epsilon_{d+1})$  for the tree  $z$  and there is a function  $f^j \in \mathcal{F}_j$  which does the same for  $\mathbf{v}$ . Depending on  $\epsilon_1 = +1$  or  $\epsilon_1 = -1$ , either  $f^i$  or  $f^j$  realize the separation over  $\epsilon$ .

We conclude that the number of subsets of  $\mathcal{F}$  with fat-shattering dimension equal to  $d$  cannot be more than two (for otherwise at least two indices will be separated by 2 or more). We have three cases:  $n = 0$ ,  $n = 1$ , or  $n = 2$ , and in the last case it must be that the indices of the two subsets differ by 1.

First, consider any  $\mathcal{F}_i$  with  $\text{fat}_2^{\text{seq}}(\mathcal{F}_i, z) \leq d-1, i \in \{0, \dots, k\}$ . By induction, there are 1/2-covers  $V^\ell$  and  $V^r$  of  $\mathcal{F}_i$  on the subtrees  $z^\ell$  and  $z^r$ , respectively, both of size at most  $g_k(d-1, n-1)$ . Informally, out of these 1/2-covers we can create a 1/2-cover  $V$  for  $\mathcal{F}_i$  on  $z$  by pairing the 1/2-covers in  $V^\ell$  and  $V^r$ . The resulting cover of  $\mathcal{F}_i$  will be of size  $g_k(d-1, n-1)$ . Formally, consider a set of pairs  $(\mathbf{v}^\ell, \mathbf{v}^r)$  of trees, with  $\mathbf{v}^\ell \in V^\ell, \mathbf{v}^r \in V^r$  and such that each tree in  $V^\ell$  and  $V^r$  appears in at least one of the pairs. Clearly, this can be done using at most  $g_k(d-1, n-1)$  pairs, and such a pairing is not unique. We join the subtrees in every pair  $(\mathbf{v}^\ell, \mathbf{v}^r)$  with a constant  $i$  as the root, thus creating a set  $V$  of trees,  $|V| \leq g_k(d-1, n-1)$ . We claim that  $V$  is a 1/2-cover for  $\mathcal{F}_i$  on  $z$ . Note that all the functions in  $\mathcal{F}_i$  take on the same value  $i$  on  $z_1$  and by construction  $\mathbf{v}_1 = i$  for any  $\mathbf{v} \in V$ . Now, consider any  $f \in \mathcal{F}_i$  and  $\epsilon \in \{\pm 1\}^n$ . Without loss of generality, assume  $\epsilon_1 = -1$ . By assumption, there is a  $\mathbf{v}^\ell \in V^\ell$  such that  $|\mathbf{v}_t^\ell(\epsilon_{2:n}) - f(z_{t+1}(\epsilon_{1:n}))| \leq 1/2$  for any  $t \in [n-1]$ . By construction  $\mathbf{v}^\ell$  appears as a left subtree of at least one tree in  $V$ , which, therefore, matches the values of  $f$  for  $\epsilon_{1:n}$ . The same argument holds for  $\epsilon_1 = +1$  by finding an appropriate subtree in  $V^r$ . We conclude that  $V$  is a 1/2-cover of  $\mathcal{F}_i$  on  $z$ , and this holds for any  $i \in \{0, \dots, k\}$  with  $\text{fat}_2^{\text{seq}}(\mathcal{F}_i, z) \leq d-1$ . Therefore, the total size of a 1/2-cover for the union  $\cup_{i: \text{fat}_2^{\text{seq}}(\mathcal{F}_i, z) \leq d-1} \mathcal{F}_i$  is at most  $(k+1-n)g_k(d-1, n-1)$ . If  $n = 0$ , the induction step is proven because  $g_k(d-1, n-1) \leq g_k(d, n-1)$  and so the total size of the constructed cover is at most

$$(k+1)g_k(d-1, n-1) \leq g_k(d, n-1) + kg_k(d-1, n-1) = g_k(d, n).$$

Now, consider the case  $n = 1$  and let  $\text{fat}_2^{\text{seq}}(\mathcal{F}_i, z) = d$ . An argument exactly as above yields a 1/2-cover for  $\mathcal{F}_i$ , and this cover is of size at most  $g_k(d, n-1)$  by induction. The total 1/2-cover is therefore of size at most

$$g_k(d, n-1) + kg_k(d-1, n-1) = g_k(d, n).$$

Lastly, for  $n = 2$ , suppose  $\text{fat}_2^{\text{seq}}(\mathcal{F}_i, z) = \text{fat}_2^{\text{seq}}(\mathcal{F}_j, z) = d$  for  $|i-j| = 1$ . Let  $\mathcal{F}' = \mathcal{F}_i \cup \mathcal{F}_j$ . Note that  $\text{fat}_2^{\text{seq}}(\mathcal{F}', z) = d$ . Just as before, the 1/2-covering for  $z$  can be constructed by considering the 1/2-covers



for the two subtrees. However, when joining any  $(\mathbf{v}^\ell, \mathbf{v}^r)$ , we take  $(i + j)/2$  as the root. It is straightforward to check that the resulting cover is indeed an  $1/2$ -cover, thanks to the relation  $|i - j| = 1$ . The size of the constructed cover is, by induction,  $g_k(d, n - 1)$ , and the induction step follows. This concludes the induction proof, yielding the main statement of the theorem.

Finally, the upper bound on  $g_k(d, n)$  is

$$\sum_{i=1}^d \binom{n}{i} k^i \leq \left(\frac{kn}{d}\right)^d \sum_{i=1}^d \binom{n}{i} \left(\frac{d}{n}\right)^i \leq \left(\frac{kn}{d}\right)^d \left(1 + \frac{d}{n}\right)^n \leq \left(\frac{ekn}{d}\right)^d$$

whenever  $n \geq d$ . □

**Proof of Theorem 33.** The proof is very close to the proof of Theorem 31, with a few key differences. As before, for any  $d \geq 0$  and  $n \geq 0$ , define the function  $g_k(d, n) = \sum_{i=0}^d \binom{n}{i} k^i$ .

The theorem claims that the size of a minimal 0-cover is at most  $g_k(d, n)$ . The proof proceeds by induction on  $n + d$ .

**Base:** For  $d = 1$  and  $n = 1$ , there is only one node in the tree, i.e. the tree is defined by the constant  $z_1 \in \mathcal{Z}$ . Functions in  $\mathcal{F}$  can take up to  $k + 1$  values on  $z_1$ , i.e.  $\mathfrak{N}^{\text{seq}}(0, \mathcal{F}, 1) \leq k + 1$ . Using the convention  $\binom{n}{0} = 1$ , we indeed verify that  $g_k(1, 1) = 1 + k = k + 1$ . The same calculation gives the base case for  $n = 1$  and any  $d \in \mathbb{N}$ . Furthermore, for any  $n \in \mathbb{N}$  if  $d = 0$ , then there is no point which is 1-shattered by  $\mathcal{F}$ . This means that all functions in  $\mathcal{F}$  are identical, proving that there is a 0-cover of size  $1 = g_k(0, n)$ .

**Induction step:** Suppose by the way of induction that the statement holds for  $(d, n - 1)$  and  $(d - 1, n - 1)$ . Consider any tree  $z$  of depth  $n$  with  $\text{fat}_1^{\text{seq}}(\mathcal{F}, z) = d$ . Define the partition  $\mathcal{F} = \mathcal{F}_0 \cup \dots \cup \mathcal{F}_k$  with  $\mathcal{F}_i = \{f \in \mathcal{F} : f(z_1) = i\}$  for  $i \in \{0, \dots, k\}$ , where  $z_1$  is the root of  $z$ .

We first argue that  $\text{fat}_1^{\text{seq}}(\mathcal{F}_i, z) = d$  for at most one value  $i \in \{0, \dots, k\}$ . By the way of contradiction, suppose we do have  $\text{fat}_1^{\text{seq}}(\mathcal{F}_i, z) = \text{fat}_1^{\text{seq}}(\mathcal{F}_j, z) = d$  for  $i \neq j$ . Then there exist two trees  $z$  and  $\mathbf{v}$  of depth  $d$  1-shattered by  $\mathcal{F}_i$  and  $\mathcal{F}_j$ , respectively, and with  $\text{Img}(z), \text{Img}(\mathbf{v}) \subseteq \text{Img}(z)$ . Since functions within each subset  $\mathcal{F}_i$  take on the same values on  $z_1$ , we conclude that  $z_1 \notin \text{Img}(z), z_1 \notin \text{Img}(\mathbf{v})$ . This follows immediately from the definition of shattering. We now *join* the two shattered  $z$  and  $\mathbf{v}$  trees with  $z_1$  at the root and observe that  $\mathcal{F}_i \cup \mathcal{F}_j$  1-shatters this resulting tree of depth  $d + 1$ , which is a contradiction. Indeed, the witness  $\mathbb{R}$ -valued tree  $\mathbf{s}$  is constructed by joining the two witnesses for the 1-shattered trees  $z$  and  $\mathbf{v}$  and by defining the root as  $\mathbf{s}_1 = (i + j)/2$ . It is easy to see that  $\mathbf{s}$  is a witness to the shattering. Given any  $\epsilon \in \{\pm 1\}^{d+1}$ , there is a function  $f^i \in \mathcal{F}_i$  which realizes the desired separation under the signs  $(\epsilon_2, \dots, \epsilon_{d+1})$  for the tree  $z$  and there is a function  $f^j \in \mathcal{F}_j$  which does the same for  $\mathbf{v}$ . Depending on  $\epsilon_1 = +1$  or  $\epsilon_1 = -1$ , either  $f^i$  or  $f^j$  realize the separation over  $\epsilon$ .

We conclude that  $\text{fat}_1^{\text{seq}}(\mathcal{F}_i, z) = d$  for at most one  $i \in \{0, \dots, k\}$ . Without loss of generality, assume  $\text{fat}_1^{\text{seq}}(\mathcal{F}_0, z) \leq d$  and  $\text{fat}_1^{\text{seq}}(\mathcal{F}_i, z) \leq d - 1$  for  $i \in \{1, \dots, k\}$ . By induction, for any  $\mathcal{F}_i, i \in \{1, \dots, k\}$ ,

there are 0-covers  $V^\ell$  and  $V^r$  of  $\mathcal{F}_i$  on the subtrees  $z^\ell$  and  $z^r$ , respectively, both of size at most  $g_k(d-1, n-1)$ . Out of these 0-covers we can create a 0-cover  $V$  for  $\mathcal{F}_i$  on  $z$  by pairing the 0-covers in  $V^\ell$  and  $V^r$ . Formally, consider a set of pairs  $(\mathbf{v}^\ell, \mathbf{v}^r)$  of trees, with  $\mathbf{v}^\ell \in V^\ell$ ,  $\mathbf{v}^r \in V^r$  and such that each tree in  $V^\ell$  and  $V^r$  appears in at least one of the pairs. Clearly, this can be done using at most  $g_k(d-1, n-1)$  pairs, and such a pairing is not unique. We join the subtrees in every pair  $(\mathbf{v}^\ell, \mathbf{v}^r)$  with a constant  $i$  as the root, thus creating a set  $V$  of trees,  $|V| \leq g_k(d-1, n-1)$ . We claim that  $V$  is a 0-cover for  $\mathcal{F}_i$  on  $z$ . Note that all the functions in  $\mathcal{F}_i$  take on the same value  $i$  on  $z_1$  and by construction  $\mathbf{v}_1 = i$  for any  $\mathbf{v} \in V$ . Now, consider any  $f \in \mathcal{F}_i$  and  $\epsilon \in \{\pm 1\}^n$ . Without loss of generality, assume  $\epsilon_1 = -1$ . By assumption, there is a  $\mathbf{v}^\ell \in V^\ell$  such that  $\mathbf{v}_t^\ell(\epsilon_{2:n}) = f(z_{t+1}(\epsilon_{1:n}))$  for any  $t \in [n-1]$ . By construction  $\mathbf{v}^\ell$  appears as a left subtree of at least one tree in  $V$ , which, therefore, matches the values of  $f$  for  $\epsilon_{1:n}$ . The same argument holds for  $\epsilon_1 = +1$  by finding an appropriate subtree in  $V^r$ . We conclude that  $V$  is a 0-cover of  $\mathcal{F}_i$  on  $z$ , and this holds for any  $i \in \{1, \dots, k\}$ . Therefore, the total size of a 0-cover for  $\mathcal{F}_1 \cup \dots \cup \mathcal{F}_k$  is at most  $kg_k(d-1, n-1)$ . A similar argument yields a 0-cover for  $\mathcal{F}_0$  on  $z$  of size at most  $g_k(d, n-1)$  by induction. Thus, the size of the resulting 0-cover of  $\mathcal{F}$  on  $z$  is at most

$$g_k(d, n-1) + kg_k(d-1, n-1) = g_k(d, n),$$

completing the induction step and yielding the main statement of the theorem.

The upper bound on  $g_k(d, n)$  appears in the proof of Theorem 31.

□

**Proof of Corollary 32.** The first two inequalities follow by simple comparison of norms. It remains to prove the bound for the  $\ell_\infty$  covering. For any  $\alpha > 0$  define an  $\alpha$ -discretization of the  $[-1, 1]$  interval as  $B_\alpha = \{-1 + \alpha/2, -1 + 3\alpha/2, \dots, -1 + (2k+1)\alpha/2, \dots\}$  for  $0 \leq k$  and  $(2k+1)\alpha \leq 4$ . Also for any  $a \in [-1, 1]$  define  $\lfloor a \rfloor_\alpha = \operatorname{argmin}_{r \in B_\alpha} |r - a|$  with ties being broken by choosing the smaller discretization point. For a function  $f : \mathcal{Z} \mapsto [-1, 1]$  let the function  $\lfloor f \rfloor_\alpha$  be defined pointwise as  $\lfloor f(x) \rfloor_\alpha$ , and let  $\lfloor \mathcal{F} \rfloor_\alpha = \{\lfloor f \rfloor_\alpha : f \in \mathcal{F}\}$ . First, we prove that  $\mathfrak{N}_\infty^{\text{seq}}(\alpha, \mathcal{F}, \mathbf{z}) \leq \mathfrak{N}_\infty^{\text{seq}}(\alpha/2, \lfloor \mathcal{F} \rfloor_\alpha, \mathbf{z})$ . Indeed, suppose the set of trees  $V$  is a minimal  $\alpha/2$ -cover of  $\lfloor \mathcal{F} \rfloor_\alpha$  on  $\mathbf{z}$ . That is,

$$\forall f_\alpha \in \lfloor \mathcal{F} \rfloor_\alpha, \forall \epsilon \in \{\pm 1\}^n \exists \mathbf{v} \in V \text{ s.t. } |\mathbf{v}_t(\epsilon) - f_\alpha(\mathbf{z}_t(\epsilon))| \leq \alpha/2$$

Pick any  $f \in \mathcal{F}$  and let  $f_\alpha = \lfloor f \rfloor_\alpha$ . Then  $\|f - f_\alpha\|_\infty \leq \alpha/2$ . Then for all  $\epsilon \in \{\pm 1\}^n$  and any  $t \in [n]$

$$|f(\mathbf{z}_t(\epsilon)) - \mathbf{v}_t(\epsilon)| \leq |f(\mathbf{z}_t(\epsilon)) - f_\alpha(\mathbf{z}_t(\epsilon))| + |f_\alpha(\mathbf{z}_t(\epsilon)) - \mathbf{v}_t(\epsilon)| \leq \alpha,$$

and so  $V$  also provides an  $L_\infty$  cover at scale  $\alpha$ .

We conclude that  $\mathfrak{N}_\infty^{\text{seq}}(\alpha, \mathcal{F}, \mathbf{z}) \leq \mathfrak{N}_\infty^{\text{seq}}(\alpha/2, \lfloor \mathcal{F} \rfloor_\alpha, \mathbf{z}) = \mathfrak{N}_\infty^{\text{seq}}(1/2, \mathcal{G}, \mathbf{z})$  where  $\mathcal{G} = \frac{1}{\alpha} \lfloor \mathcal{F} \rfloor_\alpha$ . The functions of  $\mathcal{G}$  take on a discrete set of at most  $\lfloor 2/\alpha \rfloor + 1$  values. Obviously, by adding a constant to all the functions in  $\mathcal{G}$ , we can make the set of values to be  $\{0, \dots, \lfloor 2/\alpha \rfloor\}$ . We now apply Theorem 31 with an upper

bound  $\sum_{i=0}^d \binom{n}{i} k^i \leq (ekn)^d$  which holds for any  $n > 0$ . This yields  $\mathfrak{N}_{\infty}^{\text{seq}}(1/2, \mathcal{G}, \mathbf{z}) \leq (2en/\alpha)^{\text{fat}_2^{\text{seq}}(\mathcal{G})}$ .

It remains to prove  $\text{fat}_2^{\text{seq}}(\mathcal{G}) \leq \text{fat}_{\alpha}^{\text{seq}}(\mathcal{F})$ , or, equivalently (by scaling)  $\text{fat}_{2\alpha}^{\text{seq}}(\lfloor \mathcal{F} \rfloor_{\alpha}) \leq \text{fat}_{\alpha}^{\text{seq}}(\mathcal{F})$ . To this end, suppose there exists an  $\mathbb{R}$ -valued tree  $\mathbf{z}$  of depth  $d = \text{fat}_{2\alpha}^{\text{seq}}(\lfloor \mathcal{F} \rfloor_{\alpha})$  such that there is a witness tree  $\mathbf{s}$  with

$$\forall \epsilon \in \{\pm 1\}^d, \exists f_{\alpha} \in \lfloor \mathcal{F} \rfloor_{\alpha} \text{ s.t. } \forall t \in [d], \epsilon_t(f_{\alpha}(\mathbf{z}_t(\epsilon)) - \mathbf{s}_t(\epsilon)) \geq \alpha$$

Using the fact that for any  $f \in \mathcal{F}$  and  $f_{\alpha} = \lfloor f \rfloor_{\alpha}$  we have  $\|f - f_{\alpha}\|_{\infty} \leq \alpha/2$ , it follows that

$$\forall \epsilon \in \{\pm 1\}^d, \exists f \in \mathcal{F} \text{ s.t. } \forall t \in [d], \epsilon_t(f(\mathbf{z}_t(\epsilon)) - \mathbf{s}_t(\epsilon)) \geq \alpha/2$$

That is,  $\mathbf{s}$  is a witness to  $\alpha$ -shattering by  $\mathcal{F}$ . Thus for any  $\mathbf{z}$ ,

$$\mathfrak{N}_{\infty}^{\text{seq}}(\alpha, \mathcal{F}, \mathbf{z}) \leq \mathfrak{N}_{\infty}^{\text{seq}}(\alpha/2, \lfloor \mathcal{F} \rfloor_{\alpha}, \mathbf{z}) \leq \left(\frac{2en}{\alpha}\right)^{\text{fat}_{2\alpha}^{\text{seq}}(\lfloor \mathcal{F} \rfloor_{\alpha})} \leq \left(\frac{2en}{\alpha}\right)^{\text{fat}_{\alpha}^{\text{seq}}(\mathcal{F})}$$

□

**Proof of Theorem 35.** Define  $\beta_0 = 1$  and  $\beta_j = 2^{-j}$ . For a fixed tree  $\mathbf{z}$  of depth  $n$ , let  $V_j$  be an  $\ell_2$ -cover at scale  $\beta_j$ . For any path  $\epsilon \in \{\pm 1\}^n$  and any  $f \in \mathcal{F}$ , let  $\mathbf{v}[f, \epsilon]^j \in V_j$  the element of the cover such that

$$\sqrt{\frac{1}{n} \sum_{t=1}^n |\mathbf{v}[f, \epsilon]_t^j(\epsilon) - f(\mathbf{z}_t(\epsilon))|^2} \leq \beta_j$$

By the definition such a  $\mathbf{v}[f, \epsilon]^j \in V_j$  exists, and we assume for simplicity this element is unique (ties can be broken in an arbitrary manner). Thus,  $f \mapsto \mathbf{v}[f, \epsilon]^j$  is a well-defined mapping for any fixed  $\epsilon$  and  $j$ . As before,  $\mathbf{v}[f, \epsilon]_t^j$  denotes the  $t$ -th mapping of  $\mathbf{v}[f, \epsilon]^j$ . For any  $t \in [n]$ , we have

$$f(\mathbf{z}_t(\epsilon)) = f(\mathbf{z}_t(\epsilon)) - \mathbf{v}[f, \epsilon]_t^N(\epsilon) + \sum_{j=1}^N (\mathbf{v}[f, \epsilon]_t^j(\epsilon) - \mathbf{v}[f, \epsilon]_t^{j-1}(\epsilon))$$

where  $\mathbf{v}[f, \epsilon]_t^0(\epsilon) = 0$ . Hence,

$$\begin{aligned} \mathbb{E}_{\epsilon} \left[ \sup_{f \in \mathcal{F}} \sum_{t=1}^n \epsilon_t f(\mathbf{z}_t(\epsilon)) \right] &= \mathbb{E}_{\epsilon} \left[ \sup_{f \in \mathcal{F}} \sum_{t=1}^n \epsilon_t \left( f(\mathbf{z}_t(\epsilon)) - \mathbf{v}[f, \epsilon]_t^N(\epsilon) + \sum_{j=1}^N (\mathbf{v}[f, \epsilon]_t^j(\epsilon) - \mathbf{v}[f, \epsilon]_t^{j-1}(\epsilon)) \right) \right] \\ &= \mathbb{E}_{\epsilon} \left[ \sup_{f \in \mathcal{F}} \sum_{t=1}^n \epsilon_t \left( f(\mathbf{z}_t(\epsilon)) - \mathbf{v}[f, \epsilon]_t^N(\epsilon) \right) + \sum_{t=1}^n \epsilon_t \left( \sum_{j=1}^N (\mathbf{v}[f, \epsilon]_t^j(\epsilon) - \mathbf{v}[f, \epsilon]_t^{j-1}(\epsilon)) \right) \right] \\ &\leq \mathbb{E}_{\epsilon} \left[ \sup_{f \in \mathcal{F}} \sum_{t=1}^n \epsilon_t \left( f(\mathbf{z}_t(\epsilon)) - \mathbf{v}[f, \epsilon]_t^N(\epsilon) \right) \right] + \mathbb{E}_{\epsilon} \left[ \sup_{f \in \mathcal{F}} \sum_{t=1}^n \epsilon_t \left( \sum_{j=1}^N (\mathbf{v}[f, \epsilon]_t^j(\epsilon) - \mathbf{v}[f, \epsilon]_t^{j-1}(\epsilon)) \right) \right] \end{aligned} \tag{4.8.2}$$

The first term above can be bounded via the Cauchy-Schwarz inequality as

$$\mathbb{E}_\epsilon \left[ \sup_{f \in \mathcal{F}} \sum_{t=1}^n \epsilon_t (f(\mathbf{z}_t(\epsilon)) - \mathbf{v}[f, \epsilon]_t^N(\epsilon)) \right] \leq n \mathbb{E}_\epsilon \left[ \sup_{f \in \mathcal{F}} \sum_{t=1}^n \frac{\epsilon_t}{\sqrt{n}} \frac{(f(\mathbf{z}_t(\epsilon)) - \mathbf{v}[f, \epsilon]_t^N(\epsilon))}{\sqrt{n}} \right] \leq n \beta_N.$$

The second term in Eq. (4.8.2) is bounded by considering successive refinements of the cover. The argument, however, is more delicate than in the classical case, as the trees  $\mathbf{v}[f, \epsilon]^j$ ,  $\mathbf{v}[f, \epsilon]^{j-1}$  depend on the particular path. Consider all possible pairs of  $\mathbf{v}^s \in V_j$  and  $\mathbf{v}^r \in V_{j-1}$ , for  $1 \leq s \leq |V_j|$ ,  $1 \leq r \leq |V_{j-1}|$ , where we assumed an arbitrary enumeration of elements. For each pair  $(\mathbf{v}^s, \mathbf{v}^r)$ , define a real-valued tree  $\mathbf{w}^{(s,r)}$  by

$$\mathbf{w}_t^{(s,r)}(\epsilon) = \begin{cases} \mathbf{v}_t^s(\epsilon) - \mathbf{v}_t^r(\epsilon) & \text{if there exists } f \in \mathcal{F} \text{ s.t. } \mathbf{v}^s = \mathbf{v}[f, \epsilon]^j, \mathbf{v}^r = \mathbf{v}[f, \epsilon]^{j-1} \\ 0 & \text{otherwise.} \end{cases}$$

for all  $t \in [n]$  and  $\epsilon \in \{\pm 1\}^n$ . It is crucial that  $\mathbf{w}^{(s,r)}$  can be non-zero only on those paths  $\epsilon$  for which  $\mathbf{v}^s$  and  $\mathbf{v}^r$  are indeed the members of the covers (at successive resolutions) close to  $f(\mathbf{z}(\epsilon))$  (in the  $\ell_2$  sense) for some  $f \in \mathcal{F}$ . It is easy to see that  $\mathbf{w}^{(s,r)}$  is well-defined. Let the set of trees  $W_j$  be defined as

$$W_j = \left\{ \mathbf{w}^{(s,r)} : 1 \leq s \leq |V_j|, 1 \leq r \leq |V_{j-1}| \right\}$$

Now, the second term in Eq. (4.8.2) can be written as

$$\begin{aligned} \mathbb{E}_\epsilon \left[ \sup_{f \in \mathcal{F}} \sum_{t=1}^n \epsilon_t \sum_{j=1}^N (\mathbf{v}[f, \epsilon]_t^j(\epsilon) - \mathbf{v}[f, \epsilon]_t^{j-1}(\epsilon)) \right] &\leq \sum_{j=1}^N \mathbb{E}_\epsilon \left[ \sup_{f \in \mathcal{F}} \sum_{t=1}^n \epsilon_t (\mathbf{v}[f, \epsilon]_t^j(\epsilon) - \mathbf{v}[f, \epsilon]_t^{j-1}(\epsilon)) \right] \\ &\leq \sum_{j=1}^N \mathbb{E}_\epsilon \left[ \max_{\mathbf{w} \in W_j} \sum_{t=1}^n \epsilon_t \mathbf{w}_t(\epsilon) \right] \end{aligned}$$

The last inequality holds because for any  $j \in [N]$ ,  $\epsilon \in \{\pm 1\}^n$  and  $f \in \mathcal{F}$  there is some  $\mathbf{w}^{(s,r)} \in W_j$  with  $\mathbf{v}[f, \epsilon]^j = \mathbf{v}^s$ ,  $\mathbf{v}[f, \epsilon]^{j-1} = \mathbf{v}^r$  and

$$\mathbf{v}_t^s(\epsilon) - \mathbf{v}_t^r(\epsilon) = \mathbf{w}_t^{(s,r)}(\epsilon) \quad \forall t \leq n.$$

Clearly,  $|W_j| \leq |V_j| \cdot |V_{j-1}|$ . To invoke Lemma 34, it remains to bound the magnitude of all  $\mathbf{w}^{(s,r)} \in W_j$  along all paths. For this purpose, fix  $\mathbf{w}^{(s,r)}$  and a path  $\epsilon$ . If there exists  $f \in \mathcal{F}$  for which  $\mathbf{v}^s = \mathbf{v}[f, \epsilon]^j$  and  $\mathbf{v}^r = \mathbf{v}[f, \epsilon]^{j-1}$ , then  $\mathbf{w}_t^{(s,r)}(\epsilon) = \mathbf{v}[f, \epsilon]_t^j - \mathbf{v}[f, \epsilon]_t^{j-1}$  for any  $t \in [n]$ . By triangle inequality

$$\sqrt{\sum_{t=1}^n \mathbf{w}_t^{(s,r)}(\epsilon)^2} \leq \sqrt{\sum_{t=1}^n (\mathbf{v}[f, \epsilon]_t^j(\epsilon) - f(\mathbf{z}_t(\epsilon)))^2} + \sqrt{\sum_{t=1}^n (\mathbf{v}[f, \epsilon]_t^{j-1}(\epsilon) - f(\mathbf{z}_t(\epsilon)))^2} \leq \sqrt{n}(\beta_j + \beta_{j-1}) = 3\sqrt{n}\beta_j.$$

If there exists no such  $f \in \mathcal{F}$  for the given  $\epsilon$  and  $(s, r)$ , then  $\mathbf{w}_t^{(s,r)}(\epsilon)$  is zero for all  $t \geq t_o$ , for some  $1 \leq t_o < n$ , and thus

$$\sqrt{\sum_{t=1}^n \mathbf{w}_t^{(s,r)}(\epsilon)^2} \leq \sqrt{\sum_{t=1}^n \mathbf{w}_t^{(s,r)}(\epsilon')^2}$$

for any other path  $\epsilon'$  which agrees with  $\epsilon$  up to  $t_o$ . Hence, the bound

$$\sqrt{\sum_{t=1}^n \mathbf{w}_t^{(s,r)}(\epsilon)^2} \leq 3\sqrt{n}\beta_j$$

holds for all  $\epsilon \in \{\pm 1\}^n$  and all  $\mathbf{w}^{(s,r)} \in W_j$ .

Now, back to Eq. (4.8.2), we put everything together and apply Lemma 34:

$$\begin{aligned} \mathbb{E}_\epsilon \left[ \sup_{f \in \mathcal{F}} \sum_{t=1}^n \epsilon_t f(\mathbf{z}_t(\epsilon)) \right] &\leq n \beta_N + \sqrt{n} \sum_{j=1}^N 3\beta_j \sqrt{2 \log(|V_j| |V_{j-1}|)} \\ &\leq n \beta_N + \sqrt{n} \sum_{j=1}^N 6\beta_j \sqrt{\log(|V_j|)} \\ &\leq n \beta_N + 12 \sqrt{n} \sum_{j=1}^N (\beta_j - \beta_{j+1}) \sqrt{\log \mathcal{N}_2(\beta_j, \mathcal{F}, \mathbf{z})} \\ &\leq n \beta_N + 12 \int_{\beta_{N+1}}^{\beta_0} \sqrt{n \log \mathcal{N}_2(\delta, \mathcal{F}, \mathbf{z})} d\delta \end{aligned}$$

where the last but one step is because  $2(\beta_j - \beta_{j+1}) = \beta_j$ . Now for any  $\alpha > 0$ , pick  $N = \sup\{j : \beta_j > 2\alpha\}$ . In this case we see that by our choice of  $N$ ,  $\beta_{N+1} \leq 2\alpha$  and so  $\beta_N = 2\beta_{N+1} \leq 4\alpha$ . Also note that since  $\beta_N > 2\alpha$ ,  $\beta_{N+1} = \frac{\beta_N}{2} > \alpha$ . Hence dividing throughout by  $n$  we conclude that

$$\mathfrak{R}_n^{\text{seq}}(\mathcal{F}) \leq \inf_{\alpha} \left\{ 4\alpha + 12 \int_{\alpha}^1 \sqrt{\frac{\log \mathcal{N}_2(\delta, \mathcal{F}, n)}{n}} d\delta \right\}$$

□

**Proof of Theorem 38.** Let  $(x'_1, \dots, x'_n)$  be a sequence tangent to  $(x_1, \dots, x_n)$ . Recall the notation  $\mathbb{E}_{t-1}[f(x'_t)] =$

$\mathbb{E}\{f(x'_t)|x_1, \dots, x_{t-1}\}$ . By Chebychev's inequality, for any  $f \in \mathcal{F}$ ,

$$\begin{aligned} \mathbb{P}_{\mathcal{D}} \left[ \frac{1}{n} \left| \sum_{t=1}^n (f(x'_t) - \mathbb{E}_{t-1}[f(x'_t)]) \right| > \alpha/2 \mid x_1, \dots, x_n \right] &\leq \frac{\mathbb{E} \left[ \left( \sum_{t=1}^n (f(x'_t) - \mathbb{E}_{t-1}[f(x'_t)]) \right)^2 \mid x_1, \dots, x_n \right]}{n^2 \alpha^2 / 4} \\ &= \frac{\sum_{t=1}^n \mathbb{E} \left[ (f(x'_t) - \mathbb{E}_{t-1}[f(x'_t)])^2 \mid x_1, \dots, x_n \right]}{n^2 \alpha^2 / 4} \\ &\leq \frac{4n}{n^2 \alpha^2 / 4} = \frac{16}{n \alpha^2}. \end{aligned}$$

The second step is due to the fact that the cross terms are zero:

$$\mathbb{E} \left\{ (f(x'_t) - \mathbb{E}_{t-1}[f(x'_t)]) (f(x'_s) - \mathbb{E}_{s-1}[f(x'_s)]) \mid x_1, \dots, x_n \right\} = 0.$$

Hence

$$\inf_{f \in \mathcal{F}} \mathbb{P}_{\mathcal{D}} \left[ \frac{1}{n} \left| \sum_{t=1}^n (f(x'_t) - \mathbb{E}_{t-1}[f(x'_t)]) \right| \leq \alpha/2 \mid x_1, \dots, x_n \right] \geq 1 - \frac{16}{n \alpha^2}$$

Whenever  $\alpha^2 \geq \frac{32}{n}$  we can conclude that

$$\inf_{f \in \mathcal{F}} \mathbb{P}_{\mathcal{D}} \left[ \frac{1}{n} \left| \sum_{t=1}^n (f(x'_t) - \mathbb{E}_{t-1}[f(x'_t)]) \right| \leq \alpha/2 \mid x_1, \dots, x_n \right] \geq \frac{1}{2}$$

Now given a fixed  $x_1, \dots, x_n$  let  $f^*$  be the function that maximizes  $\frac{1}{n} \left| \sum_{t=1}^n (f(x_t) - \mathbb{E}_{t-1}[f(x'_t)]) \right|$ . Note that  $f^*$  is a deterministic choice given  $x_1, \dots, x_n$ . Hence

$$\begin{aligned} \frac{1}{2} &\leq \inf_{f \in \mathcal{F}} \mathbb{P}_{\mathcal{D}} \left[ \frac{1}{n} \left| \sum_{t=1}^n (f(x'_t) - \mathbb{E}_{t-1}[f(x'_t)]) \right| \leq \alpha/2 \mid x_1, \dots, x_n \right] \\ &\leq \mathbb{P}_{\mathcal{D}} \left[ \frac{1}{n} \left| \sum_{t=1}^n (f^*(x'_t) - \mathbb{E}_{t-1}[f^*(x'_t)]) \right| \leq \alpha/2 \mid x_1, \dots, x_n \right] \end{aligned}$$

Let  $A = \{(x_1, \dots, x_n) \mid \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n f(x_t) - \mathbb{E}_{t-1}[f(x'_t)] \right| > \alpha\}$ . Since the above inequality holds for any  $x_1, \dots, x_n$  we can assert that

$$\frac{1}{2} \leq \mathbb{P}_{\mathcal{D}} \left[ \frac{1}{n} \left| \sum_{t=1}^n (f^*(x'_t) - \mathbb{E}_{t-1}[f^*(x'_t)]) \right| \leq \alpha/2 \mid (x_1, \dots, x_n) \in A \right]$$

Hence we conclude that

$$\begin{aligned}
& \frac{1}{2} \mathbb{P}_{\mathcal{D}} \left[ \sup_{f \in \mathcal{F}} \frac{1}{n} \left| \sum_{t=1}^n (f(x_t) - \mathbb{E}_{t-1}[f(x'_t)]) \right| > \alpha \right] \\
& \leq \mathbb{P}_{\mathcal{D}} \left[ \frac{1}{n} \left| \sum_{t=1}^n (f^*(x_t) - \mathbb{E}_{t-1}[f^*(x'_t)]) \right| \leq \alpha/2 \mid (x_1, \dots, x_n) \in A \right] \\
& \quad \times \mathbb{P}_{\mathcal{D}} \left[ \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n (f(x_t) - \mathbb{E}_{t-1}[f(x'_t)]) \right| > \alpha \right] \\
& \leq \mathbb{P}_{\mathcal{D}} \left[ \frac{1}{n} \left| \sum_{t=1}^n (f^*(x_t) - f^*(x'_t)) \right| > \alpha/2 \right] \\
& \leq \mathbb{P}_{\mathcal{D}} \left[ \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n (f(x_t) - f(x'_t)) \right| > \alpha/2 \right]
\end{aligned}$$

Now we apply Lemma 25 with  $\phi(u) = \mathbf{1}_{\{u > \alpha/2\}}$  and  $\Delta_f(x_t, x'_t) = f(x_t) - f(x'_t)$ ,

$$\begin{aligned}
& \mathbb{E} \left[ \mathbf{1}_{\{\sup_{f \in \mathcal{F}} |\sum_{t=1}^n f(x_t) - f(x'_t)| \geq \alpha/2\}} \right] \\
& \leq \sup_{x_1, x'_1} \left\{ \mathbb{E}_{\epsilon_1} \left[ \dots \sup_{x_n, x'_n} \left\{ \mathbb{E}_{\epsilon_n} \left[ \mathbf{1}_{\{\sup_{f \in \mathcal{F}} |\sum_{t=1}^n \epsilon_t (f(x_t) - f(x'_t))| \geq \alpha/2\}} \right] \right\} \dots \right] \right\} \quad (4.8.3)
\end{aligned}$$

The next few steps are similar to the proof of Theorem 24. Since

$$\sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n \epsilon_t (f(x_t) - f(x'_t)) \right| \leq \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n \epsilon_t f(x_t) \right| + \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n \epsilon_t f(x'_t) \right|$$

it is true that

$$\mathbf{1}_{\{\sup_{f \in \mathcal{F}} |\sum_{t=1}^n \epsilon_t (f(x_t) - f(x'_t))| \geq \alpha/2\}} \leq \mathbf{1}_{\{\sup_{f \in \mathcal{F}} |\sum_{t=1}^n \epsilon_t f(x_t)| \geq \alpha/4\}} + \mathbf{1}_{\{\sup_{f \in \mathcal{F}} |\sum_{t=1}^n \epsilon_t f(x'_t)| \geq \alpha/4\}}$$

The right-hand side of Eq. Eq. (4.8.3) then splits into two equal parts:

$$\begin{aligned}
& \sup_{x_1} \left\{ \mathbb{E}_{\epsilon_1} \left[ \dots \sup_{x_n} \left\{ \mathbb{E}_{\epsilon_n} \left[ \mathbf{1}_{\{\sup_{f \in \mathcal{F}} |\sum_{t=1}^n \epsilon_t f(x_t)| \geq \alpha/4\}} \right] \right\} \dots \right] \right\} \\
& + \sup_{x'_1} \left\{ \mathbb{E}_{\epsilon_1} \left[ \dots \sup_{x'_n} \left\{ \mathbb{E}_{\epsilon_n} \left[ \mathbf{1}_{\{\sup_{f \in \mathcal{F}} |\sum_{t=1}^n \epsilon_t f(x'_t)| \geq \alpha/4\}} \right] \right\} \dots \right] \right\} \\
& = 2 \sup_{x_1} \left\{ \mathbb{E}_{\epsilon_1} \left[ \dots \sup_{x_n} \left\{ \mathbb{E}_{\epsilon_n} \left[ \mathbf{1}_{\{\sup_{f \in \mathcal{F}} |\sum_{t=1}^n \epsilon_t f(x_t)| \geq \alpha/4\}} \right] \right\} \dots \right] \right\}
\end{aligned}$$

Moving to the tree representation,

$$\begin{aligned} \mathbb{P}_{\mathcal{D}} \left[ \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n (f(x_t) - f(x'_t)) \right| > \alpha/2 \right] &\leq 2 \sup_{\mathbf{z}} \mathbb{E}_{\epsilon} \left[ \mathbf{1}_{\left\{ \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n \epsilon_t f(\mathbf{z}_t(\epsilon)) \right| > \alpha/4 \right\}} \right] \\ &= 2 \sup_{\mathbf{z}} \mathbb{P}_{\epsilon} \left[ \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n \epsilon_t f(\mathbf{z}_t(\epsilon)) \right| > \alpha/4 \right] \end{aligned}$$

We can now conclude that

$$\mathbb{P}_{\mathcal{D}} \left[ \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n (f(x_t) - \mathbb{E}_{t-1}[f(x_t)]) \right| > \alpha \right] \leq 4 \sup_{\mathbf{z}} \mathbb{P}_{\epsilon} \left[ \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n \epsilon_t f(\mathbf{z}_t(\epsilon)) \right| > \alpha/4 \right]$$

Fix an  $\mathcal{Z}$ -valued tree  $\mathbf{z}$  of depth  $n$ . By assumption  $\text{fat}_{\alpha}^{\text{seq}}(\mathcal{F}) < \infty$  for any  $\alpha > 0$ . Let  $V$  be a minimum  $\ell_1$ -cover of  $\mathcal{F}$  over  $\mathbf{z}$  at scale  $\alpha/8$ . Corollary 32 ensures that

$$|V| = \mathcal{N}_1(\alpha/8, \mathcal{F}, \mathbf{z}) \leq \left( \frac{16en}{\alpha} \right)^{\text{fat}_{\frac{\alpha}{8}}^{\text{seq}}}$$

and for any  $f \in \mathcal{F}$  and  $\epsilon \in \{\pm 1\}^n$ , there exists  $\mathbf{v}[f, \epsilon] \in V$  such that

$$\frac{1}{n} \sum_{t=1}^n |f(\mathbf{z}_t(\epsilon)) - \mathbf{v}[f, \epsilon]_t(\epsilon)| \leq \alpha/8$$

on the given path  $\epsilon$ . Hence

$$\begin{aligned} \mathbb{P}_{\epsilon} \left[ \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n \epsilon_t f(\mathbf{z}_t(\epsilon)) \right| > \alpha/4 \right] &= \mathbb{P}_{\epsilon} \left[ \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n \epsilon_t (f(\mathbf{z}_t(\epsilon)) - \mathbf{v}[f, \epsilon]_t(\epsilon) + \mathbf{v}[f, \epsilon]_t(\epsilon)) \right| > \alpha/4 \right] \\ &\leq \mathbb{P}_{\epsilon} \left[ \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n \epsilon_t (f(\mathbf{z}_t(\epsilon)) - \mathbf{v}[f, \epsilon]_t(\epsilon)) \right| + \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n \epsilon_t \mathbf{v}[f, \epsilon]_t(\epsilon) \right| > \alpha/4 \right] \\ &\leq \mathbb{P}_{\epsilon} \left[ \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n \epsilon_t \mathbf{v}[f, \epsilon]_t(\epsilon) \right| > \alpha/8 \right] \end{aligned}$$

For fixed  $\epsilon = (\epsilon_1, \dots, \epsilon_n)$ ,

$$\frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n \epsilon_t \mathbf{v}[f, \epsilon]_t(\epsilon) \right| > \alpha/8 \implies \frac{1}{n} \max_{\mathbf{v} \in V} \left| \sum_{t=1}^n \epsilon_t \mathbf{v}_t(\epsilon) \right| > \alpha/8$$



and, therefore, for any  $z$ ,

$$\begin{aligned} \mathbb{P}_\epsilon \left[ \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n \epsilon_t f(z_t(\epsilon)) \right| > \alpha/4 \right] &\leq \mathbb{P}_\epsilon \left[ \frac{1}{n} \max_{\mathbf{v} \in V} \left| \sum_{t=1}^n \epsilon_t \mathbf{v}_t(\epsilon) \right| > \alpha/8 \right] \\ &\leq \sum_{\mathbf{v} \in V} \mathbb{P}_\epsilon \left[ \frac{1}{n} \left| \sum_{t=1}^n \epsilon_t \mathbf{v}_t(\epsilon) \right| > \alpha/8 \right] \leq 2|V|e^{-n\alpha^2/128} \leq 2 \left( \frac{16en}{\alpha} \right)^{\text{fat}_{\alpha/8}^{\text{seq}}} e^{-n\alpha^2/128} \end{aligned}$$

Hence we conclude that for any  $\mathcal{D}$

$$\mathbb{P}_{\mathcal{D}} \left[ \frac{1}{n} \sup_{f \in \mathcal{F}} \left| \sum_{t=1}^n (f(x_t) - \mathbb{E}_{t-1}[f(x_t)]) \right| > \alpha \right] \leq 8 \left( \frac{16en}{\alpha} \right)^{\text{fat}_{\alpha/8}^{\text{seq}}} e^{-n\alpha^2/128}$$

Now applying Borel-Cantelli lemma proves the required result as

$$\sum_{n=1}^{\infty} 8 \left( \frac{16en}{\alpha} \right)^{\text{fat}_{\alpha/8}^{\text{seq}}} e^{-n\alpha^2/128} < \infty.$$

□

**Proof of Proposition 42.** Using Theorem 35 we get the bound  $\mathfrak{R}_n^{\text{seq}}(\mathcal{F}_S) \leq \mathfrak{D}_n^{\text{seq}}(\mathcal{F}_S)$ . Further the fact that absolute loss is 1-Lipschitz implies that an  $\epsilon$  cover of class  $\mathcal{H}$  is also an  $\epsilon$  cover for loss class  $\mathcal{F}_S$  and so  $\mathfrak{D}_n^{\text{seq}}(\mathcal{F}_S) \leq \mathfrak{D}_n^{\text{seq}}(\mathcal{H})$ . This gives the first upper bound of value in terms of  $\mathfrak{D}_n^{\text{seq}}(\mathcal{H})$ . The second inequality in the upper bound is a direct consequence of using Corollary 32 in  $\mathfrak{D}_n^{\text{seq}}(\mathcal{H})$ . Now before we prove the final inequality in the upper bound we first prove the lower bound because we shall use ideas from the lower bound to get the final inequality in the upper bound.

For the lower bound, we use a construction similar to [35]. We construct a particular distribution which induces a lower bound on regret for any algorithm. For any  $\alpha \geq 0$  by definition of fat-shattering dimension, there exists a tree  $z$  of depth  $d = \text{fat}_{\alpha}^{\text{seq}}(\mathcal{H})$  that can be  $\alpha$ -shattered by  $\mathcal{H}$ . For simplicity, we assume  $n = kd$  where  $k$  is some non-negative integer, and the case  $n \leq d$  is discussed at the end of the proof. Now, define the  $j$ th block of time  $T_j = \{(j-1)k + 1, \dots, jk\}$ .

Now the strategy of Nature (Adversary) is to first pick  $\tilde{\epsilon} \in \{\pm 1\}^n$  independently and uniformly at random. Further let  $\epsilon \in \{\pm 1\}^d$  be defined as  $\epsilon_j = \text{sign} \left( \sum_{t \in T_j} \tilde{\epsilon}_t \right)$  for  $1 \leq j \leq d$ , the block-wise modal sign of  $\tilde{\epsilon}$ . Now note that by definition of  $\alpha$ -shattering, there exists a witness tree  $\mathbf{s}$  such that for any  $\epsilon \in \{\pm 1\}^d$  there exists  $\mathbf{h}_\epsilon \in \mathcal{H}$  with  $\epsilon_j(\mathbf{h}_\epsilon(z_j(\epsilon)) - \mathbf{s}_j(\epsilon)) \geq \alpha/2$  for all  $1 \leq j \leq d$ . Now let the random sequence  $(x_1, y_1), \dots, (x_n, y_n)$  be defined by  $x_t = z_j(\epsilon)$  for all  $t \in T_j$  and  $j \in \{1, \dots, d\}$  and  $y_t = \tilde{\epsilon}_t$ . In the remainder of the proof we show that any algorithm suffers large expected regret.

Now consider any player strategy (possibly randomized) making prediction  $\hat{y}_t \in [-1, 1]$  at round  $t$ . Note that if we consider block  $j$ ,  $y_t = \tilde{\epsilon}_t$  is  $\pm 1$  uniformly at random. This means that irrespective of what  $\hat{y}_t$  the player

plays, the expectation over  $\tilde{\epsilon}_t$  of the loss the player suffers at round  $t$  is

$$\mathbb{E}_{\tilde{\epsilon}_t} |\hat{y}_t - y_t| = 1$$

Hence on block  $j$ , the expected loss accumulated by any player is  $k$  and so for any player strategy (possibly randomized),

$$\mathbb{E} \left[ \sum_{t=1}^n |\hat{y}_t - y_t| \right] = \sum_{j=1}^d k = dk = n \quad (4.8.4)$$

On the other hand since  $x_t = z_j(\epsilon)$ , we know that there always exists a function for any  $\epsilon \in \{\pm 1\}^d$ , say  $\mathbf{h}_\epsilon$  such that  $\epsilon_j(\mathbf{h}_\epsilon(z_j(\epsilon)) - s_j(\epsilon)) \geq \alpha/2$ . Hence

$$\begin{aligned} \mathbb{E} \left[ \inf_{\mathbf{h} \in \mathcal{H}} \sum_{t=1}^n |\mathbf{h}(x_t) - y_t| \right] &\leq \sum_{j=1}^d \mathbb{E} \left[ \sum_{t \in T_j} |\mathbf{h}_\epsilon(x_t) - y_t| \right] \\ &= \sum_{j=1}^d \mathbb{E} \left[ \sum_{t \in T_j} |\mathbf{h}_\epsilon(z_j(\epsilon)) - y_t| \right] \\ &\leq \sum_{j=1}^d \mathbb{E} \left[ \max_{c_j \in [s_j(\epsilon) + \epsilon_j \frac{\alpha}{2}, \epsilon_j]} \sum_{t \in T_j} |c_j - y_t| \right] \end{aligned}$$

where the last step is because for all of block  $j$ ,  $\mathbf{h}_\epsilon(z_j(\epsilon))$  does not depend on  $t$  and lies in the interval<sup>1</sup>  $[s_j(\epsilon) + \epsilon_j \frac{\alpha}{2}, \epsilon_j]$  (i.e. the majority side) and so by replacing it by the maximal  $c_j$  in the same interval for that block we only make the quantity bigger. Now for a block  $j$ , define the number of labels that match the sign of  $\epsilon_j$  (the majority) as  $M_j = \sum_{t \in T_j} \mathbf{1}_{\{y_t = \epsilon_j\}}$ . Since  $y_t = \tilde{\epsilon}_t \in \{\pm 1\}$ , observe that the function  $g(c_j) = \sum_{t \in T_j} |c_j - y_t|$  is linear on the interval  $[-1, 1]$  with its minimum at the majority sign  $\epsilon_j$ . Hence, the maximum over  $[s_j(\epsilon) + \epsilon_j \frac{\alpha}{2}, \epsilon_j]$  must occur at  $c_j = s_j(\epsilon) + \epsilon_j \frac{\alpha}{2}$ . Substituting,

$$\begin{aligned} \max_{c_j \in [s_j(\epsilon) + \epsilon_j \frac{\alpha}{2}, \epsilon_j]} \sum_{t \in T_j} |c_j - y_t| &= M_j \left| s_j(\epsilon) + \epsilon_j \frac{\alpha}{2} - \epsilon_j \right| + (k - M_j) \left| s_j(\epsilon) + \epsilon_j \frac{\alpha}{2} + \epsilon_j \right| \\ &= M_j \left| \epsilon_j s_j(\epsilon) + \frac{\alpha}{2} - 1 \right| + (k - M_j) \left| \epsilon_j s_j(\epsilon) + \frac{\alpha}{2} + 1 \right| \\ &= M_j \left( 1 - \epsilon_j s_j(\epsilon) - \frac{\alpha}{2} \right) + (k - M_j) \left( 1 + \epsilon_j s_j(\epsilon) + \frac{\alpha}{2} \right) \\ &= k + (k - 2M_j) \left( \epsilon_j s_j(\epsilon) + \frac{\alpha}{2} \right) \end{aligned}$$

<sup>1</sup>We use the convention that  $[a, b]$  stands for  $[b, a]$  whenever  $a > b$ .

Hence,

$$\begin{aligned} \mathbb{E} \left[ \inf_{\mathbf{h} \in \mathcal{H}} \sum_{t=1}^n |\mathbf{h}(x_t) - y_t| \right] &\leq dk + \sum_{j=1}^d \mathbb{E} \left[ \epsilon_j \mathbf{s}_j(\epsilon)(k - 2M_j) + \frac{\alpha}{2}(k - 2M_j) \right] \\ &= dk + \sum_{j=1}^d \mathbb{E} [\epsilon_j \mathbf{s}_j(\epsilon)(k - 2M_j)] + \frac{\alpha}{2} \sum_{j=1}^d \mathbb{E} [k - 2M_j] \end{aligned}$$

Further note that  $k - 2M_j = -|\sum_{t \in T_j} \tilde{\epsilon}_t|$  and so  $\epsilon_j(k - 2M_j) = -\sum_{t \in T_j} \tilde{\epsilon}_t$  and so the expectation

$$\mathbb{E} [\epsilon_j \mathbf{s}_j(\epsilon)(k - 2M_j)] = \mathbb{E} [\mathbb{E}_{\tilde{\epsilon}_{k(j-1)+1:j k}} [\epsilon_j \mathbf{s}_j(\epsilon)(k - 2M_j)]] = 0$$

because  $\mathbf{s}_j(\epsilon)$  is independent of  $\tilde{\epsilon}_t$  for  $t \in T_j$ . Hence we see that

$$\mathbb{E} \left[ \inf_{\mathbf{h} \in \mathcal{H}} \sum_{t=1}^n |\mathbf{h}(x_t) - y_t| \right] \leq dk + \frac{\alpha}{2} \sum_{j=1}^d \mathbb{E} [k - 2M_j] \quad (4.8.5)$$

Combining Equations Eq. (4.8.4) and Eq. (4.8.5) we can conclude that for any player strategy,

$$\begin{aligned} \mathbb{E} \left[ \sum_{t=1}^n |\hat{y}_t - y_t| \right] - \mathbb{E} \left[ \inf_{\mathbf{h} \in \mathcal{H}} \sum_{t=1}^n |\mathbf{h}(x_t) - y_t| \right] &\geq \frac{\alpha}{2} \sum_{j=1}^d \mathbb{E} [2M_j - k] \\ &= \frac{\alpha}{2} \mathbb{E} \left[ \sum_{j=1}^d \left| \sum_{t \in T_j} \tilde{\epsilon}_t \right| \right] = \frac{\alpha}{2} \sum_{j=1}^d \mathbb{E} \left[ \sum_{t \in T_j} \tilde{\epsilon}_t \right] \geq \frac{\alpha d}{2} \sqrt{\frac{k}{2}} = \alpha \sqrt{\frac{nd}{8}} = \alpha \sqrt{\frac{n \text{fat}_\alpha^{\text{seq}}}{8}} \quad (4.8.6) \end{aligned}$$

by Khinchine's inequality (e.g. [52, Lemma A.9]), yielding the theorem statement for  $n \geq \text{fat}_\alpha^{\text{seq}}$ . For the case of  $n < \text{fat}_\alpha^{\text{seq}}$ , the proof is the same with  $k = 1$  and the depth of the shattered tree being  $n$ , yielding a lower bound of  $\alpha n / \sqrt{8}$ . Dividing throughout by  $n$  completes the lower bound.

Now we move to the final inequality in the upper bound, but before we proceed notice that since  $y_t$  are Rademacher random variables. Hence from Equation 4.8.6 we see that

$$\begin{aligned} \alpha \sqrt{\frac{\text{fat}_\alpha^{\text{seq}}}{8n}} &\leq \mathbb{E} \left[ \frac{1}{n} \sum_{t=1}^n |\hat{y}_t - y_t| \right] - \mathbb{E} \left[ \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n |\mathbf{h}(x_t) - y_t| \right] \\ &= \mathbb{E} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n (1 - |\mathbf{h}(x_t) - y_t|) \right] \\ &= \mathbb{E} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \tilde{\epsilon}_t \mathbf{h}(x_t) \right] \\ &= \mathbb{E} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \tilde{\epsilon}_t \mathbf{h}(z'_t(\tilde{\epsilon})) \right] \end{aligned}$$

Where  $z'_t(\tilde{\epsilon}) = z_{\lceil \frac{t}{k} \rceil}(\epsilon)$  where each  $\epsilon_j = \text{sign} \left( \sum_{t \in T_j} \tilde{\epsilon}_t \right)$  (ie.  $x_t$ 's can be seen as nodes of a tree formed by taking the  $z$  tree which is of depth  $d$  and making it into a depth  $d$  tree by expanding each node of the tree into a subtree of depth  $k$ ). Hence we conclude that :

$$\begin{aligned} \alpha \sqrt{\frac{\text{fat}_\alpha^{\text{seq}}}{8n}} &\leq \mathbb{E} \left[ \frac{1}{n} \sum_{t=1}^n |\hat{y}_t - y_t| \right] - \mathbb{E} \left[ \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n |\mathbf{h}(x_t) - y_t| \right] \\ &= \mathbb{E} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \tilde{\epsilon}_t \mathbf{h}(z'_t(\tilde{\epsilon})) \right] \\ &\leq \sup_{\mathbf{z}} \mathbb{E}_\epsilon \left[ \sup_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \epsilon_t \mathbf{h}(z_t(\epsilon)) \right] = \mathfrak{R}_n^{\text{seq}}(\mathcal{H}) \end{aligned}$$

Thus effectively we have shown that  $\sup_{\alpha} \alpha \sqrt{\frac{\min\{\text{fat}_\alpha^{\text{seq}}, n\}}{8n}} \leq \mathfrak{R}_n^{\text{seq}}(\mathcal{H})$ . Using this we see that if  $\hat{\alpha}$  is the solution to  $\text{fat}_{\hat{\alpha}}^{\text{seq}} = n$  then we have that  $\hat{\alpha} \leq \mathfrak{R}_n^{\text{seq}}(\mathcal{H})$  and for any  $\beta > \hat{\alpha}$ ,  $\sqrt{\frac{\text{fat}_\beta^{\text{seq}}}{n}} \leq \frac{2\sqrt{2} \mathfrak{R}_n^{\text{seq}}(\mathcal{H})}{\beta}$ . Hence using this we conclude that,

$$\begin{aligned} \inf_{\alpha} \left\{ 4\alpha + \frac{12}{\sqrt{n}} \int_{\alpha}^1 \sqrt{\text{fat}_\beta^{\text{seq}}(\mathcal{H}) \log\left(\frac{2en}{\beta}\right)} d\beta \right\} &\leq 4\hat{\alpha} + 12 \int_{\hat{\alpha}}^1 \frac{2\sqrt{2} \mathfrak{R}_n^{\text{seq}}(\mathcal{H}) \sqrt{\log\left(\frac{2en}{\beta}\right)}}{\beta} d\beta \\ &\leq 4\mathfrak{R}_n^{\text{seq}}(\mathcal{H}) + 36\sqrt{2} \mathfrak{R}_n^{\text{seq}}(\mathcal{H}) \sqrt{\log(n)} \int_{\hat{\alpha}}^1 \frac{1}{\beta} d\beta \\ &\leq 4\mathfrak{R}_n^{\text{seq}}(\mathcal{H}) + 36\sqrt{2} \mathfrak{R}_n^{\text{seq}}(\mathcal{H}) \log^{\frac{3}{2}}(n) \\ &\leq 58 \mathfrak{R}_n^{\text{seq}}(\mathcal{H}) \log^{\frac{3}{2}}(n) \end{aligned}$$

This concludes the upper bound. To show that  $\mathfrak{R}_n^{\text{seq}}(\mathcal{H}) \leq \mathcal{V}_n^S(\mathcal{H})$ , consider the adversary strategy where adversary picks three  $\mathbf{x}$  of depth  $n$ . Now at round 1 the adversary presents as instance  $x_1 = \mathbf{x}_1$  and  $y_1 = \epsilon_1$  where  $\epsilon_1$  is a Rademacher random variable. Now at round 2 the adversary presents instance  $x_2 = \mathbf{x}_2(\epsilon_1)$  and  $y_2 = \epsilon_2$  where again  $\epsilon_2$  is a Rademacher random variable. In a similar fashion at round  $t$  adversary presents instance  $x_t = \mathbf{x}_t(\epsilon_{1:t-1})$  and  $y_t = \epsilon_t$ . Therefore we see that

$$\begin{aligned} \mathcal{V}_n^S(\mathcal{H}) &\geq \mathbb{E}_\epsilon \left[ \frac{1}{n} \sum_{t=1}^n |\hat{y}_t - \epsilon_t| \right] - \mathbb{E}_\epsilon \left[ \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n |\mathbf{h}(\mathbf{x}_t(\epsilon)) - \epsilon_t| \right] \\ &= 1 - \mathbb{E}_\epsilon \left[ \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n |\mathbf{h}(\mathbf{x}_t(\epsilon)) - \epsilon_t| \right] \\ &= \mathbb{E}_\epsilon \left[ \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n (1 - |\mathbf{h}(\mathbf{x}_t(\epsilon)) - \epsilon_t|) \right] \\ &= \mathbb{E}_\epsilon \left[ \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \epsilon_t \mathbf{h}(\mathbf{x}_t(\epsilon)) \right] \end{aligned}$$

Since choice of tree  $x$  is arbitrary we conclude that  $\mathcal{V}_n^S(\mathcal{H}) \geq \mathfrak{N}_n^{\text{seq}}(\mathcal{H})$

□

**Proof of Lemma 45.** First, we claim that for any  $x \in \mathcal{X}$ ,  $\text{fat}_\alpha^{\text{seq}}(V_t(r, x)) = \text{fat}_\alpha^{\text{seq}}(V_t)$  for at most two  $r, r' \in B_\alpha$ .<sup>2</sup> Further if there are two such  $r, r' \in B_\alpha$  then  $r, r'$  are consecutive elements of  $B_\alpha$  (i.e.  $|r - r'| = \alpha$ ). Suppose, for the sake of contradiction, that  $\text{fat}_\alpha^{\text{seq}}(V_t(r, x)) = \text{fat}_\alpha^{\text{seq}}(V_t(r', x)) = \text{fat}_\alpha^{\text{seq}}(V_t)$  for distinct  $r, r' \in B_\alpha$  that are not consecutive (i.e.  $|r - r'| \geq 2\alpha$ ). Then let  $s = (r + r')/2$  and without loss of generality suppose  $r > r'$ . By definition for any  $\mathbf{h} \in V_t(r, x)$ ,

$$\mathbf{h}(x) \geq r - \alpha/2 = (r' + r)/2 + (r - r')/2 - \alpha/2 \geq s + \alpha/2$$

Also for any  $\mathbf{h}' \in V_t(r', x)$  we also have,

$$\mathbf{h}'(x) \leq r' + \alpha/2 = (r' + r)/2 + (r' - r)/2 + \alpha/2 \leq s - \alpha/2$$

Let  $\mathbf{v}$  and  $\mathbf{v}'$  be trees of depth  $\text{fat}_\alpha^{\text{seq}}(V_t)$   $\alpha$ -shattered by  $V_t(r, x)$  and  $V_t(r', x)$ , respectively. To get a contradiction, form a new tree  $\mathbf{v}''$  of depth  $\text{fat}_\alpha^{\text{seq}}(V_t) + 1$  by joining trees  $\mathbf{v}$  and  $\mathbf{v}'$  with the constant function  $\mathbf{v}''_1 = x$  as the root and  $\mathbf{v}$  and  $\mathbf{v}'$  as the left and right subtrees respectively. It is straightforward that this tree is shattered by  $V_t(r, x) \cup V_t(r', x)$ , a contradiction.

Notice that the times  $t \in [n]$  for which  $|\mathbf{h}_t(x_t) - y_t| > \alpha$  are exactly those times when we update current set  $V_{t+1}$ . We shall show that whenever an update is made,  $\text{fat}_\alpha^{\text{seq}}(V_{t+1}) < \text{fat}_\alpha^{\text{seq}}(V_t)$  and hence claim that the total number of times  $|\mathbf{h}_t(x_t) - y_t| > \alpha$  is bounded by  $\text{fat}_\alpha^{\text{seq}}(\mathcal{F})$ .

At any round we have three possibilities. First is when  $\text{fat}_\alpha^{\text{seq}}(V_t(r, x_t)) < \text{fat}_\alpha^{\text{seq}}(V_t)$  for all  $r \in B_\alpha$ . In this case, clearly, an update results in  $\text{fat}_\alpha^{\text{seq}}(V_{t+1}) = \text{fat}_\alpha^{\text{seq}}(V_t(\lfloor y_t \rfloor_\alpha, x_t)) < \text{fat}_\alpha^{\text{seq}}(V_t)$ .

The second case is when  $\text{fat}_\alpha^{\text{seq}}(V_t(r, x_t)) = \text{fat}_\alpha^{\text{seq}}(V_t)$  for exactly one  $r \in B_\alpha$ . In this case the algorithm chooses  $\mathbf{h}_t(x_t) = r$ . If the update is made,  $|\mathbf{h}_t(x_t) - y_t| > \alpha$  and thus  $\lfloor y_t \rfloor_\alpha \neq \mathbf{h}_t(x_t)$ . We can conclude that

$$\text{fat}_\alpha^{\text{seq}}(V_{t+1}) = \text{fat}_\alpha^{\text{seq}}(V_t(\lfloor y_t \rfloor_\alpha, x_t)) < \text{fat}_\alpha^{\text{seq}}(V_t(\mathbf{h}_t(x_t), x_t)) = \text{fat}_\alpha^{\text{seq}}(V_t)$$

The final case is when  $\text{fat}_\alpha^{\text{seq}}(V_t(r, x_t)) = \text{fat}_\alpha^{\text{seq}}(V_t(r', x_t)) = \text{fat}_\alpha^{\text{seq}}(V_t)$  and  $|r - r'| = \alpha$ . In this case, the algorithm chooses  $\mathbf{h}_t(x_t) = \frac{r+r'}{2}$ . Whenever  $y_t$  falls in either of these two consecutive intervals given by  $r$  or  $r'$ , we have  $|\mathbf{h}_t(x_t) - y_t| \leq \alpha$ , and hence no update is made. Thus, if an update is made,  $\lfloor y_t \rfloor_\alpha \neq r$  and  $\lfloor y_t \rfloor_\alpha \neq r'$ . However, for any element of  $B_\alpha$  other than  $r, r'$ , the fat shattering dimension is less than that of  $V_t$ . That is

$$\text{fat}_\alpha^{\text{seq}}(V_{t+1}) = \text{fat}_\alpha^{\text{seq}}(V_t(\lfloor y_t \rfloor_\alpha, x_t)) < \text{fat}_\alpha^{\text{seq}}(V_t(r, x_t)) = \text{fat}_\alpha^{\text{seq}}(V_t(r', x_t)) = \text{fat}_\alpha^{\text{seq}}(V_t).$$

<sup>2</sup>The argument should be compared to the combinatorial argument in Theorem 31.

We conclude that whenever we update,  $\text{fat}_\alpha^{\text{seq}}(V_{t+1}) < \text{fat}_\alpha^{\text{seq}}(V_t)$ , and so we can conclude that algorithm's prediction is more than  $\alpha$  away from  $y_t$  on at most  $\text{fat}_\alpha^{\text{seq}}(\mathcal{F})$  number of rounds.  $\square$

**Proof of Corollary 48.** For the choice of weights  $p_i = \frac{6}{\pi^2} i^{-2}$  we see from Proposition 53 that for any  $i$ ,

$$\mathbb{E}[\mathbf{R}_n] \leq \alpha_i + \sqrt{\frac{\text{fat}_{\alpha_i}^{\text{seq}} \log\left(\frac{2n}{\alpha_i}\right)}{n}} + \frac{1}{\sqrt{n}} (3 + 2 \log(i))$$

Now for any  $\alpha > 0$  let  $i_\alpha$  be such that  $\alpha \leq 2^{-i_\alpha}$  and for any  $i < i_\alpha$ ,  $\alpha > 2^{-i_\alpha}$ . Using the above bound on expected regret we have that

$$\mathbb{E}[\mathbf{R}_n] \leq \alpha_{i_\alpha} + \sqrt{\frac{\text{fat}_{\alpha_{i_\alpha}}^{\text{seq}} \log\left(\frac{2n}{\alpha_{i_\alpha}}\right)}{n}} + \frac{1}{\sqrt{n}} (3 + 2 \log(i_\alpha))$$

However for our choice of  $i_\alpha$  we see that  $i_\alpha \leq \log(1/\alpha)$  and further  $\alpha_{i_\alpha} \leq \alpha$ . Hence we conclude that

$$\mathbb{E}[\mathbf{R}_n] \leq \alpha + \sqrt{\frac{\text{fat}_\alpha^{\text{seq}} \log\left(\frac{2n}{\alpha}\right)}{n}} + \frac{1}{\sqrt{n}} \left(3 + 2 \log \log\left(\frac{1}{\alpha}\right)\right)$$

Since choice of  $\alpha$  was arbitrary we take infimum and get the result.  $\square$

**Proof of Proposition 49.** Fix a  $\gamma > 0$  and use loss

$$\ell(\hat{y}, y) = \begin{cases} 1 & \hat{y}y \leq 0 \\ 1 - \hat{y}y/\gamma & 0 < \hat{y}y < \gamma \\ 0 & \hat{y}y \geq \gamma \end{cases}$$

First note that since the loss is  $1/\gamma$ -Lipschitz, we can use Theorem 24 and the Rademacher contraction Lemma 27 to show that for each  $\gamma > 0$  there exists a randomized strategy  $\mathbf{A}^\gamma$  such that

$$\mathbb{E} \left[ \sum_{t=1}^n \mathbb{E}_{\mathbf{h}_t \sim \mathbf{A}_t^\gamma(z_{1:t-1})} [\ell(\mathbf{h}_t(x_t), y_t)] \right] \leq \inf_{\mathbf{h} \in \mathcal{H}} \sum_{t=1}^n \ell(\mathbf{h}(x_t), y_t) + \frac{2}{\gamma} \mathfrak{R}_n^{\text{seq}}(\mathcal{H})$$

Now note that the loss is lower bounded by the Zero-one loss  $\mathbf{1}_{\{\hat{y}y < 0\}}$  and is upper bounded by the margin Zero-one loss  $\mathbf{1}_{\{\hat{y}y < \gamma\}}$ . Hence we see that for this strategy,

$$\mathbb{E} \left[ \sum_{t=1}^n \mathbb{E}_{\mathbf{h}_t \sim \mathbf{A}_t^\gamma(z_{1:t-1})} [\mathbf{1}_{\{\mathbf{h}_t(x_t)y_t < 0\}}] \right] \leq \inf_{\mathbf{h} \in \mathcal{H}} \sum_{t=1}^n \mathbf{1}_{\{\mathbf{h}(x_t)y_t < \gamma\}} + \frac{2}{\gamma} \mathfrak{R}_n^{\text{seq}}(\mathcal{H}) \quad (4.8.7)$$

Hence for each fixed  $\gamma$  for randomized strategy given by  $\mathbf{A}^\gamma$  we have the above bound. Now we discretize over  $\gamma$ 's as  $\gamma_i = 1/2^i$  and using the output of the randomized strategies  $\mathbf{A}^{\gamma_1}, \mathbf{A}^{\gamma_2}, \dots$  that attain the regret

bounds given in Eq. (4.8.7) as experts and running experts algorithm given in Algorithm 3 with initial weight for expert  $i$  as  $p_i = \frac{6}{\pi^2 i^2}$  then using Proposition 53 we get that for this randomized strategy  $\mathbf{A}$ , such that for any  $i$

$$\mathbb{E} \left[ \sum_{t=1}^n \mathbb{E}_{\mathbf{h}_t \sim \mathbf{A}(z_{1:t-1})} [\mathbf{1}_{\{\mathbf{h}_t(x_t)y_t < 0\}}] \right] \leq \inf_{\mathbf{h} \in \mathcal{H}} \sum_{t=1}^n \mathbf{1}_{\{\mathbf{h}(x_t)y_t < \gamma_i\}} + \frac{2}{\gamma_i} \mathfrak{R}_n^{\text{seq}}(\mathcal{H}) + \sqrt{n} \left( 1 + 2 \log \left( \frac{i\pi}{\sqrt{6}} \right) \right)$$

Now for any  $\gamma > 0$  let  $i_\gamma$  be such that  $\gamma \leq 2^{-i_\gamma}$  and for any  $i < i_\gamma$ ,  $\gamma > 2^{-i_\gamma}$ . Then using the above bound we see that

$$\mathbb{E} \left[ \sum_{t=1}^n \mathbb{E}_{\mathbf{h}_t \sim \mathbf{A}(z_{1:t-1})} [\mathbf{1}_{\{\mathbf{h}_t(x_t)y_t < 0\}}] \right] \leq \inf_{\mathbf{h} \in \mathcal{H}} \sum_{t=1}^n \mathbf{1}_{\{\mathbf{h}(x_t)y_t < 2\gamma\}} + \frac{2}{\gamma} \mathfrak{R}_n^{\text{seq}}(\mathcal{H}) + \sqrt{n} \left( 1 + 2 \log \left( \frac{i\pi}{\sqrt{6}} \right) \right)$$

However note that  $i_\gamma \leq \log(1/\gamma)$  and so we can conclude that

$$\mathbb{E} \left[ \sum_{t=1}^n \mathbb{E}_{\mathbf{h}_t \sim \mathbf{A}(z_{1:t-1})} [\mathbf{1}_{\{\mathbf{h}_t(x_t)y_t < 0\}}] \right] \leq \inf_{\mathbf{h} \in \mathcal{H}} \sum_{t=1}^n \mathbf{1}_{\{\mathbf{h}(x_t)y_t < 2\gamma\}} + \frac{2}{\gamma} \mathfrak{R}_n^{\text{seq}}(\mathcal{H}) + \sqrt{n} \left( 1 + 2 \log \left( \frac{\pi \log(1/\gamma)}{\sqrt{6}} \right) \right)$$

Dividing throughout by  $n$  concludes the proof.  $\square$

**Proof of Proposition 50.** We shall prove that for any  $i \in [k]$ ,

$$\mathfrak{R}_n^{\text{seq}}(\mathcal{F}_i) \leq 2LB_i \mathfrak{R}_n^{\text{seq}}(\mathcal{F}_{i-1})$$

To see this note that

$$\begin{aligned}
\mathfrak{R}_n^{\text{seq}}(\mathcal{F}_i) &= \frac{1}{n} \sup_{\mathbf{z}} \mathbb{E}_{\epsilon} \left[ \sup_{\substack{w^i: \|w^i\|_1 \leq B_i \\ \forall j f_j \in \mathcal{F}_{i-1}}} \sum_{t=1}^n \epsilon_t \left( \sum_j w_j^i \sigma(f_j(\mathbf{z}_t(\epsilon))) \right) \right] \\
&\leq \frac{1}{n} \sup_{\mathbf{z}} \mathbb{E}_{\epsilon} \left[ \sup_{\substack{w^i: \|w^i\|_1 \leq B_i \\ \forall j f_j \in \mathcal{F}_{i-1}}} \|w^i\|_1 \max_j \left| \sum_{t=1}^n \epsilon_t \sigma(f_j(\mathbf{z}_t(\epsilon))) \right| \right] && \text{(Hölder's inequality)} \\
&\leq \frac{1}{n} \sup_{\mathbf{z}} \mathbb{E}_{\epsilon} \left[ B_i \sup_{f \in \mathcal{F}_{i-1}} \left| \sum_{t=1}^n \epsilon_t \sigma(f(\mathbf{z}_t(\epsilon))) \right| \right] \\
&= \frac{1}{n} \sup_{\mathbf{z}} \mathbb{E}_{\epsilon} \left[ B_i \sup_{f \in \mathcal{F}_{i-1}} \max \left\{ \sum_{t=1}^n \epsilon_t \sigma(f(\mathbf{z}_t(\epsilon))), -\sum_{t=1}^n \epsilon_t \sigma(f(\mathbf{z}_t(\epsilon))) \right\} \right] \\
&= \frac{1}{n} \sup_{\mathbf{z}} \mathbb{E}_{\epsilon} \left[ B_i \max \left\{ \sup_{f \in \mathcal{F}_{i-1}} \sum_{t=1}^n \epsilon_t \sigma(f(\mathbf{z}_t(\epsilon))), \sup_{f \in \mathcal{F}_{i-1}} \sum_{t=1}^n -\epsilon_t \sigma(f(\mathbf{z}_t(\epsilon))) \right\} \right] \\
&\leq \frac{1}{n} \sup_{\mathbf{z}} \mathbb{E}_{\epsilon} \left[ B_i \sup_{f \in \mathcal{F}_{i-1}} \sum_{t=1}^n \epsilon_t \sigma(f(\mathbf{z}_t(\epsilon))) \right] + \sup_{\mathbf{z}} \mathbb{E}_{\epsilon} \left[ B_i \sup_{f \in \mathcal{F}_{i-1}} \sum_{t=1}^n -\epsilon_t \sigma(f(\mathbf{z}_t(\epsilon))) \right] && (\sigma(0) = 0 \text{ and } 0 \in \mathcal{F}_i) \\
&= \frac{2B_i}{n} \sup_{\mathbf{z}} \mathbb{E}_{\epsilon} \left[ \sup_{f \in \mathcal{F}_{i-1}} \sum_{t=1}^n \epsilon_t \sigma(f(\mathbf{z}_t(\epsilon))) \right] && \text{(Proposition 29)} \\
&\leq \frac{2B_i L}{n} \sup_{\mathbf{z}} \mathbb{E}_{\epsilon} \left[ \sup_{f \in \mathcal{F}_{i-1}} \sum_{t=1}^n \epsilon_t f(\mathbf{z}_t(\epsilon)) \right] && \text{(Lemma 27)} \\
&= 2B_i L \mathfrak{R}_n^{\text{seq}}(\mathcal{F}_{i-1}) && (4.8.8)
\end{aligned}$$

To finish the proof we note that

$$\begin{aligned}
\mathfrak{R}_n^{\text{seq}}(\mathcal{F}_1) &= \sup_{\mathbf{z}} \mathbb{E}_{\epsilon} \left[ \sup_{w \in \mathbb{R}^d: \|w\|_1 \leq B_1} \frac{1}{n} \sum_{t=1}^n \epsilon_t w^\top \mathbf{z}_t(\epsilon) \right] \\
&\leq \sup_{\mathbf{z}} \mathbb{E}_{\epsilon} \left[ \sup_{w \in \mathbb{R}^d: \|w\|_1 \leq B_1} \|w\|_1 \left\| \frac{1}{n} \sum_{t=1}^n \epsilon_t \mathbf{z}_t(\epsilon) \right\|_{\infty} \right] \\
&\leq B_1 \sup_{\mathbf{z}} \mathbb{E}_{\epsilon} \left[ \max_{i \in [d]} \left\{ \frac{1}{n} \sum_{t=1}^n \epsilon_t \mathbf{z}_t(\epsilon)[i] \right\} \right]
\end{aligned}$$

Note that the instances  $x \in \mathcal{Z}$  are vectors in  $\mathbb{R}^d$  and so for a given instance tree  $\mathbf{z}$ , for any  $i \in [d]$ ,  $\mathbf{z}[i]$  given by only taking the  $i^{\text{th}}$  co-ordinate is a valid real valued tree. Hence using Lemma 34 we conclude that

$$\begin{aligned}
\mathfrak{R}_n^{\text{seq}}(\mathcal{F}_1) &\leq B_1 \sup_{\mathbf{z}} \mathbb{E}_{\epsilon} \left[ \max_{i \in [d]} \left\{ \frac{1}{n} \sum_{t=1}^n \epsilon_t \mathbf{z}_t(\epsilon)[i] \right\} \right] \\
&\leq B_1 \sqrt{\frac{2X_{\infty}^2 \log d}{n}}
\end{aligned}$$



Using the above and Equation 4.8.8 we conclude the proof.  $\square$

**Proof of Proposition 51.** For a tree of depth  $d$ , the indicator function of a leaf is a conjunction of no more than  $d$  decision functions. More specifically, if the decision tree consists of decision nodes chosen from a class  $\mathcal{C}$  of binary-valued functions, the indicator function of leaf  $l$  (which takes value 1 at a point  $x$  if  $x$  reaches  $l$ , and 0 otherwise) is a conjunction of  $d_l$  functions from  $\mathcal{C}$ , where  $d_l$  is the depth of leaf  $l$ . We can represent the function computed by the tree as the sign of

$$g(x) = \sum_l w_l \sigma_l \bigwedge_{i=1}^{d_l} c_{l,i}(x)$$

where the sum is over all leaves  $l$ ,  $w_l > 0$ ,  $\sum_l w_l = 1$ ,  $\sigma_l \in \{\pm 1\}$  is the label of leaf  $l$ ,  $c_{l,i} \in \mathcal{C}$ , and the conjunction is understood to map to  $\{0, 1\}$ . Now note that if we fix some  $L > 0$  then we see that the loss

$$\phi_L(\alpha) = \begin{cases} 1 & \text{if } \alpha \leq 0 \\ 1 - L\alpha & \text{if } 0 < \alpha \leq 1/L \\ 0 & \text{otherwise} \end{cases}$$

is  $L$ -Lipschitz and so by Theorem 24 and Lemma 27 we have that for every  $L > 0$ , there exists a randomized strategy  $\mathbf{A}^L$  for the player, such that for any sequence  $z_1 = (x_1, y_1), \dots, z_n = (x_n, y_n)$ ,

$$\mathbb{E} \left[ \sum_{t=1}^n \mathbb{E}_{\tau_t \sim \mathbf{A}^L(z_{1:t-1})} [\phi_L(y_t \tau_t(x_t))] \right] \leq \inf_{\tau \in \mathcal{T}} \sum_{t=1}^n \phi_L(y_t \tau(x_t)) + L \mathfrak{R}_n^{\text{seq}}(\mathcal{T})$$

Now note that  $\phi_L$  upper bounds the step function and so

$$\mathbb{E} \left[ \sum_{t=1}^n \mathbb{E}_{\tau_t \sim \mathbf{A}^L(z_{1:t-1})} [\mathbf{1}_{\{\tau_t(x_t) \neq y_t\}}] \right] \leq \inf_{\tau \in \mathcal{T}} \sum_{t=1}^n \phi_L(y_t \tau(x_t)) + L \mathfrak{R}_n^{\text{seq}}(\mathcal{T})$$

Now say  $\tau^* \in \mathcal{T}$  is the minimizer of  $\sum_{t=1}^n \mathbf{1}_{\{\tau(x_t) \neq y_t\}}$  then note that

$$\begin{aligned} \sum_{t=1}^n \phi_L(y_t \tau^*(x_t)) &= \sum_{t=1}^n \mathbf{1}_{\{\tau^*(x_t) \neq y_t\}} + \sum_l \tilde{C}_n(l) \phi_L(w_l) \\ &\leq \sum_{t=1}^n \mathbf{1}_{\{\tau^*(x_t) \neq y_t\}} + \sum_l \tilde{C}_n(l) \max(0, 1 - Lw_l) \\ &\leq \sum_{t=1}^n \mathbf{1}_{\{\tau^*(x_t) \neq y_t\}} + \sum_l \max(0, (1 - Lw_l) \tilde{C}_n(l)) \\ &= \inf_{\tau \in \mathcal{T}} \sum_{t=1}^n \mathbf{1}_{\{\tau(x_t) \neq y_t\}} + \sum_l \max(0, (1 - Lw_l) \tilde{C}_n(l)) \end{aligned}$$

Hence we see that

$$\mathbb{E} \left[ \sum_{t=1}^n \mathbb{E}_{\tau_t \sim \mathbf{A}^L(z_{1:t-1})} [\mathbf{1}_{\{\tau_t(x_t) \neq y_t\}}] \right] \leq \inf_{\tau \in \mathcal{T}} \sum_{t=1}^n \mathbf{1}_{\{\tau(x_t) \neq y_t\}} + \sum_l \max \left( 0, (1 - Lw_l) \tilde{C}_n(l) \right)$$

Now if we discretize over  $L$  as  $L_i = i$  for all  $i \in \mathbb{N}$  and run experts algorithm 3 with output of randomized strategies,  $\mathbf{A}^{L_1}, \mathbf{A}^{L_2}, \dots$  as our experts and weight of expert  $i$  with  $p_i = \frac{6}{\pi^2} i^{-2}$  so that  $\sum_i p_i = 1$  then we get that for this randomized strategy  $\mathbf{A}$ , we have from Proposition 53 that for all  $L \in \mathbb{N}$ ,

$$\begin{aligned} & \mathbb{E} \left[ \sum_{t=1}^n \mathbb{E}_{\tau_t \sim \mathbf{A}(z_{1:t-1})} [\mathbf{1}_{\{\tau_t(x_t) \neq y_t\}}] \right] \\ & \leq \inf_{\tau \in \mathcal{T}} \sum_{t=1}^n \mathbf{1}_{\{\tau(x_t) \neq y_t\}} + \sum_l \max \left( 0, (1 - Lw_l) \tilde{C}_n(l) \right) + L \mathfrak{R}_n^{\text{seq}}(\mathcal{T}) + \sqrt{n} + 2\sqrt{n} \log(L\pi/\sqrt{6}) \end{aligned}$$

Now we pick  $L = |\{l : \tilde{C}_n(l) > 2\mathfrak{R}_n^{\text{seq}}(\mathcal{T})\}| =: N_{\text{leaf}}$  and also pick  $w_l = 0$  if  $\tilde{C}_n(l) \leq 2\mathfrak{R}_n^{\text{seq}}(\mathcal{T})$  and  $w_l = 1/L$  otherwise. Hence we see that

$$\begin{aligned} \mathbb{E} \left[ \sum_{t=1}^n \mathbb{E}_{\tau_t \sim \mathbf{A}(z_{1:t-1})} [\mathbf{1}_{\{\tau_t(x_t) \neq y_t\}}] \right] & \leq \inf_{\tau \in \mathcal{T}} \sum_{t=1}^n \mathbf{1}_{\{\tau(x_t) \neq y_t\}} + \sum_l \tilde{C}_n(l) \mathbf{1}_{\{\tilde{C}_n(l) \leq 2\mathfrak{R}_n^{\text{seq}}(\mathcal{T})\}} \\ & \quad + 2\mathfrak{R}_n^{\text{seq}}(\mathcal{T}) \sum_l \mathbf{1}_{\{\tilde{C}_n(l) > 2\mathfrak{R}_n^{\text{seq}}(\mathcal{T})\}} + \sqrt{n} + 2\sqrt{n} \log(N_{\text{leaf}}\pi/\sqrt{6}) \\ & = \inf_{\tau \in \mathcal{T}} \sum_{t=1}^n \mathbf{1}_{\{\tau(x_t) \neq y_t\}} + \sum_l \min(\tilde{C}_n(l), 2\mathfrak{R}_n^{\text{seq}}(\mathcal{T})) + \sqrt{n} \left( 1 + 2 \log(N_{\text{leaf}}\pi/\sqrt{6}) \right) \end{aligned}$$

Now finally we can apply Corollary 28 to bound  $\mathfrak{R}_n^{\text{seq}}(\mathcal{T}) \leq d\mathcal{O}(\log^{3/2} n) \mathfrak{R}_n^{\text{seq}}(\mathcal{H})$  and thus conclude the proof by plugging this into the above.  $\square$

## 4.8.2 Exponentially Weighted Average (EWA) Algorithm on Countable Experts

We consider here a version of the exponentially weighted experts algorithm for countable (possibly infinite) number of experts and provide a bound on the expected regret of the randomized algorithm. The proof of the result closely follows the finite case (e.g. [52, Theorem 2.2]).

Say we are provided with countable experts  $E_1, E_2, \dots$  where each expert can herself be thought of as a randomized/deterministic player strategy which, given history, produces an element of  $\mathcal{F}$  at round  $t$ . Here we also assume that  $\mathcal{F} \subset [0, 1]^{\mathcal{Z}}$  contains only non-negative functions (corresponds to loss class). Denote by  $f_t^i$  the function output by expert  $i$  at round  $t$  given the history. The EWA algorithm we consider needs access to the countable set of experts and also needs an initial weighting on each expert  $p_1, p_2, \dots$  such that  $\sum_i p_i = 1$ .

---

**Algorithm 3** EWA ( $E_1, E_2, \dots, p_1, p_2, \dots$ )

---

Initialize each  $w_i^1 \leftarrow p_i$   
**for**  $t = 1$  to  $n$  **do**  
    Pick randomly an expert  $i$  with probability  $w_i^t$   
    Play  $f_t = f_i^t$   
    Receive  $z_t$   
    Update for each  $i$ ,  $w_i^{t+1} = \frac{w_i^t e^{-\eta f_i^t(z_t)}}{\sum_i w_i^t e^{-\eta f_i^t(z_t)}}$   
**end for**

---

**Proposition 53.** For the exponentially weighted average forecaster (Algorithm 3) with  $\eta = n^{-1/2}$  yields

$$\mathbb{E} \left[ \sum_{t=1}^n f_t(z_t) \right] \leq \sum_{t=1}^n f_i^t(z_t) + \frac{\sqrt{n}}{8} + \sqrt{n} \log(1/p_i)$$

for any  $i \in \mathbb{N}$ .

*Proof.* Define  $W_t = \sum_i p_i e^{-\eta \sum_{j=1}^t f_i^j(z_j)}$ . Then note that

$$\log \left( \frac{W_t}{W_{t-1}} \right) = \log \left( \frac{\sum_i p_i e^{-\eta \sum_{j=1}^t f_i^j(z_j)}}{W_{t-1}} \right) = \log \left( \sum_i w_i^{t-1} e^{-\eta f_i^t(z_t)} \right)$$

Now using Hoeffding's inequality (see [52, Lemma 2.2]) we have that

$$\log \left( \frac{W_t}{W_{t-1}} \right) \leq -\eta \sum_i w_i^{t-1} f_i^t(z_t) + \frac{\eta^2}{8} = -\eta \mathbb{E} [f_t(z_t)] + \frac{\eta^2}{8}$$

Summing over  $t$  we get

$$\log(W_n) - \log(W_0) = \sum_{t=1}^n \log \left( \frac{W_t}{W_{t-1}} \right) \leq -\eta \mathbb{E} \left[ \sum_{t=1}^n f_t(z_t) \right] + \frac{n\eta^2}{8} \quad (4.8.9)$$

Note that  $W_0 = \sum_i p_i = 1$  and so  $\log(W_0) = 0$ . Also note that for any  $i \in \mathbb{N}$ ,

$$\log(W_n) = \log \left( \sum_i p_i e^{-\eta \sum_{t=1}^n f_i^t(z_t)} \right) \geq \log \left( p_i^{-\eta \sum_{t=1}^n f_i^t(z_t)} \right) = \log(p_i) - \eta \sum_{t=1}^n f_i^t(z_t)$$

Hence using this with Equation 4.8.9 we see that

$$\log(p_i) - \eta \sum_{t=1}^n f_i^t(z_t) \leq -\eta \mathbb{E} \left[ \sum_{t=1}^n f_t(z_t) \right] + \frac{n\eta^2}{8}$$

Rearranging we get

$$\mathbb{E} \left[ \sum_{t=1}^n f_t(z_t) \right] \leq \sum_{t=1}^n f_i^t(z_t) + \frac{\eta n}{8} + \frac{1}{\eta} \log \left( \frac{1}{p_i} \right)$$

Using  $\eta = \frac{1}{\sqrt{n}}$  we get the desired bound.  $\square$

*Proof of Proposition 52.* First, by the classical result of Kolmogorov and Tihomirov [58], the class  $\mathcal{G}$  of all bounded Lipschitz functions has small metric entropy:  $\log \widehat{\mathcal{N}}_\infty(\alpha, \mathcal{G}) = \Theta(1/\alpha)$ . For the particular class of non-decreasing 1-Lipschitz functions, it is trivial to verify that the entropy is in fact bounded by  $2/\alpha$ .

Next, consider the class  $\mathcal{F} = \{\langle w, x \rangle \mid \|w\|_2 \leq 1\}$  over the Euclidean ball. By Proposition ??,  $\mathfrak{N}_n^{\text{seq}}(\mathcal{F}) \leq \sqrt{\frac{2}{n}}$ . Using the lower bound of Proposition 42,  $\text{fat}_\alpha^{\text{seq}} \leq 64/\alpha^2$  whenever  $\alpha > 8/\sqrt{n}$ . This implies that  $\mathfrak{N}_\infty^{\text{seq}}(\alpha, \mathcal{F}, n) \leq (2en/\alpha)^{64/\alpha^2}$  whenever  $\alpha > 8/\sqrt{n}$ . Note that this bound does not depend on the ambient dimension of  $\mathcal{Z}$ .

Next, we show that a composition of  $\mathcal{G}$  with any small class  $\mathcal{F} \subset [-1, 1]^{\mathcal{Z}}$  also has a small cover. To this end, suppose  $\mathfrak{N}_\infty^{\text{seq}}(\alpha, \mathcal{F}, n)$  is the covering number for  $\mathcal{F}$ . Fix a particular tree  $z$  and let  $V = \{\mathbf{v}_1, \dots, \mathbf{v}_N\}$  be an  $\ell_\infty$  cover of  $\mathcal{F}$  on  $z$  at scale  $\alpha$ . Analogously, let  $W = \{g_1, \dots, g_M\}$  be an  $\ell_\infty$  cover of  $\mathcal{G}$  with  $M = \widehat{\mathcal{N}}_\infty(\alpha, \mathcal{G})$ . Consider the class  $\mathcal{G} \circ \mathcal{F} = \{g \circ f : g \in \mathcal{G}, f \in \mathcal{F}\}$ . The claim is that  $\{g(\mathbf{v}) : \mathbf{v} \in V, g \in W\}$  provides an  $\ell_\infty$  cover for  $\mathcal{G} \circ \mathcal{F}$  on  $z$ . Fix any  $f \in \mathcal{F}, g \in \mathcal{G}$  and  $\epsilon \in \{\pm 1\}^n$ . Let  $\mathbf{v} \in V$  be such that  $\max_{t \in [n]} |f(z_t(\epsilon)) - \mathbf{v}_t(\epsilon)| \leq \alpha$ , and let  $g' \in W$  be such that  $\|g - g'\|_\infty \leq \alpha$ . Then, using the fact that functions in  $\mathcal{G}$  are 1-Lipschitz, for any  $t \in [n]$ ,

$$|g(f(z_t(\epsilon))) - g'(\mathbf{v}_t(\epsilon))| \leq |g(f(z_t(\epsilon))) - g'(f(z_t(\epsilon)))| + |g'(f(z_t(\epsilon))) - g'(\mathbf{v}_t(\epsilon))| \leq 2\alpha.$$

Hence,  $\mathfrak{N}_\infty^{\text{seq}}(2\alpha, \mathcal{G} \circ \mathcal{F}, n) \leq \widehat{\mathcal{N}}_\infty(\alpha, \mathcal{G}) \times \mathfrak{N}_\infty^{\text{seq}}(\alpha, \mathcal{F}, n)$ .

Finally, we put all the pieces together. By Lemma 27, the Sequential Rademacher complexity of  $\mathcal{H}$  is bounded by 4 times the Sequential Rademacher complexity of the class

$$\mathcal{G} \circ \mathcal{F} = \{u(\langle w, x \rangle) \mid u : [-1, 1] \mapsto [-1, 1] \text{ is 1-Lipschitz}, \|w\|_2 \leq 1\}$$

since the squared loss is 4-Lipschitz on the space of possible values. The latter complexity is then bounded by

$$\mathfrak{D}_n^{\text{seq}}(\mathcal{G} \circ \mathcal{F}) \leq \frac{32}{\sqrt{n}} + 12 \int_{8/\sqrt{n}}^1 \sqrt{\frac{\log \mathfrak{N}_\infty^{\text{seq}}(\delta, \mathcal{G} \circ \mathcal{F}, n)}{n}} d\delta \leq \frac{32}{\sqrt{n}} + 12 \int_{8/\sqrt{n}}^1 \sqrt{\frac{2}{n\delta} + \frac{64}{n\delta^2} \log(2en)} d\delta.$$

We conclude that the value of the game  $\mathcal{V}_n(\mathcal{H}, \mathcal{Z} \times \mathcal{Y}) = O(\sqrt{\frac{\log^3 n}{n}})$ .  $\square$

## 4.9 Discussion

While in this chapter we introduced tools for analyzing rates for online learning problems analogous to the various complexity measures for statistical learning framework, as we saw in the previous chapter, for statistical learning framework these tools fail to characterize learnability in general. Similar situation is true for the online setting too. While these tools characterize learnability of online supervised learning problems and can also be used to obtain rates for online convex learning problems indirectly, in general they could fail to characterize learnability general of online learning problems.

In the previous chapter we then turned to the notion of online stability to characterize learnability and even provided a generic learning algorithm. Is there some notion of stability that can be used to characterize learnability in the online learning framework? Can we provide a generic algorithm for general learning problems in the online framework?

Another interesting avenue to explore is the question of fast rates for online learning problems. In the statistical learning framework the notion of Localized Rademacher complexity introduced in [59] can often be used to obtain fast rates. Just like we provided an analog to Rademacher complexity for online learning, can we provide an analog of localized complexity measures, specifically a local sequential Rademacher complexity that can then be used to obtain fast rates for online learnign problems?

## **Part II**

# **Convex Problems : Oracle Efficient Learning/Optimization**

## Chapter 5

# Convex Learning and Optimization

## Problem Setup

In the first part of this dissertation we mainly focused on the question of learnability (and learning rates) in both statistical and online settings, that is whether the problems were at all learnable using some algorithm. We did not take into consideration tractability of the learning rules we considered and whether the problem is learnable using some efficient learning algorithm. The generic learning rules/algorithms presented in the first part are not at all tractable. In this part of the dissertation, we try to address the issue of tractability of learning algorithms in both statistical and online learning settings. To do this, we restrict ourselves to so called convex learning or optimization problems. In this chapter we introduce the convex learning problems we will encounter in the second part of this dissertation and associated notations.

### 5.1 Convex Problems

Let us now give the basic setup for the convex learning and optimization problems we consider in the second part of this dissertation. Of course when we say convex problem, we mean that the set of target hypothesis  $\mathcal{H}$  is a convex set and for each given instance  $z \in \mathcal{Z}$ , the loss function  $\ell(\mathbf{h}, z)$  is convex in  $\mathbf{h}$ . To describe more formal, we consider an arbitrary real vector space  $\mathcal{B}$  and denote its dual by  $\mathcal{B}^*$ . Now the target hypothesis class  $\mathcal{H} \subset \mathcal{B}$  we shall consider throughout will be a **convex and centrally symmetric** subset of  $\mathcal{B}$ . Also consider the set  $\mathcal{X} \subset \mathcal{B}^*$  to be a **bounded, convex and centrally symmetric** subset of the dual  $\mathcal{B}^*$ . The role of set  $\mathcal{X}$  will become clear in the following paragraph. It will be convenient for us, to relate the notion of a convex centrally symmetric sets to their corresponding (semi)norms. To do this, recall the definition of the Minkowski functional of a set  $\mathcal{K}$  of a real vector space  $\mathcal{B}$ . It is defined as

$$\|\mathbf{v}\|_{\mathcal{K}} := \inf \{ \alpha > 0 : \mathbf{v} \in \alpha \mathcal{K} \}$$

Now it can be seen that if  $\mathcal{K}$  is convex and centrally symmetric (i.e.  $\mathcal{K} = -\mathcal{K}$ ), then  $\|\cdot\|_{\mathcal{K}}$  is a semi-norm. Further, for instance in  $\mathbb{R}^d$ , if the set  $\mathcal{K}$  is bounded then  $\|\cdot\|_{\mathcal{K}}$  is in fact a norm. Our assumption on the sets  $\mathcal{H}$  and  $\mathcal{X}$  ensure that  $\|\cdot\|_{\mathcal{H}}$  and  $\|\cdot\|_{\mathcal{X}}$  (the Minkowski functionals of the sets  $\mathcal{H}$  and  $\mathcal{X}$ ) are semi-norms. For simplicity we shall further assume that  $\|\cdot\|_{\mathcal{H}}$  and  $\|\cdot\|_{\mathcal{X}}$  are in fact norms. Even though we do this for simplicity, we remark that all our results go through for semi-norms too. We use  $\mathcal{X}^*$  and  $\mathcal{H}^*$  to represent the duals of  $\mathcal{X}$  and  $\mathcal{H}$  respectively, i.e. the unit balls of the dual norms  $\|\cdot\|_{\mathcal{X}^*}$  and  $\|\cdot\|_{\mathcal{H}^*}$ .

As mentioned we consider convex learning and optimization problem where target set  $\mathcal{H}$  is the unit ball of norm  $\|\cdot\|_{\mathcal{H}}$  and for each instance  $z \in \mathcal{Z}$  the loss  $\ell(\cdot, z)$  is convex. Now the convex problems we consider are of three flavors. The first case we consider is the one where, for each instance  $z \in \mathcal{Z}$ , the sub-gradients of the convex function  $\ell(\cdot, z)$  belong to the set  $\mathcal{X}$ . Notice that when  $\mathcal{X} = \mathcal{H}^*$  this case exactly corresponds to convex 1-lipschitz problems. Most prior work on online learning, statistical learning and convex optimization problems considers this case when  $\mathcal{H}$  is the unit ball of some Banach space, and  $\mathcal{X}$  is the unit ball of the dual space. However, we analyze the general problem where  $\mathcal{X} \in \mathcal{B}^*$  is not necessarily the dual ball of  $\mathcal{H}$ . The second flavor of problems we consider are the ones where for each  $z \in \mathcal{Z}$ , the function  $\ell(\cdot, z)$  is in fact uniformly convex w.r.t. norm  $\|\cdot\|_{\mathcal{X}^*}$ . The final flavor of problems we consider are problems where for each  $z \in \mathcal{Z}$ , the function  $\ell(\cdot, z)$  are non-negative, convex and smooth w.r.t. norm  $\|\cdot\|_{\mathcal{X}}$ . The next section formally defines the key set of convex function classes we consider for the convex learning and optimization problems throughout the second part.

## 5.2 Various Convex Learning/Optimization Problems

Below we provide examples of various convex learning/optimization problems we consider in this work.

**Example 6** (Lipschitz Convex functions). The set  $\mathcal{Z} = \mathcal{Z}_{\text{Lip}}(\mathcal{X})$  correspond to the set of all convex loss functions such that for any  $z \in \mathcal{Z}$  and  $\mathbf{h} \in \mathcal{H}$ ,  $\nabla_{\mathbf{h}}\ell(\mathbf{h}, z) \in \mathcal{X}$  where  $\mathcal{X}$  is some set in the dual vector space. In short

$$\mathcal{Z}_{\text{Lip}}(\mathcal{X}) = \{z : \ell(\cdot, z) \text{ is convex and } \forall \mathbf{h} \in \mathcal{H}, \nabla\ell(\mathbf{h}, z) \in \mathcal{X}\}$$

**Example 7** (Linear functions). The set  $\mathcal{Z} = \mathcal{Z}_{\text{lin}}$  consists of linear functions on  $\bar{\mathcal{H}}$  from the set  $\mathcal{X}$ .

$$\mathcal{Z}_{\text{lin}}(\mathcal{X}) = \{z : \ell(\cdot, z) = \langle \mathbf{x}, \cdot \rangle \text{ where } \mathbf{x} \in \mathcal{X}\}$$

**Example 8** (Supervised Learning with Linear predictors). The set  $\mathcal{Z} = \mathcal{Z}_{\text{sup}}(\mathcal{X})$  consists of functions of form

$$\mathcal{Z}_{\text{sup}}(\mathcal{X}) = \{z : \ell(\mathbf{h}, z) = |\langle \mathbf{h}, \mathbf{x} \rangle - y| \text{ where } \mathbf{x} \in \mathcal{X}, y \in [-b, b]\}$$

In the first part we introduced generic supervised learning problem with absolute loss and arbitrary function class for prediction that mapped input  $x$  to reals. At first glance the above class might seem specific given that the predictor is always linear. However at second glance, if we consider supervised learning with absolute loss and we require loss function for every instance to be convex in  $\mathbf{h}$ , then necessarily the predictor has to



be linear. To see this note that we need  $|h(\mathbf{x}) - y|$  to be convex in  $\mathbf{h}$ , when both  $y = 1$  and  $y = -1$  which basically means predictor has to be both convex and concave in  $\mathbf{h}$  and so is linear. The same argument can be extended to other common margin losses like squared loss, logistic loss etc. Given a convex loss function  $\phi$  one can also more generally define a class  $\mathcal{Z}_\phi = \{z : \ell(\mathbf{h}, z) = \phi(\langle \mathbf{x}, \mathbf{h} \rangle, y) : \mathbf{x} \in \mathcal{X}, y \in [-b, b], \}$  for any 1-Lipschitz loss function  $\phi : \mathbb{R} \times \mathbb{R} \mapsto \mathbb{R}$ , and this class would also be a subset of  $\mathcal{Z}_{\text{Lip}}$ . In fact, this setting includes supervised learning fairly generally, including problems such as multitask learning and matrix completion, where in all cases  $\mathcal{X}$  specifies the data domain<sup>1</sup>.

**Example 9 (Non-Negative Smooth Convex Loss).** The set  $\mathcal{Z}$  corresponds to non-negative convex functions that are smooth w.r.t. to the norm  $\|\cdot\|_{\mathcal{X}^*}$ , that is

$$\mathcal{Z}_{\text{smt}(H)}(\mathcal{X}) = \{z : \ell(\cdot, z) \geq 0 \text{ is convex and } \forall \mathbf{h}, \mathbf{h}' \in \bar{\mathcal{H}}, \|\nabla \ell(\mathbf{h}, z) - \nabla \ell(\mathbf{h}', z)\|_{\mathcal{X}} \leq H \|\mathbf{h} - \mathbf{h}'\|_{\mathcal{X}^*}\}$$

A subset of the above instance class are cases of supervised learning problems  $\mathcal{Z}_\phi$  where  $\phi$  is non-negative and smooth function on the reals like logistic loss, smoothed hinge loss and squared loss. In the chapters to come we show how for non-negative smooth convex instances one can get faster learning rates in both online and statistical convex learning frameworks when optimal loss itself is small.

**Example 10 (Regularized Convex Loss).** The set  $\mathcal{Z}$  consists of functions of form

$$\mathcal{Z}_{\text{reg}}(\mathcal{X}) = \{z : \ell(\cdot, z) = \phi(\cdot, z) + R(\mathbf{h}) \text{ where } \phi(\cdot, z) \text{ is convex, } \forall \mathbf{h} \in \mathcal{H}, \nabla \phi(\mathbf{h}, z) \in \mathcal{X} \text{ and } R \text{ is convex}\}$$

The class  $\mathcal{Z}_{\text{reg}}(\mathcal{X})$  captures regularized convex objective classes where  $R : \bar{\mathcal{H}} \mapsto \mathbb{R}$  is a convex regularizer used to enforce structure or prior into the learning problem. Commonly regularizer chosen are strongly convex or more generally uniformly convex. The following instance class captures more specifically these classes.

**Example 11 (Uniformly Convex Loss).** The set  $\mathcal{Z}$  corresponds to uniformly convex functions,

$$\mathcal{Z}_{\text{ucvx}(\sigma, q)}(\mathcal{X}) = \{z : \ell(\cdot, z) = \phi(\cdot, z) + R(\cdot) \text{ is } (\sigma, q)\text{-uniformly convex w.r.t. } \|\cdot\|_{\mathcal{X}^*}, \forall z \in \mathcal{Z} \nabla \phi(\cdot, z) \in \mathcal{X}\}$$

In the above example, we assume that  $\ell$  is uniformly convex but however we assume that  $\ell$  can be decomposed as  $\ell(\cdot, z) = \phi(\cdot, z) + R(\cdot)$  and only assume that  $\nabla \phi(\cdot, z) \in \mathcal{X}$ . The reason we did not directly assume in the above that  $\nabla \ell$  is itself in  $\mathcal{X}$  is so that we can capture many regularized learning problems where the regularizer  $R$  is not Lipschitz and so  $\ell$  is not Lipschitz but however the loss function of interest  $\phi$  in these examples are Lipschitz.

Another class of important loss functions are non-negative smooth convex loss functions. Linear predictors with logistic loss, squared loss are common examples of such problems. The following example captures such classes.

---

<sup>1</sup>Note that any convex supervised learning problem can necessarily be viewed as linear classification with some convex constraint  $\mathcal{H}$  on the predictors.

**Example 12** (Bounded Convex functions). The set  $\mathcal{Z} = \mathcal{Z}_{\text{bnd}}$  corresponds to the set of all convex loss functions that are bounded by  $b$  on  $\mathcal{H}$ , that is

$$\mathcal{Z}_{\text{bnd}} = \{z : \ell(\cdot, z) \text{ is convex and bounded by } b \text{ on } \mathcal{H}\}$$

**Remark 54.** *It must be noted that for convex optimization problem, both in the online and statistical setting, we can use Jensen's inequality to show that for every randomized algorithm there exists a deterministic algorithm (that play the expected action of the randomized algorithm) that achieves learning rates that are at most as bad as that of the randomized algorithm. Hence for the convex optimization problem it suffices to only consider deterministic learning algorithms.*

Owing to the above remark, we see that since we only need to consider deterministic learning algorithms, both in statistical and online convex learning settings, any learning algorithm  $\mathbf{A}$  is specified by a mapping  $\mathbf{A} : \bigcup_{n \in \mathbb{N}} \mathcal{Z}^{n-1} \mapsto \bar{\mathcal{H}}$ . Also note that while we assumed that the set  $\mathcal{H}$  was convex and centrally symmetric and for ease even assumed that  $\mathcal{H}$  is the unit ball of norm  $\|\cdot\|_{\mathcal{H}}$ , no such assumptions are made on hypothesis set  $\bar{\mathcal{H}}$  other than that it contains  $\mathcal{H}$  and so it could even be all of  $\mathcal{B}$ .

### 5.3 Discussion

Notice that the sets  $\mathcal{H}$  and  $\mathcal{X}$  that we consider are arbitrary convex centrally symmetric sets and need not be related to each other a priori. While the special case of when  $\mathcal{H} = \mathcal{X}^*$  is what is usually encountered in majority of the theoretical analysis existing literature, in many applications  $\mathcal{H}$  and  $\mathcal{X}$  are not dual to each other. Here we provide a generic theoretical analysis of the non-dual case. Note that when  $\mathcal{H} = \mathcal{X}^*$ ,  $\mathcal{Z}_{\text{Lip}}(\mathcal{X})$  corresponds to usual convex Lipschitz problem,  $\mathcal{Z}_{\text{smt}}(\mathcal{X})$  corresponds to usual smooth convex class and finally  $\mathcal{Z}_{\text{ucvx}(\sigma, 2)}(\mathcal{X})$  corresponds to  $\sigma$ -strongly convex functions. Overall, the convex learning and optimization problems that we introduced in this chapter and will study in the remaining chapters cover majority of the problems considered in previous works. Perhaps the one case not covered in this thesis is the case of exp-concave loss functions for which in finite (low) dimensional cases, one can get faster learning rates in both online and statistical learning cases.

# Chapter 6

## Mirror Descent Methods

Perhaps one of the most popular convex optimization algorithm most readers would be familiar with is the gradient descent algorithm. The gradient descent algorithm proceeds by starting with an initial point and iteratively updating it by taking steps in the direction of the negative gradient of the function to optimize at the current point. The gradient descent method is a natural algorithm for problems in Euclidean space. The mirror descent algorithm [1] is a natural generalization of gradient descent method for general convex learning problems. Section 6.1 describes the basic update step of the mirror descent algorithm. Section ?? provides bounds on regret of mirror descent method for generic online convex learning problems, online smooth convex learning problems and online uniformly convex learning problems. The section 6.3 which follows shows how mirror descent algorithm can be used for statistical convex learning problem and associated learning guarantee. Section 6.4 shows how mirror descent can also be used for offline convex optimization problems. Following that Section 6.5 provides proofs of all the results of this chapter and finally we conclude this chapter with some discussion in Section 6.6.

### 6.1 The Mirror Descent Update

Given a strictly convex function  $\Psi : \mathcal{B} \mapsto \mathbb{R}$ , the Mirror Descent algorithm,  $\mathbf{A}_{\text{MD}}$  is given by the update

$$\mathbf{h}_{t+1} = \underset{\mathbf{h} \in \mathcal{H}}{\operatorname{argmin}} \Delta_{\Psi}(\mathbf{h} | \mathbf{h}_t) + \eta \langle \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h} - \mathbf{h}_t \rangle \quad (6.1.1)$$

$$\text{or equivalently } \mathbf{h}'_{t+1} = \nabla \Psi^*(\nabla \Psi(\mathbf{h}_t) - \eta \nabla \ell(\mathbf{h}_t, z_t)), \quad \mathbf{h}_{t+1} = \underset{\mathbf{h} \in \mathcal{H}}{\operatorname{argmin}} \Delta_{\Psi}(\mathbf{h} | \mathbf{h}'_{t+1}) \quad (6.1.2)$$

where  $\Delta_{\Psi}(\mathbf{h} | \mathbf{h}') := \Psi(\mathbf{h}) - \Psi(\mathbf{h}') - \langle \nabla \Psi(\mathbf{h}'), \mathbf{h} - \mathbf{h}' \rangle$  is the Bregman divergence and  $\Psi^*$  is the convex conjugate of  $\Psi$ . As an example notice that when  $\Psi(\mathbf{h}) = \frac{1}{2} \|\mathbf{h}\|_2^2$  then we get back the online gradient descent algorithm. It is worth noting that the perceptron algorithm can be viewed as a conservative variant of online gradient descent with hinge loss function. Also when  $\mathcal{H}$  is the  $d$  dimensional simplex and  $\Psi(\mathbf{h}) =$

$\sum_{i=1}^d \mathbf{h}_i \log(1/\mathbf{h}_i)$ , then we get the multiplicative weights update algorithm. In general the function  $\Psi$  used in mirror descent is often referred to as the proxy-function.

The mirror descent algorithm is an  $O(1)$ , memory single pass, first order method that only needs some sub-gradient for each update. Often times in practice, each mirror descent update step has time complexity same as that of calculating a single gradient and hence overall runtime is linear in number of rounds (for online learning case) or number of samples (statistical learning case). This makes the mirror descent algorithm attractive from a computational viewpoint.

Before we proceed we would like to point out that in the setting we consider, the hypothesis set from which learner is allowed to pick,  $\bar{\mathcal{H}} \subset \mathcal{B}$ , need not be the same as the target hypothesis set  $\mathcal{H}$  and only needs to be a superset. Of course when  $\bar{\mathcal{H}} = \mathcal{H}$  then the update above corresponds to the usual mirror descent update. However when  $\bar{\mathcal{H}}$  is all of  $\mathcal{B}$  then notice that the projection step is mute and essentially the update becomes,

$$\mathbf{h}_{t+1} = \nabla \Psi^* (\nabla \Psi(\mathbf{h}_t) - \eta \nabla \ell(\mathbf{h}_t, z_t)) .$$

This is especially attractive because in many machine learning applications while we would like to do as well as the best hypothesis from some target class, we don't really care if the hypothesis learner picks itself is selected from this target hypothesis set as long as it gives good results. Hence if we set up the problem such that  $\bar{\mathcal{H}} = \mathcal{B}$  then we avoid extra computational time on projection step which could at times be expensive.

## 6.2 Online Mirror Descent

In this section we describe the online mirror descent algorithm  $\mathcal{A}_{\text{MD}} : \bigcup_{n \in \mathbb{N}} \mathcal{Z}^{n-1} \mapsto \bar{\mathcal{H}}$  which simply uses the mirror descent update given in the previous section and returns the  $h_t$ 's in each round. That is,  $\mathcal{A}_{\text{MD}}(\{\}) = \mathbf{h}_1$  and further for any  $t \in \mathbb{N}$  and any  $z_1, \dots, z_{t-1} \in \mathcal{Z}$ ,

$$\mathcal{A}_{\text{MD}}(z_1, \dots, z_{t-1}) = h_t .$$

A key tool in the analysis mirror descent is the notion of strong convexity or more generally uniform convexity of the function  $\Psi$ . Recall the definition of uniform convexity :

**Definition 31.** A function  $\Psi : \mathcal{B} \rightarrow \mathbb{R}$  is said to be  $q$ -uniformly convex w.r.t.  $\|\cdot\|$  if for any  $\mathbf{h}, \mathbf{h}' \in \mathcal{B}$ :

$$\forall_{\alpha \in [0,1]} \Psi(\alpha \mathbf{h} + (1-\alpha)\mathbf{h}') \leq \alpha \Psi(\mathbf{h}) + (1-\alpha)\Psi(\mathbf{h}') - \frac{\alpha(1-\alpha)}{q} \|\mathbf{h} - \mathbf{h}'\|^q$$

We are interested in bounding the regret of the of the mirror descent algorithm given by,

$$\begin{aligned} \mathbf{R}_n(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n) &:= \frac{1}{n} \sum_{t=1}^n \ell(\mathcal{A}_{\text{MD}}(z_{1:t-1}), z_t) - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \\ &= \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t). \end{aligned}$$

Now assuming we can find an appropriate  $q$ -uniformly convex function on  $\mathcal{B}$ , below we provide bounds on regret of mirror descent method for generic online convex learning problems, non-negative smooth convex learning problems and uniformly convex learning problems.

**Convex Losses with Sub-gradients in  $\mathcal{X}$**  : We first start with the case when the convex costs at each round are such that their sub-gradients lie in the set  $\mathcal{X}$ . Note that instance sets  $\mathcal{Z}_{\text{Lip}}(\mathcal{X})$ ,  $\mathcal{Z}_{\text{supp}}(\mathcal{X})$  and  $\mathcal{Z}_{\text{lin}}(\mathcal{X})$  are examples of instance classes that fall in this set.

**Lemma 55.** *Let  $\Psi : \mathcal{B} \mapsto \mathbb{R}$  be non-negative and  $q$ -uniformly convex w.r.t. norm  $\|\cdot\|_{\mathcal{X}^*}$ . For the Mirror Descent algorithm with this  $\Psi$ , using  $\mathbf{h}_1 = \underset{\mathbf{h} \in \mathcal{H}}{\text{argmin}} \Psi(\mathbf{h})$  and  $\eta = \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{nB} \right)^{1/p}$  we can guarantee that for any  $z_1, \dots, z_n$  s.t.  $\frac{1}{n} \sum_{t=1}^n \|\nabla \ell(\cdot, z_t)\|_{\mathcal{X}}^p \leq 1$  (where  $p = \frac{q}{q-1}$ ),*

$$\mathbf{R}(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n) \leq 2 \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{n} \right)^{\frac{1}{q}}.$$

Note that in our case we have that for each  $z \in \mathcal{Z}$ ,  $\nabla \ell(\cdot, z) \in \mathcal{X}$ , i.e.  $\|\nabla \ell(\cdot, z)\|_{\mathcal{X}} \leq 1$ , and so certainly  $\frac{1}{n} \sum_{t=1}^n \|\nabla \ell(\cdot, z_t)\|_{\mathcal{X}}^p \leq 1$ .

**Non-Negative Smooth Convex Losses** : Next we deal with the case when the convex costs in each round need not be such that their sub gradients are from set  $\mathcal{X}$  but rather are such that the costs are non-negative and  $H$ -smooth w.r.t. to the norm  $\|\cdot\|_{\mathcal{X}}$ . That is :

$$\forall z \in \mathcal{Z}, \forall \mathbf{h}, \mathbf{h}' \in \bar{\mathcal{H}}, \|\nabla \ell(\mathbf{h}, z) - \nabla \ell(\mathbf{h}', z)\|_{\mathcal{X}} \leq H \|\mathbf{h} - \mathbf{h}'\|_{\mathcal{X}^*}.$$

For non-negative smooth losses, one has a property called self-bounding property that for any  $\mathbf{h}$  and any  $z$ ,

$$\|\nabla \ell(\mathbf{h}, z)\|_{\mathcal{X}} \leq \sqrt{4H\ell(\mathbf{h}, z)}$$

In [60] we had made the observation that any non-negative smooth convex loss satisfies this above self bounding property. Shalev-Shwartz [61] showed that the self bounding property can be used this to provide optimistic rates on regret of mirror gradient descent for the dual case and using strongly convex function. By optimistic rates we refer to rates that improve when average loss of best hypothesis is small. The following lemma provides regret bounds for mirror descent algorithm with optimistic rates using similar lines of

reasoning as in [60, 61].

**Lemma 56.** *Let  $\Psi : \mathcal{B} \mapsto \mathbb{R}$  be non-negative and  $q$ -uniformly convex w.r.t. norm  $\|\cdot\|_{\mathcal{X}^*}$ . For any  $\bar{L}^* \geq 0$ , using the Mirror Descent algorithm with function  $\Psi$  for making updates, using  $\mathbf{h}_1 = \operatorname{argmin}_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})$  and*

$$\eta = \begin{cases} \left( \frac{p \sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{n} \right)^{1/p} \frac{1}{\sqrt{4HL^*}} & \text{if } \bar{L}^* \geq \frac{16H}{p^{2/p}} \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{n} \right)^{2/q} \\ \left( \frac{p}{2} \right)^{\frac{p}{2}} \frac{1}{4H} \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{n} \right)^{\frac{2-p}{p}} & \text{otherwise} \end{cases}$$

we can guarantee that for any sequence  $z_1, \dots, z_n \in \mathcal{Z}$  and any  $\mathbf{h}^* \in \mathcal{H}$  such that  $\frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \leq \bar{L}^*$ ,

$$\mathbf{R}(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n) \leq \sqrt{64HL^*} \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{n} \right)^{1/q} + 40H \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{n} \right)^{2/q}.$$

**Uniformly Convex Losses :** We now consider the case when loss functions are uniformly convex, specifically  $(\sigma, q')$ -uniformly convex. Up to now in this chapter we considered  $\bar{\mathcal{H}}$  to be any superset of  $\mathcal{H}$  including  $\bar{\mathcal{H}} = \mathcal{H}$ . For the case of uniformly convex loss instance class  $\mathcal{Z}_{\text{ucvx}(\sigma, q)}(\mathcal{X})$  alone we will assume that  $\bar{\mathcal{H}} = \mathcal{B}$  so that we don't need a projection step and  $\mathbf{h}'_t = \mathbf{h}_t$  in the mirror descent update. Under this setting we have the following upper bound for regret of mirror descent algorithm.

**Lemma 57.** *Let  $\Psi : \mathcal{B} \mapsto \mathbb{R}$  be non-negative and  $q$ -uniformly convex w.r.t. norm  $\|\cdot\|_{\mathcal{X}^*}$ . Let  $\psi : \mathcal{B} \mapsto \mathbb{R}$  be non-negative and  $q'$  uniformly convex w.r.t. norm  $\|\cdot\|_{\mathcal{X}^*}$ . For each  $t \in [n]$ , define*

$$\tilde{\Psi}_t(\cdot) = \frac{1}{\eta} \Psi(\cdot) + R(\cdot) + \sigma t \psi(\cdot)$$

Then for the choice of

$$\eta = \begin{cases} \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{n} \right)^{1/p} & \text{if } n \geq \left( (2-p')\sigma^{p'-1} \sup_{\mathbf{h} \in \mathcal{H}} \Psi^{1/q}(\mathbf{h}) \right)^{\frac{1}{2-p'-1/p}} \\ \infty & \text{otherwise} \end{cases}$$

we can guarantee that for any sequence  $z_1, \dots, z_n \in \mathcal{Z}_{\text{ucvx}(\sigma, q')}$ , if  $q' > 2$ :

$$\mathbf{R}(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n) \leq \min \left\{ \frac{2 (\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h}^*))^{1/q}}{n^{1/q}}, \frac{2}{(2-p')\sigma^{p'-1}n^{p'-1}} \right\} + \frac{\sup_{\mathbf{h} \in \mathcal{H}} R(\mathbf{h})}{n}.$$

where  $p' = \frac{q'}{q'-1}$  and for  $q' = 2$ :  $\mathbf{R}(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n) \leq \frac{2 \log n}{\sigma n} + \frac{\sup_{\mathbf{h} \in \mathcal{H}} R(\mathbf{h})}{n}$ .

**Choice of  $\Psi$  : Construction I** The Mirror Descent bound suggests that as long as we can find an appropriate function  $\Psi$  that is uniformly convex w.r.t.  $\|\cdot\|_{\mathcal{X}^*}$  we can get a diminishing regret guarantee using Mirror

Descent. This suggests constructing the following function:

$$\tilde{\Psi}_q := \underset{\substack{\psi: \psi \text{ is } q\text{-uniformly convex} \\ \text{w.r.t. } \|\cdot\|_{\mathcal{X}^*} \text{ on } \mathcal{H} \text{ and } \psi \geq 0}}{\operatorname{argmin}} \sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h}). \quad (6.2.1)$$

If no  $q$ -uniformly convex function exists then  $\tilde{\Psi}_q = \infty$  is assumed by default. The above function is in a sense the best choice for the Mirror Descent bound in Eq. (55). The question then is: when can we find such appropriate functions and what is the best rate we can guarantee using Mirror Descent?

### 6.3 Stochastic Mirror Descent

In the previous section we saw that mirror descent can be used successfully for online convex learning. In general, especially for convex problems, any online method can be converted into an algorithm for statistical learning with same learning rate guarantee and this is often referred to as online to batch conversion. Refer to [62] for more details about online to batch conversion. Hence one can also use the mirror descent method for statistical convex learning problems. This algorithm is often referred to as stochastic mirror descent algorithm. The stochastic mirror descent algorithm is given as follows :

$$\overline{\mathcal{A}_{\text{MD}}}(\{\}) = \mathbf{h}_1, \quad \forall t \in [n], \quad \overline{\mathcal{A}_{\text{MD}}}(z_1, \dots, z_t) = \frac{1}{t} \sum_{i=1}^t \mathbf{h}_i.$$

**Proposition 58.** *For any hypothesis set  $\bar{\mathcal{H}} \subset \mathcal{B}$ , target hypothesis class  $\mathcal{H} \subset \mathcal{B}$  and convex learning problem specified by instance space  $\mathcal{Z}$  and any fixed distribution  $\mathcal{D}$  on instance space  $\mathcal{Z}$  :*

$$\mathbb{E}_{S \sim \mathcal{D}^n} \left[ L_{\mathcal{D}}(\overline{\mathcal{A}_{\text{MD}}}(z_{1:n})) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] \leq \mathbb{E}_{S \sim \mathcal{D}^n} [\mathbf{R}_n(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n)]$$

*Proof.* Note that

$$\begin{aligned} \mathbb{E}_{S \sim \mathcal{D}^n} [\mathbf{R}_n(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n)] &= \mathbb{E}_{S \sim \mathcal{D}^n} \left[ \frac{1}{n} \sum_{t=1}^n \ell(\mathcal{A}_{\text{MD}}(z_{1:t-1}), z_t) - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \right] \\ &= \frac{1}{n} \sum_{t=1}^n \mathbb{E}_{S \sim \mathcal{D}^n} [\ell(\mathcal{A}_{\text{MD}}(z_{1:t-1}), z_t)] - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \\ &= \frac{1}{n} \sum_{t=1}^n \mathbb{E}_{S \sim \mathcal{D}^n} [\mathbb{E}_{z_t \sim \mathcal{D}} [\ell(\mathcal{A}_{\text{MD}}(z_{1:t-1}), z_t)]] - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \\ &= \frac{1}{n} \sum_{t=1}^n \mathbb{E}_{S \sim \mathcal{D}^n} [L_{\mathcal{D}}(\mathcal{A}_{\text{MD}}(z_{1:t-1}))] - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \\ &= \mathbb{E}_{S \sim \mathcal{D}^n} \left[ \frac{1}{n} \sum_{t=1}^n L_{\mathcal{D}}(\mathcal{A}_{\text{MD}}(z_{1:t-1})) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] \end{aligned}$$

$$\begin{aligned}
&\geq \mathbb{E}_{S \sim \mathcal{D}^n} \left[ L_{\mathcal{D}} \left( \frac{1}{n} \sum_{t=1}^n \mathcal{A}_{\text{MD}}(z_{1:t-1}) \right) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] \\
&= \mathbb{E}_{S \sim \mathcal{D}^n} \left[ L_{\mathcal{D}}(\overline{\mathcal{A}_{\text{MD}}}(z_{1:n})) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right]
\end{aligned}$$

where the inequality step above is due to Jensen's inequality. Thus we get the statement of the proposition.  $\square$

Owing to the above proposition we see that one can get the same learning guarantees for statistical convex learning problems as for the online counterpart. Specifically we get the following lemmas for statistical convex learning problems.

**Lemma 59.** *Let  $\Psi : \mathcal{B} \mapsto \mathbb{R}$  be non-negative and  $q$ -uniformly convex w.r.t. norm  $\|\cdot\|_{\mathcal{X}^*}$ . For the Mirror Descent algorithm with this  $\Psi$ , using  $\mathbf{h}_1 = \underset{\mathbf{h} \in \mathcal{H}}{\operatorname{argmin}} \Psi(\mathbf{h})$  and  $\eta = \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{nB} \right)^{1/p}$  we can guarantee that for any distribution  $\mathcal{D}$  over  $\mathcal{Z}$  s.t.  $\mathbb{E}_{z \sim \mathcal{D}} [\|\nabla \ell(\cdot, z)\|_{\mathcal{X}}^p] \leq 1$  we have that :*

$$\mathbf{R}(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n) \leq 2 \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{n} \right)^{\frac{1}{q}}. \quad (\text{where } p = \frac{q}{q-1})$$

*Proof.* The statement follows by using Proposition 58 along with line of proof in Lemma 55.  $\square$

Similar bound can be given for stochastic convex learning of non-negative smooth convex losses as follows :

**Lemma 60.** *Let  $\Psi : \mathcal{B} \mapsto \mathbb{R}$  be non-negative and  $q$ -uniformly convex w.r.t. norm  $\|\cdot\|_{\mathcal{X}^*}$ . For any  $L^* \geq 0$ , using the stochastic Mirror Descent algorithm with function  $\Psi$  for making updates, using  $\mathbf{h}_1 = \underset{\mathbf{h} \in \mathcal{H}}{\operatorname{argmin}} \Psi(\mathbf{h})$  and*

$$\eta = \begin{cases} \left( \frac{p \sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{n} \right)^{1/p} \frac{1}{\sqrt{4HL^*}} & \text{if } L^* \geq \frac{16H}{p^{2/p}} \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{n} \right)^{2/q} \\ \left( \frac{p}{2} \right)^{\frac{p}{2}} \frac{1}{4H} \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{n} \right)^{\frac{2-p}{p}} & \text{otherwise} \end{cases}$$

*we can guarantee that for any distribution  $\mathcal{D}$  over  $\mathcal{Z}$  such that  $\inf_{\mathbf{h} \in \mathcal{H}} L(\mathbf{h}) \leq L^*$ ,*

$$L(\overline{\mathcal{A}_{\text{MD}}}) - \inf_{\mathbf{h} \in \mathcal{H}} L(\mathbf{h}) \leq \sqrt{64HL^*} \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{n} \right)^{1/q} + 40H \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{n} \right)^{2/q}.$$

*Proof.* The statement follows by using Proposition 58 along with line of proof in Lemma 56.  $\square$

Finally bounds for stochastic learning of uniformly convex losses can also be given.



**Lemma 61.** Let  $\Psi : \mathcal{B} \mapsto \mathbb{R}$  be non-negative and  $q$ -uniformly convex w.r.t. norm  $\|\cdot\|_{\mathcal{X}^*}$ . Let  $\psi : \mathcal{B} \mapsto \mathbb{R}$  be non-negative and  $q'$  uniformly convex w.r.t. norm  $\|\cdot\|_{\mathcal{X}^*}$ . For each  $t \in [n]$ , define

$$\tilde{\Psi}_t(\cdot) = \frac{1}{\eta} \Psi(\cdot) + R(\cdot) + \sigma t \psi(\cdot)$$

Then for the choice of

$$\eta = \begin{cases} \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{n} \right)^{1/p} & \text{if } n \geq \left( (2-p')\sigma^{p'-1} \sup_{\mathbf{h} \in \mathcal{H}} \Psi^{1/q}(\mathbf{h}) \right)^{\frac{1}{2-p'-1/p}} \\ \infty & \text{otherwise} \end{cases}$$

we can guarantee that for any distribution  $\mathcal{D}$  over instance space  $\mathcal{Z}_{\text{ucvx}(\sigma, q')}(\mathcal{X})$ , if  $q' > 2$ :

$$L(\overline{\mathcal{A}_{\text{MD}}}) - \inf_{\mathbf{h} \in \mathcal{H}} L(\mathbf{h}) \leq \min \left\{ \frac{2 (\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h}^*))^{1/q}}{n^{1/q}}, \frac{2}{(2-p')\sigma^{p'-1}n^{p'-1}} \right\} + \frac{\sup_{\mathbf{h} \in \mathcal{H}} R(\mathbf{h})}{n}.$$

where  $p' = \frac{q'}{q'-1}$  and for  $q' = 2$ :  $L(\overline{\mathcal{A}_{\text{MD}}}) - \inf_{\mathbf{h} \in \mathcal{H}} L(\mathbf{h}) \leq \frac{2 \log n}{\sigma n} + \frac{\sup_{\mathbf{h} \in \mathcal{H}} R(\mathbf{h})}{n}$ .

*Proof.* The statement follows by using Proposition 58 along with line of proof in Lemma 57.  $\square$

## 6.4 Mirror Descent for Offline Optimization

Notice that one can think of offline convex optimization as a special case of stochastic convex optimization where the set of distributions we consider are point distributions on a single instance of the instance space  $\mathcal{Z}$ . Using this observation we see that  $\overline{\mathcal{A}_{\text{MD}}}$  can directly be used for offline convex optimization problems with same guarantee on sub-optimality. In this thesis, for the problem of offline convex optimization we only consider the instance class  $\mathcal{Z}_{\text{Lip}}(\mathcal{X})$ . The following corollary is a direct consequence of Lemma 59.

**Corollary 62.** Let  $\Psi : \mathcal{B} \mapsto \mathbb{R}$  be non-negative and  $q$ -uniformly convex w.r.t. norm  $\|\cdot\|_{\mathcal{X}^*}$ . For the Stochastic Mirror Descent algorithm with this  $\Psi$ , using  $\mathbf{h}_1 = \underset{\mathbf{h} \in \mathcal{H}}{\text{argmin}} \Psi(\mathbf{h})$  and  $\eta = \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{nB} \right)^{1/p}$  we can guarantee that for any instance  $z \in \mathcal{Z}_{\text{Lip}}$ , we have that :

$$\ell(\overline{\mathcal{A}_{\text{MD}}}(\nabla \ell(h_1, z), \dots, \nabla \ell(h_n, z)), z) - \inf_{\mathbf{h} \in \mathcal{H}} \ell(\mathbf{h}, z) \leq 2 \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{n} \right)^{\frac{1}{q}}. \quad (\text{where } p = \frac{q}{q-1})$$

## 6.5 Detailed Proofs

*Proof of Lemma 55 (generalized MD guarantee).* Note that for any  $\mathbf{h}^* \in \mathcal{H}$ ,

$$\begin{aligned}
\eta \left( \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \right) &\leq \sum_{t=1}^n \langle \eta \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}^* \rangle \\
&= \sum_{t=1}^n \left( \langle \eta \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}'_{t+1} \rangle + \langle \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}'_{t+1} - \mathbf{h}^* \rangle \right) \\
&= \sum_{t=1}^n \left( \langle \eta \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}'_{t+1} \rangle + \langle \nabla \Psi(\mathbf{h}_t) - \nabla \Psi(\mathbf{h}'_{t+1}), \mathbf{h}'_{t+1} - \mathbf{h}^* \rangle \right) \\
&\leq \sum_{t=1}^n \left( \|\eta \nabla \ell(\mathbf{h}_t, z_t)\|_{\mathcal{X}} \|\mathbf{h}_t - \mathbf{h}'_{t+1}\|_{\mathcal{X}^*} + \langle \nabla \Psi(\mathbf{h}_t) - \nabla \Psi(\mathbf{h}'_{t+1}), \mathbf{h}'_{t+1} - \mathbf{h}^* \rangle \right) \\
&\leq \sum_{t=1}^n \left( \frac{\eta^p}{p} \|\nabla \ell(\mathbf{h}_t, z_t)\|_{\mathcal{X}}^p + \frac{1}{q} \|\mathbf{h}_t - \mathbf{h}'_{t+1}\|_{\mathcal{X}^*}^q + \langle \nabla \Psi(\mathbf{h}_t) - \nabla \Psi(\mathbf{h}'_{t+1}), \mathbf{h}'_{t+1} - \mathbf{h}^* \rangle \right)
\end{aligned}$$

Using simple manipulation we can show that

$$\langle \nabla \Psi(\mathbf{h}_t) - \nabla \Psi(\mathbf{h}'_{t+1}), \mathbf{h}'_{t+1} - \mathbf{h}^* \rangle = \Delta_{\Psi}(\mathbf{h}^* | \mathbf{h}_t) - \Delta_{\Psi}(\mathbf{h}^* | \mathbf{h}'_{t+1}) - \Delta_{\Psi}(\mathbf{h}'_{t+1} | \mathbf{h}_t)$$

where given any  $\mathbf{h}, \mathbf{h}' \in \mathcal{B}$ ,

$$\Delta_{\Psi}(\mathbf{h} | \mathbf{h}') := \Psi(\mathbf{h}) - \Psi(\mathbf{h}') - \langle \nabla \Psi(\mathbf{h}'), \mathbf{h} - \mathbf{h}' \rangle$$

is the Bregman divergence between  $\mathbf{h}$  and  $\mathbf{h}'$  w.r.t. function  $\Psi$ . Hence,

$$\begin{aligned}
&\eta \left( \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \right) \\
&\leq \sum_{t=1}^n \left( \frac{\eta^p}{p} \|\nabla \ell(\mathbf{h}_t, z_t)\|_{\mathcal{X}}^p + \frac{1}{q} \|\mathbf{h}_t - \mathbf{h}'_{t+1}\|_{\mathcal{X}^*}^q + \langle \nabla \Psi(\mathbf{h}_t) - \nabla \Psi(\mathbf{h}'_{t+1}), \mathbf{h}'_{t+1} - \mathbf{h}^* \rangle \right) \\
&= \sum_{t=1}^n \left( \frac{\eta^p}{p} \|\nabla \ell(\mathbf{h}_t, z_t)\|_{\mathcal{X}}^p + \frac{1}{q} \|\mathbf{h}_t - \mathbf{h}'_{t+1}\|_{\mathcal{X}^*}^q + \Delta_{\Psi}(\mathbf{h}^* | \mathbf{h}_t) - \Delta_{\Psi}(\mathbf{h}^* | \mathbf{h}'_{t+1}) - \Delta_{\Psi}(\mathbf{h}'_{t+1} | \mathbf{h}_t) \right) \\
&\leq \sum_{t=1}^n \left( \frac{\eta^p}{p} \|\nabla \ell(\mathbf{h}_t, z_t)\|_{\mathcal{X}}^p + \frac{1}{q} \|\mathbf{h}_t - \mathbf{h}'_{t+1}\|_{\mathcal{X}^*}^q + \Delta_{\Psi}(\mathbf{h}^* | \mathbf{h}_t) - \Delta_{\Psi}(\mathbf{h}^* | \mathbf{h}'_{t+1}) - \Delta_{\Psi}(\mathbf{h}'_{t+1} | \mathbf{h}_t) \right) \\
&= \sum_{t=1}^n \left( \frac{\eta^p}{p} \|\nabla \ell(\mathbf{h}_t, z_t)\|_{\mathcal{X}}^p + \frac{1}{q} \|\mathbf{h}_t - \mathbf{h}'_{t+1}\|_{\mathcal{X}^*}^q - \Delta_{\Psi}(\mathbf{h}'_{t+1} | \mathbf{h}_t) \right) + \Delta_{\Psi}(\mathbf{h}^* | \mathbf{h}_1) - \Delta_{\Psi}(\mathbf{h}^* | \mathbf{h}_{n+1}) \\
&\leq \sum_{t=1}^n \left( \frac{\eta^p}{p} \|\nabla \ell(\mathbf{h}_t, z_t)\|_{\mathcal{X}}^p + \frac{1}{q} \|\mathbf{h}_t - \mathbf{h}'_{t+1}\|_{\mathcal{X}^*}^q - \Delta_{\Psi}(\mathbf{h}'_{t+1} | \mathbf{h}_t) \right) + \Psi(\mathbf{h}^*)
\end{aligned}$$

Now since  $\Psi$  is  $q$ -uniformly convex w.r.t.  $\|\cdot\|_{\mathcal{X}^*}$ , for any  $\mathbf{h}, \mathbf{h}' \in \mathcal{B}^*$ ,  $\Delta_\Psi(\mathbf{h}'|\mathbf{h}) \geq \frac{1}{q} \|\mathbf{h} - \mathbf{h}'\|_{\mathcal{X}^*}^q$ . Hence we conclude that

$$\begin{aligned} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) &\leq \frac{\eta^{p-1}}{p} \sum_{t=1}^n \|\nabla \ell(\mathbf{h}_t, z_t)\|_{\mathcal{X}}^p + \frac{\Psi(\mathbf{h}^*)}{\eta} \\ &\leq \frac{\eta^{p-1} B n}{p} + \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{\eta} \\ &\leq \frac{\eta^{p-1} B n}{p} + \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{\eta} \end{aligned}$$

Plugging in the value of  $\eta = \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{nB} \right)^{1/p}$  we get :

$$\sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \leq 2 \left( \sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h}) \right)^{1/q} (Bn)^{1/p}$$

dividing throughout by  $n$  conclude the proof.  $\square$

*Proof of Lemma 56.* Note that for any  $\mathbf{h}^* \in \mathcal{H}$ ,

$$\begin{aligned} \eta \left( \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \right) &\leq \sum_{t=1}^n \langle \eta \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}^* \rangle \\ &= \sum_{t=1}^n \left( \langle \eta \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}'_{t+1} \rangle + \langle \eta \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}'_{t+1} - \mathbf{h}^* \rangle \right) \\ &= \sum_{t=1}^n \left( \langle \eta \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}'_{t+1} \rangle + \langle \nabla \Psi(\mathbf{h}_t) - \nabla \Psi(\mathbf{h}'_{t+1}), \mathbf{h}'_{t+1} - \mathbf{h}^* \rangle \right) \\ &\leq \sum_{t=1}^n \left( \|\eta \nabla \ell(\mathbf{h}_t, z_t)\|_{\mathcal{X}} \|\mathbf{h}_t - \mathbf{h}'_{t+1}\|_{\mathcal{X}^*} + \langle \nabla \Psi(\mathbf{h}_t) - \nabla \Psi(\mathbf{h}'_{t+1}), \mathbf{h}'_{t+1} - \mathbf{h}^* \rangle \right) \\ &\leq \sum_{t=1}^n \left( \frac{\eta^p}{p} \|\nabla \ell(\mathbf{h}_t, z_t)\|_{\mathcal{X}}^p + \frac{1}{q} \|\mathbf{h}_t - \mathbf{h}'_{t+1}\|_{\mathcal{X}^*}^q + \langle \nabla \Psi(\mathbf{h}_t) - \nabla \Psi(\mathbf{h}'_{t+1}), \mathbf{h}'_{t+1} - \mathbf{h}^* \rangle \right) \end{aligned}$$

Using simple manipulation we can show that

$$\langle \nabla \Psi(\mathbf{h}_t) - \nabla \Psi(\mathbf{h}'_{t+1}), \mathbf{h}'_{t+1} - \mathbf{h}^* \rangle = \Delta_\Psi(\mathbf{h}^*|\mathbf{h}_t) - \Delta_\Psi(\mathbf{h}^*|\mathbf{h}'_{t+1}) - \Delta_\Psi(\mathbf{h}'_{t+1}|\mathbf{h}_t)$$

where given any  $\mathbf{h}, \mathbf{h}' \in \mathcal{B}$ ,

$$\Delta_\Psi(\mathbf{h}|\mathbf{h}') := \Psi(\mathbf{h}) - \Psi(\mathbf{h}') - \langle \nabla \Psi(\mathbf{h}'), \mathbf{h} - \mathbf{h}' \rangle$$

is the Bregman divergence between  $\mathbf{h}$  and  $\mathbf{h}'$  w.r.t. function  $\Psi$ . Hence,

$$\begin{aligned}
& \eta \left( \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \right) \\
& \leq \sum_{t=1}^n \left( \frac{\eta^p}{p} \|\nabla \ell(\mathbf{h}_t, z_t)\|_{\mathcal{X}}^p + \frac{1}{q} \|\mathbf{h}_t - \mathbf{h}'_{t+1}\|_{\mathcal{X}^*}^q + \langle \nabla \Psi(\mathbf{h}_t) - \nabla \Psi(\mathbf{h}'_{t+1}), \mathbf{h}'_{t+1} - \mathbf{h}^* \rangle \right) \\
& = \sum_{t=1}^n \left( \frac{\eta^p}{p} \|\nabla \ell(\mathbf{h}_t, z_t)\|_{\mathcal{X}}^p + \frac{1}{q} \|\mathbf{h}_t - \mathbf{h}'_{t+1}\|_{\mathcal{X}^*}^q + \Delta_{\Psi}(\mathbf{h}^* | \mathbf{h}_t) - \Delta_{\Psi}(\mathbf{h}^* | \mathbf{h}'_{t+1}) - \Delta_{\Psi}(\mathbf{h}'_{t+1} | \mathbf{h}_t) \right) \\
& \leq \sum_{t=1}^n \left( \frac{\eta^p}{p} \|\nabla \ell(\mathbf{h}_t, z_t)\|_{\mathcal{X}}^p + \frac{1}{q} \|\mathbf{h}_t - \mathbf{h}'_{t+1}\|_{\mathcal{X}^*}^q + \Delta_{\Psi}(\mathbf{h}^* | \mathbf{h}_t) - \Delta_{\Psi}(\mathbf{h}^* | \mathbf{h}'_{t+1}) - \Delta_{\Psi}(\mathbf{h}'_{t+1} | \mathbf{h}_t) \right) \\
& = \sum_{t=1}^n \left( \frac{\eta^p}{p} \|\nabla \ell(\mathbf{h}_t, z_t)\|_{\mathcal{X}}^p + \frac{1}{q} \|\mathbf{h}_t - \mathbf{h}'_{t+1}\|_{\mathcal{X}^*}^q - \Delta_{\Psi}(\mathbf{h}'_{t+1} | \mathbf{h}_t) \right) + \Delta_{\Psi}(\mathbf{h}^* | \mathbf{h}_1) - \Delta_{\Psi}(\mathbf{h}^* | \mathbf{h}_{n+1}) \\
& \leq \sum_{t=1}^n \left( \frac{\eta^p}{p} \|\nabla \ell(\mathbf{h}_t, z_t)\|_{\mathcal{X}}^p + \frac{1}{q} \|\mathbf{h}_t - \mathbf{h}'_{t+1}\|_{\mathcal{X}^*}^q - \Delta_{\Psi}(\mathbf{h}'_{t+1} | \mathbf{h}_t) \right) + \Psi(\mathbf{h}^*)
\end{aligned}$$

Now since  $\Psi$  is  $q$ -uniformly convex w.r.t.  $\|\cdot\|_{\mathcal{X}^*}$ , for any  $\mathbf{h}, \mathbf{h}' \in \mathcal{B}^*$ ,  $\Delta_{\Psi}(\mathbf{h}' | \mathbf{h}) \geq \frac{1}{q} \|\mathbf{h} - \mathbf{h}'\|_{\mathcal{X}^*}^q$ . Hence we conclude that

$$\sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \leq \frac{\eta^{p-1}}{p} \sum_{t=1}^n \|\nabla \ell(\mathbf{h}_t, z_t)\|_{\mathcal{X}}^p + \frac{\Psi(\mathbf{h}^*)}{\eta}$$

Now by smoothness of objective, using the Lemma ?? we have that each  $\|\nabla \ell(\mathbf{h}_t, z_t)\|_{\mathcal{X}} \leq \sqrt{4H\ell(\mathbf{h}_t, z_t)}$ .

Using this in the above we get that,

$$\begin{aligned}
\frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) & \leq \frac{\eta^{p-1}}{p} \frac{1}{n} \sum_{t=1}^n (4H\ell(\mathbf{h}_t, z_t))^{\frac{p}{2}} + \frac{1}{n} \frac{\Psi(\mathbf{h}^*)}{\eta} \\
& \leq \frac{(4H)^{\frac{p}{2}} \eta^{p-1}}{p} \left( \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) \right)^{\frac{p}{2}} + \frac{1}{n} \frac{\Psi(\mathbf{h}^*)}{\eta} \\
& \leq \frac{(4H)^{\frac{p}{2}} \eta^{p-1}}{p} \left( \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \right)^{\frac{p}{2}} + \frac{(4H)^{\frac{p}{2}} \eta^{p-1}}{p} \left( \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \right)^{\frac{p}{2}}
\end{aligned}$$

We now use the fact that if for any  $x, B \geq 0$  if  $x \leq Bx^\alpha + A$  for some  $\alpha \in (\frac{1}{2}, 1]$  then as long as  $B < A^{1-\alpha}$ ,  $x \leq \frac{A}{1-BA^{\alpha-1}}$ . Using this with  $x = \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t)$ ,  $B = \frac{(4H)^{\frac{p}{2}} \eta^{p-1}}{p}$  and  $A = \frac{1}{n} \frac{\Psi(\mathbf{h}^*)}{\eta} + \frac{(4H)^{\frac{p}{2}} \eta^{p-1}}{p} \left( \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \right)^{\frac{p}{2}}$  we conclude that for any  $\eta$  such that,

$$\frac{(4H)^{\frac{p}{2}} \eta^{p-1}}{p} \leq \frac{1}{2} \left( \frac{1}{n} \frac{\Psi(\mathbf{h}^*)}{\eta} + \frac{(4H)^{\frac{p}{2}} \eta^{p-1}}{p} \left( \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \right)^{\frac{p}{2}} \right)^{1-\frac{p}{2}}, \quad (6.5.1)$$

we will have that :

$$\frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \leq \frac{\frac{1}{n} \frac{\Psi(\mathbf{h}^*)}{\eta} + \frac{(4H)^{\frac{p}{2}} \eta^{p-1}}{p} \left( \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \right)^{p/2}}{1 - \frac{(4H)^{\frac{p}{2}} \eta^{p-1}}{p} \left( \frac{1}{n} \frac{\Psi(\mathbf{h}^*)}{\eta} + \frac{(4H)^{\frac{p}{2}} \eta^{p-1}}{p} \left( \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \right)^{p/2} \right)^{\frac{p}{2}-1}} \quad (6.5.2)$$

$$\leq 2 \left( \frac{1}{n} \frac{\Psi(\mathbf{h}^*)}{\eta} + \frac{(4H)^{\frac{p}{2}} \eta^{p-1}}{p} \left( \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \right)^{p/2} \right) \quad (6.5.3)$$

To this end we now choose the step size as

$$\eta = \begin{cases} \left( \frac{p\Psi(\mathbf{h}^*)}{n} \right)^{1/p} \frac{1}{\sqrt{\frac{4H}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t)}} & \text{if } \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \geq \frac{16H}{p^{2/p}} \left( \frac{\Psi(\mathbf{h}^*)}{n} \right)^{2/q} \\ \left( \frac{p}{2} \right)^{\frac{p}{2}} \frac{1}{4H} \left( \frac{\Psi(\mathbf{h}^*)}{n} \right)^{\frac{2-p}{p}} & \text{otherwise} \end{cases}$$

It is easy to verify that the choice of  $\eta$  above satisfies condition in Equation 6.5.1. To see this note that if we plug in  $\eta = \left( \frac{p\Psi(\mathbf{h}^*)}{n} \right)^{1/p} \frac{1}{\sqrt{\frac{4H}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t)}}$  into condition in Equation 6.5.1 and rearrange we get that  $\frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \geq \frac{16H}{p^{2/p}} \left( \frac{\Psi(\mathbf{h}^*)}{n} \right)^{2/q}$ . Thus whenever  $\frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \geq \frac{16H}{p^{2/p}} \left( \frac{\Psi(\mathbf{h}^*)}{n} \right)^{2/q}$ , the choice of  $\eta$  satisfies the condition. On the other hand, the choice  $\eta = \left( \frac{p}{2} \right)^{\frac{p}{2}} \frac{1}{4H} \left( \frac{\Psi(\mathbf{h}^*)}{n} \right)^{\frac{2-p}{p}}$  was in the first place derived by making

$$\frac{(4H)^{\frac{p}{2}} \eta^{p-1}}{p} \leq \frac{1}{2} \left( \frac{1}{n} \frac{\Psi(\mathbf{h}^*)}{\eta} \right)$$

which is got by zeroing out the second term in the condition. Hence we conclude that this choice of  $\eta$  always satisfies the condition in Equation 6.5.1. Now we plug in this choice of  $\eta$  into the bound in Equation 6.5.2. Note that whenever  $\eta = \left( \frac{p\Psi(\mathbf{h}^*)}{n} \right)^{1/p} \frac{1}{\sqrt{\frac{4H}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t)}}$ , then simply plugging in this choice of  $\eta$  into the bound in Equation 6.5.2, we get

$$\begin{aligned} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) &\leq \frac{\sqrt{64H} \Psi^{1/q}(\mathbf{h}^*) \sqrt{\frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t)}}{p^{1/p} n^{1/q}} \\ &\leq \sqrt{\frac{64H}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t)} \left( \frac{\Psi(\mathbf{h}^*)}{n} \right)^{1/q} \end{aligned} \quad (6.5.4)$$

On the other hand when

$$\eta = \left( \frac{p}{2} \right)^{\frac{p}{2}} \frac{1}{4H} \left( \frac{\Psi(\mathbf{h}^*)}{n} \right)^{\frac{2-p}{p}}$$

plugging into bound in Equation 6.5.2 we get,

$$\frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \leq 2 \left( \frac{4H}{\left(\frac{p}{2}\right)^{\frac{p}{2}}} \left(\frac{\Psi(\mathbf{h}^*)}{n}\right)^{2/q} + \frac{(4H)^{\frac{2-p}{2}} \left(\frac{p}{2}\right)^{\frac{p(p-1)}{2}} \left(\frac{\Psi(\mathbf{h}^*)}{n}\right)^{\frac{(2-p)(p-1)}{p}}}{p} \left(\frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t)\right)^{p/q} \right)$$

However note that we pick  $\eta = \left(\frac{p}{2}\right)^{\frac{p}{2}} \frac{1}{4H} \left(\frac{\Psi(\mathbf{h}^*)}{n}\right)^{\frac{2-p}{p}}$  only when  $\frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) < \frac{16H}{p^{2/p}} \left(\frac{\Psi(\mathbf{h}^*)}{n}\right)^{2/q}$  and so plugging this inequality we get,

$$\begin{aligned} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) &\leq 2 \left( \frac{4H}{\left(\frac{p}{2}\right)^{\frac{p}{2}}} \left(\frac{\Psi(\mathbf{h}^*)}{n}\right)^{2/q} + \frac{2^{p-2} 4H}{\left(\frac{p}{2}\right)^{\frac{4-p(p-1)}{2}}} \left(\frac{\Psi(\mathbf{h}^*)}{n}\right)^{2/q} \right) \\ &\leq 40H \left(\frac{\Psi(\mathbf{h}^*)}{n}\right)^{2/q} \end{aligned} \quad (6.5.5)$$

Combining the above and bound in Equation 6.5.4 we conclude that for this choice of  $\eta$ ,

$$\begin{aligned} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) &\leq \max \left\{ \sqrt{\frac{64H}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t)} \left(\frac{\Psi(\mathbf{h}^*)}{n}\right)^{1/q}, 40H \left(\frac{\Psi(\mathbf{h}^*)}{n}\right)^{2/q} \right\} \\ &\leq \sqrt{\frac{64H}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t)} \left(\frac{\Psi(\mathbf{h}^*)}{n}\right)^{1/q} + 40H \left(\frac{\Psi(\mathbf{h}^*)}{n}\right)^{2/q} \end{aligned}$$

□

*Proof of Lemma 57.* Note that for any  $\mathbf{h}^* \in \mathcal{H}$ ,

$$\begin{aligned} \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) &\leq \sum_{t=1}^n \langle \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}^* \rangle - \frac{\sigma}{q'} \|\mathbf{h}_t - \mathbf{h}^*\|_{\mathcal{X}^*}^{q'} \\ &= \sum_{t=1}^n \langle \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}'_{t+1} \rangle + \langle \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}'_{t+1} - \mathbf{h}^* \rangle - \frac{\sigma}{q'} \|\mathbf{h}_t - \mathbf{h}^*\|_{\mathcal{X}^*}^{q'} \\ &= \sum_{t=1}^n \langle \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}'_{t+1} \rangle + \langle \nabla \tilde{\Psi}_t(\mathbf{h}_t) - \nabla \tilde{\Psi}_t(\mathbf{h}'_{t+1}), \mathbf{h}'_{t+1} - \mathbf{h}^* \rangle - \frac{\sigma}{q'} \|\mathbf{h}_t - \mathbf{h}^*\|_{\mathcal{X}^*}^{q'} \end{aligned}$$

Using simple manipulation we can show that

$$\langle \nabla \tilde{\Psi}_t(\mathbf{h}_t) - \nabla \tilde{\Psi}_t(\mathbf{h}'_{t+1}), \mathbf{h}'_{t+1} - \mathbf{h}^* \rangle = \Delta_{\tilde{\Psi}_t}(\mathbf{h}^* | \mathbf{h}_t) - \Delta_{\tilde{\Psi}_t}(\mathbf{h}^* | \mathbf{h}'_{t+1}) - \Delta_{\tilde{\Psi}_t}(\mathbf{h}'_{t+1} | \mathbf{h}_t)$$

Hence,

$$\begin{aligned}
& \sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \leq \sum_{t=1}^n \langle \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}'_{t+1} \rangle + \left\langle \nabla \tilde{\Psi}_t(\mathbf{h}_t) - \nabla \tilde{\Psi}_t(\mathbf{h}'_{t+1}), \mathbf{h}'_{t+1} - \mathbf{h}^* \right\rangle - \frac{\sigma}{q'} \|\mathbf{h}_t - \mathbf{h}^*\|_{\mathcal{X}^*}^{q'} \\
& \leq \sum_{t=1}^n \langle \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}'_{t+1} \rangle + \Delta_{\tilde{\Psi}_t}(\mathbf{h}^* | \mathbf{h}_t) - \Delta_{\tilde{\Psi}_t}(\mathbf{h}^* | \mathbf{h}'_{t+1}) - \Delta_{\tilde{\Psi}_t}(\mathbf{h}'_{t+1} | \mathbf{h}_t) - \frac{\sigma}{q'} \|\mathbf{h}_t - \mathbf{h}^*\|_{\mathcal{X}^*}^{q'} \\
& \leq \sum_{t=1}^n \langle \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}'_{t+1} \rangle + \Delta_{\tilde{\Psi}_t}(\mathbf{h}^* | \mathbf{h}_t) - \Delta_{\tilde{\Psi}_t}(\mathbf{h}^* | \mathbf{h}_{t+1}) - \Delta_{\tilde{\Psi}_t}(\mathbf{h}'_{t+1} | \mathbf{h}_t) - \frac{\sigma}{q'} \|\mathbf{h}_t - \mathbf{h}^*\|_{\mathcal{X}^*}^{q'} \\
& \leq \sum_{t=1}^n (\langle \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}'_{t+1} \rangle - \Delta_{\tilde{\Psi}_t}(\mathbf{h}'_{t+1} | \mathbf{h}_t)) + \Delta_{\tilde{\Psi}_1}(\mathbf{h}^* | \mathbf{h}_1) - \frac{\sigma}{q'} \|\mathbf{h}_1 - \mathbf{h}^*\|_{\mathcal{X}^*}^{q'} \\
& \quad + \sum_{t=2}^n \left( \Delta_{\tilde{\Psi}_t}(\mathbf{h}^* | \mathbf{h}_t) - \Delta_{\tilde{\Psi}_{t-1}}(\mathbf{h}^* | \mathbf{h}_t) - \frac{\sigma}{q'} \|\mathbf{h}_t - \mathbf{h}^*\|_{\mathcal{X}^*}^{q'} \right) \\
& \leq \sum_{t=1}^n (\langle \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}'_{t+1} \rangle - \Delta_{\tilde{\Psi}_t}(\mathbf{h}'_{t+1} | \mathbf{h}_t)) + \Delta_{\tilde{\Psi}_1}(\mathbf{h}^* | \mathbf{h}_1) - \sigma \Delta_{\psi}(\mathbf{h}^* | \mathbf{h}_1) \\
& \quad + \sum_{t=2}^n \left( \Delta_{\tilde{\Psi}_t}(\mathbf{h}^* | \mathbf{h}_t) - \Delta_{\tilde{\Psi}_{t-1}}(\mathbf{h}^* | \mathbf{h}_t) - \sigma \Delta_{\psi}(\mathbf{h}^* | \mathbf{h}_t) \right) \\
& = \sum_{t=1}^n (\langle \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}'_{t+1} \rangle - \Delta_{\tilde{\Psi}_t}(\mathbf{h}'_{t+1} | \mathbf{h}_t)) + \frac{1}{\eta} \Delta_{\Psi}(\mathbf{h}^* | \mathbf{h}_1) + \Delta_R(\mathbf{h}^* | \mathbf{h}_1) \\
& = \sum_{t=1}^n (\langle \nabla \ell(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}_{t+1} \rangle - \Delta_{\tilde{\Psi}_t}(\mathbf{h}_{t+1} | \mathbf{h}_t)) + \frac{1}{\eta} \Delta_{\Psi}(\mathbf{h}^* | \mathbf{h}_1) + \Delta_R(\mathbf{h}^* | \mathbf{h}_1) \\
& = \sum_{t=1}^n \left( \langle \nabla \phi(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}_{t+1} \rangle - \Delta_{\frac{\Psi}{\eta} + \sigma t \psi}(\mathbf{h}_{t+1} | \mathbf{h}_t) + R(\mathbf{h}_t) - R(\mathbf{h}_{t+1}) \right) + \frac{1}{\eta} \Delta_{\Psi}(\mathbf{h}^* | \mathbf{h}_1) \\
& \quad + \Delta_R(\mathbf{h}^* | \mathbf{h}_1) \\
& \leq \sum_{t=1}^n \left( \langle \nabla \phi(\mathbf{h}_t, z_t), \mathbf{h}_t - \mathbf{h}'_{t+1} \rangle - \frac{1}{\eta q} \|\mathbf{h}^* - \mathbf{h}'_{t+1}\|_{\mathcal{X}^*}^q - \frac{\sigma t}{q'} \|\mathbf{h}^* - \mathbf{h}'_{t+1}\|_{\mathcal{X}^*}^{q'} \right) + \frac{\Psi(\mathbf{h}^*)}{\eta} \\
& \quad + R(\mathbf{h}^*) - R(\mathbf{h}_{n+1}) \\
& \leq \sum_{t=1}^n \inf_{\mathbf{u}_t + \mathbf{v}_t = \nabla \phi(\mathbf{h}_t, z_t)} \left\{ \frac{\eta^{p-1}}{p} \|\mathbf{u}_t\|_{\mathcal{X}}^p + \frac{1}{p' \sigma^{p'-1} t^{p'-1}} \|\mathbf{v}_t\|_{\mathcal{X}}^{p'} \right\} + \frac{\Psi(\mathbf{h}^*)}{\eta} + R(\mathbf{h}^*)
\end{aligned}$$

Where in the steps above we used the fact that for any functions  $F$  and  $G$ ,  $\Delta_{G+F}(\cdot | \cdot) = \Delta_G(\cdot | \cdot) + \Delta_F(\cdot | \cdot)$  and that fact that for any function  $F$  that is  $q$ -uniformly convex w.r.t.  $\|\cdot\|_{\mathcal{X}^*}$ , for any  $\mathbf{h}, \mathbf{h}' \in \mathcal{B}^*$ ,  $\Delta_F(\mathbf{h}' | \mathbf{h}) \geq \frac{1}{q} \|\mathbf{h} - \mathbf{h}'\|_{\mathcal{X}^*}^q$ . The final step is due to Fenchel Young inequality. Now we upper bound the summation term by replacing each infimum over decompositions of  $\nabla \ell(\mathbf{h}_t, z_t)$  into any arbitrary vectors  $\mathbf{u}_t$  and  $\mathbf{v}_t$  to vectors of specific form,  $\mathbf{u}_t = (\mathbf{1} - \alpha) \nabla \phi(\mathbf{h}_t, z_t)$  and  $\mathbf{v}_t = \alpha \nabla \phi(\mathbf{h}_t, z_t)$  for some  $\alpha \in [0, 1]$ . Hence we get for any

$\alpha \in [0, 1]$ ,

$$\begin{aligned}
\sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) &\leq \sum_{t=1}^n \left( \frac{\eta^{p-1}(1-\alpha)^p}{p} \|\nabla\phi(\mathbf{h}_t, z_t)\|_{\mathcal{X}}^p + \frac{\alpha^{p'}}{p' \sigma^{p'-1} t^{p'-1}} \|\nabla\phi(\mathbf{h}_t, z_t)\|_{\mathcal{X}}^{p'} \right) \\
&\quad + \frac{\Psi(\mathbf{h}^*)}{\eta} + R(\mathbf{h}^*) \\
&\leq \frac{\eta^{p-1}(1-\alpha)^p n}{p} + \frac{\alpha^{p'}}{p' \sigma^{p'-1}} \sum_{t=1}^n \frac{1}{t^{p'-1}} + \frac{\Psi(\mathbf{h}^*)}{\eta} + R(\mathbf{h}^*) \\
&\leq \frac{\eta^{p-1}(1-\alpha)^p n}{p} + \frac{\alpha^{p'}}{p' \sigma^{p'-1}} \frac{n^{2-p'}}{2-p'} + \frac{\Psi(\mathbf{h}^*)}{\eta} + R(\mathbf{h}^*) \\
&\leq \eta^{p-1}(1-\alpha)^p n + \frac{\alpha^{p'}}{\sigma^{p'-1}} \frac{n^{2-p'}}{2-p'} + \frac{\Psi(\mathbf{h}^*)}{\eta} + R(\mathbf{h}^*) \\
&\leq \eta^{p-1}(1-\alpha)^p n + \frac{\alpha^{p'}}{\sigma^{p'-1}} \frac{n^{2-p'}}{2-p'} + \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{\eta} + \sup_{\mathbf{h} \in \mathcal{H}} R(\mathbf{h})
\end{aligned}$$

Using  $\alpha = 1$  whenever  $n \geq \left( (2-p')\sigma^{p'-1}\Psi^{1/q}(\mathbf{h}^*) \right)^{\frac{1}{2-p'-1/p}}$  and  $\alpha = 0$  otherwise and picking

$$\eta = \begin{cases} \left( \frac{\sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h})}{n} \right)^{1/p} & \text{if } n \geq \left( (2-p')\sigma^{p'-1} \sup_{\mathbf{h} \in \mathcal{H}} \Psi^{1/q}(\mathbf{h}) \right)^{\frac{1}{2-p'-1/p}} \\ \infty & \text{otherwise} \end{cases}$$

we get that

$$\sum_{t=1}^n \ell(\mathbf{h}_t, z_t) - \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \leq \min \left\{ 2 \sup_{\mathbf{h} \in \mathcal{H}} \Psi^{1/q}(\mathbf{h}) n^{1/p} + \sup_{\mathbf{h} \in \mathcal{H}} R(\mathbf{h}), \frac{2}{(2-p')\sigma^{p'-1}} n^{2-p'} + \sup_{\mathbf{h} \in \mathcal{H}} R(\mathbf{h}) \right\}$$

Dividing throughout by  $n$  concludes the proof.  $\square$

## 6.6 Discussion

The mirror descent algorithm with uniformly convex  $\Psi$  functions were introduced by Nemirovski and Yudin in [1] for offline convex optimization. Specific upper bounds for offline convex optimization of  $\mathcal{Z}_{\text{Lip}}$  dual case when  $\mathcal{H}$  is the unit  $\ell_p$  ball and  $\mathcal{X}$  is the dual of  $\mathcal{H}$  are provided in [1]. For online convex optimization problem, online gradient descent (Euclidean case) was proposed by Zinkevich in [30]. Faster rates when the losses are strongly convex in the Euclidean case for online gradient descent was proposed in [63]. Mirror descent for general strongly convex objectives with  $\log n/n$  rates was proposed and analyzed in [64]. While in all the above the set  $\mathcal{X}$  in the corresponding problems are same as dual of set  $\mathcal{H}$  and  $\mathcal{H} = \tilde{\mathcal{H}}$ , in this chapter we consider the generic case and provide bounds for the non-dual case for online and statistical convex learning and for offline convex optimization. One fact to pay attention to is that the upper bounds are provided assuming one can find appropriate function  $\Psi$  that is  $q$ -uniformly convex w.r.t. norm  $\|\cdot\|_{\mathcal{X}^*}$ .



Ofcourse the immediate question that arises is “When can one find such functions  $\Psi$  and are the bounds got using such  $\Psi$  optimal?”. The next three chapters deals with this question for online and statistical convex learning problems and for offline convex optimization problems.

## Chapter 7

# Optimality of Mirror Descent for Online Convex Learning Problem

In this chapter we will show that the mirror descent method is universal and near optimal for online convex learning problems. Very roughly, the main result we show in this chapter can be stated as :

**For any online convex learning problem, if some online learning algorithm can guarantee a regret bound of  $\text{Rate}_n$ , then mirror descent algorithm can guarantee regret bounded as  $\tilde{O}(\text{Rate}_n)$ .** Of course in the remainder of the chapter we will exactly qualify this result and show optimality of mirror descent for online learning when losses are s.t. gradients are in set  $\mathcal{X}$ , for non-negative smooth convex losses and for uniformly convex losses.

Before we proceed we start by noticing that owing to Remark 54, it suffices to only consider Deterministic learning algorithms. As a result, the online convex learning problem can be viewed as a multi-round game where on round  $t$ , the learner first picks a vector  $\mathbf{h}_t \in \mathcal{H}$ . Next, the adversary picks instance  $z_t \in \mathcal{Z}$  where  $\mathcal{Z}$  is a class of instances specifying some set of convex functions. At the end of the round, the learner pays instantaneous cost  $\ell(\mathbf{h}_t, z_t)$ . Recall that a deterministic online learning algorithm  $\mathbf{A}$  for the problem is specified by the mapping  $\mathbf{A} : \bigcup_{n \in \mathbb{N}} \mathcal{Z}^{n-1} \mapsto \bar{\mathcal{H}}$ . We shall represent the regret of an algorithm  $\mathbf{A}$  for a given sequence of instances  $z_1, \dots, z_n$  by the shorthand :

$$\mathbf{R}_n(\mathbf{A}, z_1, \dots, z_n) := \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{A}(z_{1:t-1}), z_t) - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) .$$

The goal of the learner as before, is to minimize the regret at the end of  $n$  rounds.

In Chapter 4 since we considered randomized online learning algorithms we had to be careful in defining the value of the game in Equation 4.2.1. Since for convex learning problems it suffices to only consider deterministic learning algorithms, it is easier to write down the value of the game for these problems. The value of the online convex learning problem can be written as the best possible guarantee on regret against

any sequence of instances that any algorithm can enjoy. Formally the value can be written as :

$$\mathcal{V}_n(\mathcal{H}, \mathcal{Z}) = \inf_{\mathbf{A}} \sup_{z_{1:n} \in \mathcal{Z}} \mathbf{R}_n(\mathbf{A}, z_1, \dots, z_n) \quad (7.0.1)$$

In the Section 7.1 we will see that for online convex learning problems introduced in previous chapter, a problem is online learnable if and only if it is learnable using a gradient-based online learning algorithm. We will also see that value of the linear game acts as key in characterizing optimal learning rates of various other convex learning problems and hence will focus on that. In Chapter 6 we describe the online mirror descent algorithm and provide guarantees for various problems. However these guarantees relied on our ability to be able to pick appropriate uniformly convex function to use with the mirror descent algorithm. In section 7.2 we show how the concept of martingale type (a generalization of it as per our need) captures closely the value of linear game and hence can be used to closely characterize rates for the various convex learning problems. In Section 7.3 we extend Pisier’s result [65] to show that martingale type of the problem can be used to ensure existence of an appropriate uniformly convex function. Subsequently in Section 7.4 we put it all together and establish that for the convex problems we consider, there exists appropriate uniformly convex function so that mirror descent with this function and right step size is always near optimal (upto to log factors). Thus we establish universality and near optimality of mirror descent. This is also shown for non-negative smooth convex losses and certain uniformly convex losses. In Section 7.5 we provide several examples of commonly encountered convex learning problems and establish rates for these problems using mirror descent. Section 7.6 provides detailed proofs for results in this chapter and finally we conclude with some discussion in Section 7.7.

## 7.1 Value of the Linear Game

The value of online learning learning problem plays an important role in characterizing optimal rates of various online convex learning problems. In this section we show how the value of the linear online learning game is related to value of other online convex learning games and also introduce some necessary definitions to build towards showing the main result of this chapter. The following lemma shows how one can use online algorithms for linear problems for other online convex learning problems (problems where sub-gradients are in  $\mathcal{X}$ ).

**Lemma 63.** *Let  $\mathbf{A}$  be any online learning algorithm for linear learning problems specified by instance set  $\mathcal{Z}_{\text{lin}}$ . Using this, for any convex learning problem specified by instance set  $\mathcal{Z}$  such that for any  $\mathbf{h} \in \mathcal{H}$ ,  $\nabla_{\mathbf{h}} \ell(\mathbf{h}, z) \in \mathcal{X}$ , one can construct a new gradient-based learning algorithm  $\mathbf{A}^{\text{O}1\text{st}}$  such that for any  $z_1, \dots, z_n \in \mathcal{Z}$ ,*

$$\mathbf{R}_n(\mathbf{A}^{\text{O}1\text{st}}, z_1, \dots, z_n) \leq \sup_{z_1^*, \dots, z_n^* \in \mathcal{Z}_{\text{lin}}} \mathbf{R}_n(\mathbf{A}, z_1^*, \dots, z_n^*)$$

A direct consequence of the above lemma is the following corollary that shows that value of the linear game upper bounds value of other online convex learning problems and is in fact equal to the value of the supervised

learning game and online convex Lipschitz learning game.

**Corollary 64.** *For any convex learning problem specified by instance set  $\mathcal{Z}$  which is s.t.  $\forall \mathbf{h} \in \mathcal{H}, z \in \mathcal{Z} : \nabla_{\mathbf{h}} \ell(\mathbf{h}, z) \in \mathcal{X}$ , we have that,*

$$\mathcal{V}_n(\mathcal{H}, \mathcal{Z}) \leq \mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X}))$$

Furthermore  $\mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{Lip}}(\mathcal{X})) = \mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{sup}}(\mathcal{X})) = \mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X}))$

The equality in the above corollary can also be extended to most other commonly occurring convex loss function classes like say the hinge loss class and logistic learning loss class with some extra constant factors. As we see from the above result, the value of the online learning problem for linear instance class  $\mathcal{Z}_{\text{lin}}(\mathcal{X})$  is critical in upper and lower bounds on rates of various other convex learning problems. In fact as we will see later the value of the linear game also plays an important role in characterizing rates of smooth and uniformly convex online learning problems. Owing to this, for any  $p \in [1, 2]$  we define constant :

$$V_p := \inf \left\{ V \mid \forall n \in \mathbb{N}, \mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X})) \leq V n^{-(1-\frac{1}{p})} \right\} \quad (7.1.1)$$

Notice that  $V_p$  characterizes optimal rate for online linear learning problems (and hence supervised and Lipschitz classes too) up to polynomial.

The main aim of this chapter is to show near optimality of mirror descent algorithm. To this end, similar to  $V_p$  for each  $p \in [1, 2]$  we can define:

$$\text{MD}_p := \inf \left\{ D : \exists \Psi, \eta \text{ s.t. } \forall n \in \mathbb{N}, \sup_{z_{1:n} \in \mathcal{Z}} \mathbf{R}_n(\mathcal{A}_{\text{MD}}, z_{1:n}) \leq D n^{-(1-\frac{1}{p})} \right\} \quad (7.1.2)$$

where the Mirror Descent algorithm in the above definition is run with the corresponding  $\Psi$  and  $\eta$ . The constant  $\text{MD}_p$  is a characterization of the best guarantee the Mirror Descent algorithm can provide by choosing the best  $\Psi$  and  $\eta$ .

A simple consequence of the definitions of  $V_p$  and  $\text{MD}_p$  is the following proposition.

**Proposition 65.** *For any  $p \in [1, 2]$ :*

$$V_p \leq \text{MD}_p$$

## 7.2 Value and Martingale Type

In [7], it was shown that the concept of the *Martingale type* (also sometimes called the *Haar type*) of a Banach space and optimal rates for online convex optimization problem, where  $\mathcal{X}$  and  $\mathcal{H}$  are duals of each other, are closely related. In this section we extend the classic notion of Martingale type of a Banach space (see for instance [65]) to one that accounts for the pair  $(\mathcal{H}^*, \mathcal{X})$ . Before we proceed with the definitions we would like to introduce a few necessary notations. First, throughout we shall use  $\epsilon \in \{\pm 1\}^{\mathbb{N}}$  to represent infinite sequence of signs drawn uniformly at random (i.e. each  $\epsilon_i$  has equal probability of being +1 or -1). Also

throughout  $(\mathbf{x}_n)_{n \in \mathbb{N}}$  represents an  $\mathcal{B}^*$  valued tree of infinite depth, that is a sequence of mappings where each  $\mathbf{x}_n : \{\pm 1\}^{n-1} \mapsto \mathcal{B}^*$ . We are now ready to give the extended definition of Martingale type (or M-type) of a pair  $(\mathcal{H}^*, \mathcal{X})$ .

**Definition 32.** A pair  $(\mathcal{H}^*, \mathcal{X})$  of subsets of a vector space  $\mathcal{B}^*$  is said to be of M-type  $p$  if there exists a constant  $C \geq 1$  such that for all  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth and any  $\mathbf{x}_0 \in \mathcal{B}^*$  :

$$\sup_n \mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] \leq C^p \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{n \geq 1} \mathbb{E} [\|\mathbf{x}_n(\epsilon)\|_{\mathcal{X}}^p] \right) \quad (7.2.1)$$

The concept is called Martingale type because  $(\epsilon_n \mathbf{x}_n(\epsilon))_{n \in \mathbb{N}}$  is a martingale difference sequence and it can be shown that rate of convergence of martingales in Banach spaces is governed by the rate of convergence of martingales of the form  $Z_n = \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon)$  (which are incidentally called Walsh-Paley martingales). We point the reader to [65, 66] for more details. Further, for any  $p \in [1, 2]$  we also define,

$$C_p := \inf \left\{ C \mid \forall \mathbf{x}_0 \in \mathcal{B}^*, \forall \mathbf{x}, \sup_n \mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] \leq C^p \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{n \geq 1} \mathbb{E} \|\mathbf{x}_n(\epsilon)\|_{\mathcal{X}}^p \right) \right\}$$

$C_p$  is useful in determining if the pair  $(\mathcal{H}^*, \mathcal{X})$  has Martingale type  $p$ .

By the results in Chapter ?? (using it specifically for linear class) we have the following theorem:

**Theorem 66.** For any  $\mathcal{H} \in \mathcal{B}$  and any  $\mathcal{X} \in \mathcal{B}^*$  and any  $n \geq 1$ ,

$$\sup_{\mathbf{x}} \mathbb{E} \left[ \left\| \frac{1}{n} \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} \right] \leq \mathcal{V}_n(\mathcal{H}, \mathcal{X}) \leq 2 \sup_{\mathbf{x}} \mathbb{E} \left[ \left\| \frac{1}{n} \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} \right]$$

where the supremum above is over  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth.

Our main interest here will be in establishing that low regret implies Martingale type. To do so, we start with the above theorem to relate value of the online convex optimization game to rate of convergence of martingales in the Banach space. We then extend the result of Pisier in [65] to the “non-matching” setting combining it with the above theorem to finally get :

**Lemma 67.** If for some  $r \in (1, 2]$  there exists a constant  $D > 0$  such that for any  $n$ ,

$$\mathcal{V}_n(\mathcal{H}, \mathcal{X}) \leq D n^{-(1-\frac{1}{r})}$$

then for all  $s < r$ , we can conclude that any  $\mathbf{x}_0 \in \mathcal{B}^*$  and any  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth will satisfy :

$$\sup_n \mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^s \right] \leq \left( \frac{1104 D}{(r-s)^2} \right)^s \left( \|\mathbf{x}_0\|_{\mathcal{X}}^s + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^s] \right)$$

That is, the pair  $(\mathcal{H}, \mathcal{X})$  is of martingale type  $s$ .

The following corollary is an easy consequence of the above lemma.

**Corollary 68.** *For any  $p \in [1, 2]$  and any  $p' < p$  :  $C_{p'} \leq \frac{1104 V_p}{(p-p')^2}$*

### 7.3 Martingale Type and Uniform Convexity

The classical notion of Martingale type plays a central role in the study of geometry of Banach spaces. In [65], it was shown that a Banach space has Martingale type  $p$  (the classical notion) if and only if uniformly convex functions with certain properties exist on that space (w.r.t. the norm of that Banach space). In this section, we extend this result and show how the Martingale type of a pair  $(\mathcal{H}^*, \mathcal{X})$  are related to existence of certain uniformly convex functions. Specifically, the following theorem shows that the notion of Martingale type of pair  $(\mathcal{H}^*, \mathcal{X})$  is equivalent to the existence of a non-negative function that is uniformly convex w.r.t. the norm  $\|\cdot\|_{\mathcal{X}^*}$ .

**Lemma 69.** *If, for some  $p \in (1, 2]$ , there exists a constant  $C > 0$ , such that for all  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth and any  $\mathbf{x}_0 \in \mathcal{B}^*$ :*

$$\sup_n \mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] \leq C^p \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{n \geq 1} \mathbb{E} [\|\mathbf{x}_n(\epsilon)\|_{\mathcal{X}}^p] \right)$$

(i.e.  $(\mathcal{H}^*, \mathcal{X})$  has Martingale type  $p$ ), then there exists a convex function  $\Psi : \mathcal{B} \mapsto \mathbb{R}^+$  with  $\Psi(0) = 0$ , that is  $q$ -uniformly convex w.r.t. norm  $\|\cdot\|_{\mathcal{X}^*}$  s.t.  $\forall \mathbf{h} \in \mathcal{B}$ ,  $\frac{1}{q} \|\mathbf{h}\|_{\mathcal{X}^*}^q \leq \Psi(\mathbf{h}) \leq \frac{C^q}{q} \|\mathbf{h}\|_{\mathcal{H}}^q$ .

Define,

$$D_p := \inf \left\{ \left( \sup_{\mathbf{h} \in \mathcal{H}} \Psi(\mathbf{h}) \right)^{\frac{p-1}{p}} \mid \Psi : \mathcal{H} \mapsto \mathbb{R}^+ \text{ is } \frac{p}{p-1}\text{-uniformly convex w.r.t. } \|\cdot\|_{\mathcal{X}^*}, \Psi(0) = 0 \right\}$$

The following corollary follows directly from the above lemma.

**Corollary 70.** *For any  $p \in [1, 2]$ ,  $D_p \leq C_p$ .*

The proof of Lemma 69 goes further and gives a specific uniformly convex function  $\Psi$  satisfying the desired requirement (i.e. establishing  $D_p \leq C_p$ ) under the assumptions of the previous lemma:

$$\Psi_q^*(\mathbf{x}) := \sup \left\{ \frac{1}{C^p} \sup_n \mathbb{E} \left[ \left\| \mathbf{x} + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] - \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \right\}, \quad \Psi_q := (\Psi_q^*)^*. \quad (7.3.1)$$

where the supremum above is over  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth and  $p = \frac{q}{q-1}$ .

## 7.4 Main Result : Optimality of Online Mirror Descent

In previous chapter we argued that if we can find an appropriate uniformly convex function to use with the mirror descent algorithm, one can guarantee diminishing regret. However the pending question there was when one can find such a function and what is the rate one can guarantee. In Section ?? we introduced the extended notion of Martingale type of a pair  $(\mathcal{H}^*, \mathcal{X})$  and how it related to the value of the game. In Section ??, we saw how the concept of M-type related to existence of certain uniformly convex functions. We can now combine these results to show that the mirror descent algorithm is a universal online learning algorithm for convex learning problems. Specifically we show that whenever a problem is online learnable, the mirror descent algorithm can guarantee near optimal rates:

**Theorem 71.** *If for some constant  $V > 0$  and some  $q \in [2, \infty)$ ,  $\mathcal{V}_n(\mathcal{H}, \mathcal{X}) \leq V n^{-\frac{1}{q}}$  for all  $n$ , then for any  $n > e^{q-1}$ , there exists regularizer function  $\Psi$  and step-size  $\eta$ , such that the regret of the mirror descent algorithm using  $\Psi$  against any  $z_1, \dots, z_n \in \mathcal{Z}_{\text{Lip}}$  chosen by the adversary is bounded as:*

$$\mathbf{R}_n(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n) \leq 6002 V \log^2(n) n^{-\frac{1}{q}} \quad (7.4.1)$$

*Proof.* Combining Mirror descent guarantee in Lemma 55, Lemma 69 and the lower bound in Lemma 67 with  $s = \frac{q}{q-1} - \frac{1}{\log(n)}$  we get the above statement.  $\square$

The above Theorem tells us that, with appropriate  $\Psi$  and learning rate  $\eta$ , mirror descent will obtain regret at most a factor of  $6002 \log^2(n)$  from the best possible worst-case upper bound. We would like to point out that the constant  $V$  in the value of the game appears linearly and there is no other problem or space related hidden constants in the bound.

The following figure summarizes the relationship between the various constants. The arrow mark from  $C_{p'}$  to  $C_p$  indicates that for any  $n$ , all the quantities are within  $\log^2 n$  factor of each other.




Figure 7.1: Relationship between the various constants

We now provide some general guidelines that will help us in picking out appropriate function  $\Psi$  for mirror descent. First we note that though the function  $\Psi_q$  in the construction Eq. (7.3.1) need not be such that  $(q\Psi_q(\mathbf{h}))^{1/q}$  is a norm, with a simple modification as noted in [66] we can make it a norm. This basically tells us that the pair  $(\mathcal{H}, \mathcal{X})$  is online learnable, if and only if we can sandwich a  $q$ -uniformly convex norm in-between  $\mathcal{X}^*$  and a scaled version of  $\mathcal{H}$  (for some  $q < \infty$ ). Also note that by definition of uniform convexity, if any function  $\Psi$  is  $q$ -uniformly convex w.r.t. some norm  $\|\cdot\|$  and we have that  $\|\cdot\| \geq c \|\cdot\|_{\mathcal{X}}$ , then  $\frac{\Psi(\cdot)}{c^q}$  is  $q$ -uniformly convex w.r.t. norm  $\|\cdot\|_{\mathcal{X}}$ . These two observations together suggest that given pair  $(\mathcal{H}, \mathcal{X})$  what

we need to do is find a norm  $\|\cdot\|$  in between  $\|\cdot\|_{\mathcal{X}}^*$  and  $C\|\cdot\|_{\mathcal{H}}$  ( $C < \infty$ , smaller the  $C$  better the bound) such that  $\|\cdot\|^q$  is  $q$ -uniformly convex w.r.t  $\|\cdot\|$ .

### 7.4.1 Smooth Loss Case

**Lemma 72.** *Given any  $\bar{L}^* \in (0, \frac{3}{4}]$  and any  $n \in \mathbb{N}$ , for any online convex learning algorithm  $\mathbf{A}$ , there exists instances  $z_1, \dots, z_n \in \mathcal{Z}_{\text{smt}(1)}$  such that,  $\inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \leq \bar{L}^*$  and*

$$\mathbf{R}_n(\mathbf{A}, z_{1:n}) \geq \frac{1}{n} \sup_{\mathbf{x}} \mathbb{E}_{\epsilon} \left[ \left\| \sum_{t=1}^{n\bar{L}^*} \epsilon_t \mathbf{x}_t(\epsilon) \right\|_{\mathcal{H}^*} \right] \geq \frac{\bar{L}^*}{2} \mathcal{V}_{n\bar{L}^*}(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X}))$$

The following theorem shows the almost optimality of mirror descent for non-negative smooth convex objectives. The upper bound on the regret of the mirror descent algorithm also shows how having small  $\bar{L}^*$  helps reduce regret.

**Theorem 73.** *If for some  $V > 0$  and  $q \in [2, \infty)$  we have that*

$$\mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{smt}(1)}) \leq \frac{V}{n^{1/q}}$$

*then for any  $\bar{L}^* > 0$  and  $n > e^{q-1}$  there exists regularizer function  $\Psi$  and step-size  $\eta$  such that regret of the mirror descent algorithm with these  $\Psi$  and  $\eta$  against any  $z_1, \dots, z_n \in \mathcal{Z}_{\text{smt}(H)}$  s.t.  $\inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \leq \bar{L}^*$  is bounded as :*

$$\mathbf{R}_n(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n) \leq \frac{48016V\sqrt{H\bar{L}^*} \log^2 n}{n^{\frac{1}{q}}} + \frac{(37960V)^2 H \log^4 n}{n^{\frac{2}{q}}}$$

*Proof.* Using Lemma 72 with  $\bar{L}^* = \frac{3}{4}$  with the premise of the theorem we have that for any  $n$ ,

$$\mathcal{V}_{3n/4}(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X})) \leq \frac{8V}{3n^{1/q}} \leq \frac{8V}{3(3n/4)^{1/q}}$$

and so we can conclude that for any  $m$ ,  $\mathcal{V}_m(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X})) \leq \frac{8V}{3m^{1/q}}$ . Combining Mirror descent guarantee for smooth loss in Lemma 56, Lemma 69 and the lower bound in Lemma 67 with  $s = \frac{q}{q-1} - \frac{1}{\log(n)}$  we get the above statement.  $\square$

The above theorem shows that the mirror descent can achieve a near optimal rate in terms of dependence on  $n$ . However a closer look at the above Theorem and Lemma also reveals that we can in fact also capture a tighter dependence on  $\bar{L}^*$ . In fact when  $q = 2$ , (that is when we get  $1/\sqrt{n}$  rates), we can show that the dependence on  $\bar{L}^*$  of mirror descent is tight.

To see this assume that algorithm  $\mathbf{A}$  is a minimax optimal algorithm and is such that for any  $\bar{L}^* \in (0, 3/4]$  and  $n$  the regret of the algorithm against any  $z_1, \dots, z_n \in \mathcal{Z}_{\text{smt}(1)}$  s.t.  $\inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \leq \bar{L}^*$  is



bounded as :

$$\mathbf{R}(\mathbf{A}, z_1, \dots, z_n) \leq \frac{V\overline{L}^{*a}}{n^{1/q}} + \frac{V}{n^{1/b}}$$

where  $V, a, b > 0$  and  $q \in [2, \infty)$  with  $b < q$ .

Using the guarantee of the algorithm and Lemma 72 we see that

$$\mathcal{V}_{nL^*}(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X})) \leq \frac{2V\overline{L}^{*a-1}}{n^{1/q}} + \frac{2V}{\overline{L}^*n^{1/b}}$$

Hence using notation  $m = n\overline{L}^*$  we conclude that for any  $m$ ,

$$\mathcal{V}_m(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X})) \leq \frac{2V\overline{L}^{*a-1+\frac{1}{q}}}{m^{1/q}} + \frac{2V}{\overline{L}^{*1-\frac{1}{b}}m^{1/b}} \quad (7.4.2)$$

From the above inequality we can conclude that  $a \geq 1 - \frac{1}{q} = \frac{1}{p}$ . This is because if  $a < \frac{1}{p}$ , then since the above inequality holds for all  $\overline{L}^*$ , by picking  $\overline{L}^*$  to minimize the above inequality we can conclude that  $\mathcal{V}_m(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X})) = o\left(\frac{1}{m^{1/q}}\right)$ . Hence Theorem 73 we can conclude that regret of the mirror descent algorithm for any  $z_1, \dots, z_n \in \mathcal{Z}_{\text{smt}(1)}(\mathcal{X})$  is bounded by  $o\left(\frac{1}{m^{1/q}}\right)$  which is a contradiction since  $\mathbf{A}$  is minimax optimal but can only guarantee bound of order  $1/n^{1/q}$ . Hence we can conclude that  $a \leq 1/p$ .

Now note that by the guarantee of algorithm  $\mathbf{A}$  we can conclude that

$$\mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{smt}(1)}(\mathcal{X})) \leq \frac{2V}{n^{1/q}}$$

and so by Theorem 73 we can conclude that there exists regularizer function  $\Psi$  and step-size  $\eta$  such that regret of the mirror descent algorithm with these  $\Psi$  and  $\eta$  against any  $z_1, \dots, z_n \in \mathcal{Z}_{\text{smt}(H)}$  s.t.  $\inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \leq \overline{L}^*$  is bounded as :

$$\mathbf{R}_n(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n) \leq \frac{96032V\sqrt{H\overline{L}^*} \log^2 n}{n^{\frac{1}{q}}} + \frac{(75920V)^2 H \log^4 n}{n^{\frac{2}{q}}}$$

Thus we can conclude that while the best guarantee of any algorithm  $\mathbf{A}$  is,

$$\mathbf{R}(\mathbf{A}, z_1, \dots, z_n) \leq \frac{V\overline{L}^{*1/p}}{n^{1/q}} + \frac{V}{n^{1/b}},$$

the regret of mirror descent can be bounded as

$$\mathbf{R}_n(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n) \leq \frac{96032V\sqrt{H\overline{L}^*} \log^2 n}{n^{\frac{1}{q}}} + \frac{(75920V)^2 H \log^4 n}{n^{\frac{2}{q}}}$$

Notice that when  $q = 2$ , for any  $\overline{L}^* > \frac{V}{n^{1/q}}$ , guarantee of mirror descent is not only optimal in terms of  $n$  but also in terms of  $\overline{L}^*$ .

## 7.4.2 Uniformly Convex Loss Case

**Lemma 74.** For any  $\sigma > 0$  and  $q' \in (2, \infty)$  we have that,

$$\mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{ucvx}(\sigma, q')}(\mathcal{X})) \geq \frac{1}{2\sigma^{p'-1} D_{p'}^{p'}} \left( \sup_{\mathbf{u}} \mathbb{E}_{\epsilon} \left[ \left\| \frac{1}{n} \sum_{t=1}^n \epsilon_t \mathbf{u}_t(\epsilon) \right\|_{\mathcal{H}^*} \right] \right)^{p'} \geq \frac{1}{2\sigma^{p'-1} D_{p'}^{p'}} (\mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X})))^{p'}$$

where the supremum is over all  $\mathcal{X}$ -valued trees  $\mathbf{u}$  of depth  $n$ .

For simplicity we are going to assume now that  $\sigma = 1$ . The following theorem shows that for certain  $q'$ , mirror descent achieves optimal rate.

**Theorem 75.** Given a  $q'$ , if there exists  $V > 0$  and  $a > 0$  such that :

$$\mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{ucvx}(1, q')}) \leq \frac{V}{n^a}$$

then for any  $n > e^{q'-1}$  there exists regularizer function  $\Psi$  and step-size  $\eta$  such that regret of the mirror descent algorithm with these  $\Psi$  and  $\eta$  against any  $z_1, \dots, z_n \in \mathcal{Z}_{\text{ucvx}(1, q')}$  or the case when  $q' > 2$  is bounded as :

$$\mathbf{R}_n(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n) \leq \min \left\{ \frac{12008(2V)^{\frac{1}{p'}} D_{p'} \log^2 n}{n^{\frac{a}{p'}}}, \frac{2}{(2-p')n^{p'-1}} \right\} + \frac{\sup_{\mathbf{h} \in \mathcal{H}} R(\mathbf{h})}{n}.$$

and for  $q' = 2$  is bounded as :  $\mathbf{R}_n(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n) \leq \frac{2 \log n}{n} + \frac{\sup_{\mathbf{h} \in \mathcal{H}} R(\mathbf{h})}{n}$ .

Notice that when  $q' = \infty$ ,  $D_{p'} = D_1 = 1$  and so by above theorem,

$$\mathbf{R}_n(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n) \leq \frac{24016V \log^2 n}{n^a} + \frac{\sup_{\mathbf{h} \in \mathcal{H}} R(\mathbf{h})}{n}.$$

Hence for this case mirror descent is near optimal.

Next notice that by lower bound in lemma 74 and assumption of above theorem, we have

$$\mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X})) \leq \frac{(2V)^{\frac{1}{p'}} D_{p'}}{n^{\frac{a}{p'}}$$

Now define  $p^* = \sup\{p : V_p < \infty\}$ . By above inequality and definition of  $p^*$  we see that if for any  $p'$ ,  $\frac{a}{p'} > \frac{1}{q^*}$ , then it should be true that  $(2V)^{\frac{1}{p'}} D_{p'} = \infty$ . In other words, it has to be true that for any  $p'$ ,  $a \leq \frac{p'}{q^*}$ . Hence we can conclude that as  $p' \rightarrow p^*$ , any algorithm will have a regret of order at least  $\frac{1}{n^{\frac{p^*}{q^*}}} = \frac{1}{n^{p^*-1}}$ .

However mirror descent guarantee in Theorem 75 shows that mirror descent algorithm achieves regret at most

$$\mathbf{R}_n(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n) \leq \frac{2}{(2-p^*)n^{p^*-1}} + \frac{\sup_{\mathbf{h} \in \mathcal{H}} R(\mathbf{h})}{n}$$

when  $p^* < 2$  and

$$\mathbf{R}_n(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n) \leq \frac{2 \log n}{n} + \frac{\sup_{\mathbf{h} \in \mathcal{H}} R(\mathbf{h})}{n}$$

when  $p^* = 2$ . Hence we can conclude that for the case when  $p' \rightarrow p^*$  again mirror descent is optimal. (The reason we are taking limit of  $p' \rightarrow p^*$  is because one may not even be able to find  $q^*$  uniformly convex functions on the given space that is bounded). This specifically shows that mirror descent is near optimal for strongly convex objectives.

## 7.5 Examples

We demonstrate our results on several online learning problems, specified by  $\mathcal{H}$  and  $\mathcal{X}$ .

### 7.5.1 Example : $\ell_p$ non-dual pairs

It is usual in the literature to consider the case when  $\mathcal{H}$  is the unit ball of the  $\ell_p$  norm in some finite dimension  $d$  while  $\mathcal{X}$  is taken to be the unit ball of the dual norm  $\ell_q$  where  $p, q$  are Hölder conjugate exponents. Using the machinery developed in this chapter, it becomes effortless to consider the non-dual case when  $\mathcal{H}$  is the unit ball  $B_{p_1}$  of some  $\ell_{p_1}$  norm while  $\mathcal{X}$  is the unit ball  $B_{p_2}$  for arbitrary  $p_1, p_2$  in  $[1, \infty]$ . We shall use  $q_1$  and  $q_2$  to represent Holder conjugates of  $p_1$  and  $p_2$ . Before we proceed we first note that for any  $r \in (1, 2]$ ,  $\psi_r(\mathbf{h}) := \frac{1}{2(r-1)} \|\mathbf{h}\|_r^2$  is 2-uniformly w.r.t. norm  $\|\cdot\|_r$  (see for instance [61]). On the other hand by Clarkson's inequality, we have that for  $r \in (2, \infty)$ ,  $\psi_r(\mathbf{h}) := \frac{2^r}{r} \|\mathbf{h}\|_r^r$  is  $r$ -uniformly convex w.r.t.  $\|\cdot\|_r$ . Putting it together we see that for any  $r \in (1, \infty)$ , the function  $\psi_r$  defined above, is  $Q$ -uniformly convex w.r.t  $\|\cdot\|_r$  for  $Q = \max\{r, 2\}$ . The basic technique idea is to be to select  $\psi_r$  based on the guidelines in the end of the previous section. Finally we show that using  $\tilde{\psi}_r := d^{Q \max\{\frac{1}{q_2} - \frac{1}{r}, 0\}} \psi_r$  in Mirror descent Lemma 55 yields the bound that for any  $z_1, \dots, z_n \in \mathcal{Z}_{\text{Lip}}$ :

$$\mathbf{R}_n(\mathcal{A}_{\text{MD}}, z_1, \dots, z_n) \leq \frac{2 \max\{2, \frac{1}{\sqrt{2(r-1)}}\} d^{\max\{\frac{1}{q_2} - \frac{1}{r}, 0\} + \max\{\frac{1}{r} - \frac{1}{p_1}, 0\}}}{n^{1/\max\{r, 2\}}}$$

The following table summarizes the scenarios where a value of  $r = 2$ , i.e. a rate of  $D_2/\sqrt{n}$ , is possible, and lists the corresponding values of  $D_2$  (up to numeric constant of at most 16):

$p_1$ Range	$p_2$ Range	$D_2$	Treating as dual in $\ \cdot\ _{p_1}$
$1 \leq p_1 \leq 2$	$p_2 < 2$	1	$1/\sqrt{p_1 - 1}$
$1 \leq p_1 \leq 2$	$2 \leq p_2 \leq \frac{p_1}{p_1 - 1}$	$\sqrt{p_2 - 1}$	$1/\sqrt{p_1 - 1}$
$1 \leq p_1 \leq 2$	$\frac{p_1}{p_1 - 1} < p_2$	$d^{\frac{p_2 - 1}{p_2} - \frac{1}{p_1}} / \sqrt{p_1 - 1}$	$d^{\frac{p_2 - 1}{p_2} - \frac{1}{p_1}} / \sqrt{p_1 - 1}$
$p_1 > 2$	$p_2 < 2$	$d^{\frac{1}{2} - \frac{1}{p_1}}$	$d^{\frac{1}{2} - \frac{1}{p_1}}$
$p_1 > 2$	$p_2 \geq 2$	$d^{\frac{p_2 - 1}{p_2} - \frac{1}{p_1}}$	$d^{\frac{p_2 - 1}{p_2} - \frac{1}{p_1}}$
$1 \leq p_1 \leq 2$	$p_2 = \infty$	$\sqrt{\log(d)}$	$\sqrt{\log(d)}$

Note that the first two rows are dimension free, and so apply also in infinite-dimensional settings, whereas in

the other scenarios,  $D_2$  is finite only when the dimension is finite. An interesting phenomena occurs when  $d$  is  $\infty$ ,  $p_1 > 2$  and  $q_2 \geq p_1$ . In this case  $D_2 = \infty$  and so one cant expect a rate of  $O(\frac{1}{\sqrt{n}})$ . However we have  $D_{p_2} < 16$  and so can still get a rate of  $n^{-\frac{1}{q_2}}$ .

Ball et al [67] tightly calculate the constants of strong convexity of squared  $\ell_p$  norms, establishing the tightness of  $D_2$  when  $p_1 = p_2$ . By extending their constructions it is also possible to show tightness (up to a factor of 16) for all other values in the table. Also, Agarwal et al [68] recently showed lower bounds on the sample complexity of stochastic optimization when  $p_1 = \infty$  and  $p_2$  is arbitrary—their lower bounds match the last two rows in the table.

### 7.5.2 Example : Non-dual Schatten norm pairs in finite dimensions

Exactly the same analysis as above can be carried out for Schatten  $p$ -norms, i.e. when  $\mathcal{H} = B_{S(p_1)}$ ,  $\mathcal{X} = B_{S(p_2)}$  are the unit balls of Schatten  $p$ -norm (the  $p$ -norm of the singular values) for matrix of dimensions  $d_1 \times d_2$ . We get the same results as in the table above (as upper bounds on  $D_2$ ), with  $d = \min\{d_1, d_2\}$ . These results again follow using similar arguments as  $\ell_p$  case and tight constants for strong convexity parameters of the Schatten norm from [67].

### 7.5.3 Example : Non-dual group norm pairs in finite dimensions

In applications such as multitask learning, groups norms such as  $\|\mathbf{h}\|_{q,1}$  are often used on matrices  $\mathbf{h} \in \mathbb{R}^{k \times d}$  where  $(q, 1)$  norm means taking the  $\ell_1$ -norm of the  $\ell_q$ -norms of the columns of  $\mathbf{h}$ . Popular choices include  $q = 2, \infty$ . Here, it may be quite unnatural to use the dual norm  $(p, \infty)$  to define the space  $\mathcal{X}$  where the data lives. For instance, we might want to consider  $\mathcal{H} = B_{(q,1)}$  and  $\mathcal{X} = B_{(\infty,\infty)} = B_\infty$ . In such a case we can calculate that  $D_2(\mathcal{H}, \mathcal{X}) = \Theta(k^{1-\frac{1}{q}} \sqrt{\log(d)})$  using  $\Psi(\mathbf{h}) = \frac{1}{q+r-2} \|\mathbf{h}\|_{q,r}^2$  where  $r = \frac{\log d}{\log d-1}$ .

### 7.5.4 Example : Max Norm

Max-norm has been proposed as a convex matrix regularizer for application such as matrix completion [69]. In the online version of the matrix completion problem at each time step one element of the matrix is revealed, corresponding to  $\mathcal{X}$  being the set of all matrices with a single element being 1 and the rest 0. Since we need  $\mathcal{X}$  to be convex we can take the absolute convex hull of this set and use  $\mathcal{X}$  to be the unit element-wise  $\ell_1$  ball. Its dual is  $\|W\|_{\mathcal{X}^*} = \max_{i,j} |W_{i,j}|$ . On the other hand given a matrix  $W$ , its max-norm is given by  $\|W\|_{\max} = \min_{U,V:W=UV^T} (\max_i \|U_i\|_2) (\max_j \|V_j\|_2)$ . The set  $\mathcal{H}$  is the unit ball under the max norm. As noted in [70] the max-norm ball is equivalent, up to a factor two, to the convex hull of all rank one sign matrices. Let us now make a more general observation.

**Proposition 76.** *Let  $\mathcal{H} = \text{abscvx}(\{\mathbf{h}_1, \dots, \mathbf{h}_K\})$ . The Minkowski norm for this  $\mathcal{H}$  is given by*

$$\|\mathbf{h}\|_{\mathcal{H}} := \inf_{\alpha_1, \dots, \alpha_K: \mathbf{h} = \sum_{i=1}^K \alpha_i \mathbf{h}_i} \sum_{i=1}^K |\alpha_i|$$

In this case, for any  $q \in (1, 2]$ , if we define the norm :

$$\|\mathbf{h}\|_{\mathcal{H},q} = \inf_{\alpha_1, \dots, \alpha_K: \mathbf{h} = \sum_{i=1}^K \alpha_i \mathbf{h}_i} \left( \sum_{i=1}^K |\alpha_i|^q \right)^{1/q} \quad (7.5.1)$$

then the function  $\Psi(\mathbf{h}) = \frac{1}{2(q-1)} \|\mathbf{h}\|_{\mathcal{H},q}^2$  is 2-uniformly convex w.r.t.  $\|\cdot\|_{\mathcal{H},q}$ . Further if we use  $q = \frac{\log K}{\log K - 1}$ , then  $\sup_{\mathbf{h} \in \mathcal{H}} \sqrt{\Psi(\mathbf{h})} = O(\sqrt{\log K})$ .

Proof of the above proposition is similar to proof of strong convexity of  $\ell_q$  norms. For the max norm case as noted before the norm is equivalent to the norm got by the taking the absolute convex hull of the set of all rank one sign matrices. Cardinality of this set is of course  $2^{N+M}$ . Hence using the above proposition and noting that  $\mathcal{X}^*$  is the unit ball of  $\|\cdot\|_\infty$  we see that  $\Psi$  is obviously 2-uniformly convex w.r.t.  $\|\cdot\|_{\mathcal{X}^*}$  and so we get a regret bound  $O\left(\sqrt{\frac{M+N}{n}}\right)$ . This matches the stochastic (PAC) learning guarantee [70], and is the first guarantee we are aware of for the max norm matrix completion problem in the online setting.

### 7.5.5 Example : Interpolation Norms

Another interesting setting is when the set  $\mathcal{H}$  is got by interpolating between unit balls of two other norms  $\|\cdot\|_{\mathcal{H}_1}$  and  $\|\cdot\|_{\mathcal{H}_2}$ . Specifically one can consider  $\mathcal{H}$  to be the unit ball of two such interpolated norms, the first type of interpolation norm is given by,

$$\|\mathbf{h}\|_{\mathcal{H}} = \|\mathbf{h}\|_{\mathcal{H}_1} + \|\mathbf{h}\|_{\mathcal{H}_2} \quad (7.5.2)$$

The second type of interpolation norm one can consider is given by

$$\|\mathbf{h}\|_{\mathcal{H}} = \inf_{\mathbf{h}_1 + \mathbf{h}_2 = \mathbf{h}} (\|\mathbf{h}_1\|_{\mathcal{H}_1} + \|\mathbf{h}_2\|_{\mathcal{H}_2}) \quad (7.5.3)$$

In learning problems such interpolation norms are often used to induce certain structures or properties into the regularization. For instance one might want sparsity along with grouping effect in the linear predictors for which elastic-net type regularization introduced by Zou and Hastie [71] (this is captured by interpolation of the first type between  $\ell_1$  and  $\ell_2$  norms). Another example is in matrix completion problems when we would like the predictor matrix to be decomposable into sum of sparse and low rank matrices as done by Chandrasekaran et. al [72] (here one can use the interpolation norm of second type to interpolate between trace norm and element wise  $\ell_1$  norm). Another example where interpolation norms of type two are useful are in multi-task learning problems (with linear predictors) as done by Jalali et. al [73]. The basic idea is that the matrix of linear predictors can be decomposed into sum of two matrices one with for instance low entry-wise  $\ell_1$  norm and other with low  $B_{(2,\infty)}$  group norm (group sparsity).

While in these applications the set  $\mathcal{H}$  used is obtained through interpolation norms, it is typically not natural for the set  $\mathcal{X}$  to be the dual ball of  $\mathcal{H}$  but rather something more suited to the problem at hand. For instance, for the elastic net regularization case, the set  $\mathcal{X}$  usually considered are either the vectors with bounded  $\ell_\infty$

norm or bounded  $\ell_2$ . Similarly for the [73] case  $\mathcal{X}$  could be either matrices with bounded entries or some other natural assumption that suits the problem.

It can be shown that in general for any interpolation norm of first type specified in Equation 7.5.2,

$$D_2(\mathcal{H}, \mathcal{X}) \leq 2 \min\{D_2(\mathcal{H}_1, \mathcal{X}), D_2(\mathcal{H}_2, \mathcal{X})\} \quad (7.5.4)$$

Similarly for the interpolation norm of type two one can in general show that,

$$D_2(\mathcal{H}, \mathcal{X}) \leq \frac{1}{2} \max\{D_2(\mathcal{H}_1, \mathcal{X}), D_2(\mathcal{H}_2, \mathcal{X})\} \quad (7.5.5)$$

Using the above bounds one can get regret bounds for mirror descent algorithm with appropriate  $\Psi$  and step size  $\eta$  for specific examples like the ones mentioned.

The bounds given in Equations Eq. (7.5.4) and Eq. (7.5.5) are only upper bounds and it would be interesting to analyze these cases in more detail and also to analyze interpolation between several norms instead of just two.

## 7.6 Detailed Proofs

*Proof of Lemma 63.* Let  $\mathbf{A} : \bigcup_{n \in \mathbb{N}} \mathcal{X}^n \mapsto \mathcal{H}$  be any proper learning algorithm for the linear learning problem specified by instance set  $\mathcal{Z}_{\text{lin}}(\mathcal{X})$ . One can use this learning algorithm to construct a new gradient-based learning algorithm  $\mathbf{A}^{\mathcal{O}^{\text{1st}}} : \bigcup_{n \in \mathbb{N}} \mathcal{X}^n \mapsto \mathcal{H}$  as follows, for any  $t \in \mathbb{N}$  and any  $\mathbf{x}_1, \dots, \mathbf{x}_t \in \mathcal{X}$ ,

$$\mathbf{A}^{\mathcal{O}^{\text{1st}}}(z_1, \dots, z_t) = \mathbf{A}(\nabla \ell(\mathbf{h}_1, z_1), \dots, \nabla \ell(\mathbf{h}_t, z_t)).$$

That is, at the end of each round the learning algorithm for linear problem is fed with a sub-gradient of the loss at the hypothesis  $\mathbf{h}_t$  played on that round  $t$  and thus  $\mathbf{h}_{t+1}$  is selected using this algorithm. Note that for

any  $\mathbf{h}^* \in \mathcal{H}$ , by convexity of the loss function,

$$\begin{aligned}
& \frac{1}{n} \sum_{t=1}^n \ell \left( \mathbf{A}^{\mathcal{O}^{\text{1st}}} (\nabla \ell(\mathbf{h}_1, z_1), \dots, \nabla \ell(\mathbf{h}_{t-1}, z_{t-1})), z_t \right) - \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}^*, z_t) \\
& \leq \frac{1}{n} \sum_{t=1}^n \left\langle \nabla \ell \left( \mathbf{A}^{\mathcal{O}^{\text{1st}}} (\nabla \ell(\mathbf{h}_1, z_1), \dots, \nabla \ell(\mathbf{h}_{t-1}, z_{t-1})), z_t \right), \mathbf{A}^{\mathcal{O}^{\text{1st}}} (\nabla \ell(\mathbf{h}_1, z_1), \dots, \nabla \ell(\mathbf{h}_{t-1}, z_{t-1})) - \mathbf{h}^* \right\rangle \\
& = \frac{1}{n} \sum_{t=1}^n \left\langle \mathbf{x}_t^*, \mathbf{A}(\mathbf{x}_1^*, \dots, \mathbf{x}_{t-1}^*) - \mathbf{h}^* \right\rangle \\
& \leq \frac{1}{n} \sum_{t=1}^n \left\langle \mathbf{x}_t^*, \mathbf{A}(\mathbf{x}_1^*, \dots, \mathbf{x}_{t-1}^*) \right\rangle - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \left\langle \mathbf{x}_t^*, \mathbf{h} \right\rangle \\
& \leq \sup_{\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathcal{X}} \left\{ \frac{1}{n} \sum_{t=1}^n \left\langle \mathbf{x}_t, \mathbf{A}(\mathbf{x}_1, \dots, \mathbf{x}_{t-1}) \right\rangle - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \left\langle \mathbf{x}_t, \mathbf{h} \right\rangle \right\}
\end{aligned}$$

where in the above we used the notation  $\mathbf{x}_t^* = \nabla \ell(\mathbf{h}_t, z_t)$ . Thus we can conclude that the new first-order oracle-based learning algorithm enjoys the same regret guarantee as the algorithm  $\mathbf{A}$  enjoys on linear learning problems. From this we conclude the lemma statement.  $\square$

*Proof of Corollary 64.* First by Lemma 63 we have that for any algorithm  $\mathbf{A}$  for linear learning problem, there exists an oracle based learning algorithm  $\mathbf{A}^{\mathcal{O}^{\text{1st}}}$  such that

$$\sup_{z_1, \dots, z_n \in \mathcal{Z}} \mathbf{R}_n(\mathbf{A}^{\mathcal{O}^{\text{1st}}}, z_{1:n}) \leq \sup_{z_1^*, \dots, z_n \in \mathcal{Z}_{\text{lin}}} \mathbf{R}_n(\mathbf{A}, z_{1:n}^*)$$

and so

$$\mathcal{V}_n(\mathcal{H}, \mathcal{Z}) \leq \mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X}))$$

As for the set of equalities for specific classes,  $\mathcal{Z}_{\text{sup}}$ ,  $\mathcal{Z}_{\text{Lip}}$  and  $\mathcal{Z}_{\text{lin}}$ , note that since each of these classes are such that sub-gradients of losses belong to  $\mathcal{X}$ , we have that

$$\mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{sup}}(\mathcal{X})) \leq \mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{Lip}}(\mathcal{X})) \leq \mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X}))$$

On the other hand,  $\mathcal{Z}_{\text{lin}}(\mathcal{X}) \subset \mathcal{Z}_{\text{Lip}}(\mathcal{X})$  and so  $\mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{Lip}}(\mathcal{X})) \geq \mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X}))$  and hence we can conclude that

$$\mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{Lip}}(\mathcal{X})) = \mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X})).$$

Similarly for the supervised learning problem specified by instance set  $\mathcal{Z}_{\text{sup}}$ , note that if adversary always

pick targets  $y_t = -b = -\sup_{\mathbf{h} \in \mathcal{H}, \mathbf{x} \in \mathcal{X}} \langle \mathbf{h}, \mathbf{x} \rangle$  then for any  $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathcal{X}$ ,

$$\begin{aligned} \frac{1}{n} \sum_{t=1}^n |\langle \mathbf{x}_t, \mathbf{h}_t \rangle - y_t| - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n |\langle \mathbf{x}_t, \mathbf{h} \rangle - y_t| &= \frac{1}{n} \sum_{t=1}^n \langle \mathbf{x}_t, \mathbf{h}_t \rangle + b - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \left( \sum_{t=1}^n \langle \mathbf{x}_t, \mathbf{h} \rangle + b \right) \\ &= \frac{1}{n} \sum_{t=1}^n \langle \mathbf{x}_t, \mathbf{h}_t \rangle - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \langle \mathbf{x}_t, \mathbf{h} \rangle \end{aligned}$$

Thus we can conclude that  $\mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{sup}}(\mathcal{X})) \geq \mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X}))$  and so,  $\mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{sup}}(\mathcal{X})) = \mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X}))$ . Thus we conclude the corollary.  $\square$

**Lemma 77.** *Let  $1 < p \leq 2$  and  $C > 0$  be fixed constants, the following statements are equivalent :*

1. *For all  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth and any  $\mathbf{x}_0 \in \mathcal{B}^*$ :*

$$\sup_n \mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] \leq C^p \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{n \geq 1} \mathbb{E} [\|\mathbf{x}_n(\epsilon)\|_{\mathcal{X}}^p] \right)$$

2. *There exist a non-negative convex function  $\Psi$  on  $\mathcal{B}$  with  $\Psi(0) = 0$ , that is  $q$ -uniformly convex w.r.t. norm  $\|\cdot\|_{\mathcal{X}^*}$  and for any  $\mathbf{h} \in \mathcal{B}$ ,  $\frac{1}{q} \|\mathbf{h}\|_{\mathcal{X}^*}^q \leq \Psi(\mathbf{h}) \leq \frac{C^q}{q} \|\mathbf{h}\|_{\mathcal{H}}^q$ .*

*Proof.* For any  $\mathbf{x} \in \mathcal{B}^*$  define  $\Psi^* : \mathcal{B}^* \mapsto \mathbb{R}$  as

$$\Psi^*(\mathbf{x}) := \sup \left\{ \left( \frac{1}{C^p} \sup_n \mathbb{E} \left[ \left\| \mathbf{x} + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] - \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \right) \right\}$$

where the supremum is over  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth such that,  $\sup_n \mathbb{E} [\|\mathbf{x} + \sum_{i=1}^n \mathbf{x}_i\|_{\mathcal{H}^*}^p] < \infty$ . Since supremum of convex functions is a convex function, it is easily verified that  $\Psi^*(\cdot)$  is convex. Note that by the definition of M-type in Equation 7.2.1, we have that for any  $\mathbf{x}_0 \in \mathcal{B}^*$ ,  $\Psi^*(\mathbf{x}_0) \leq \|\mathbf{x}_0\|_{\mathcal{X}}^p$ . On the other hand, note that by considering the sequence of constant mappings,  $\mathbf{x}_i = 0$  for all  $i \geq 1$ , we get that for any  $\mathbf{x}_0 \in \mathcal{B}^*$ ,

$$\Psi^*(\mathbf{x}_0) = \sup \left\{ \left( \frac{1}{C^p} \sup_n \mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] - \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \right) \right\} \geq \frac{1}{C^p} \|\mathbf{x}_0\|_{\mathcal{H}^*}^p$$

Thus we can conclude that for any  $\mathbf{x} \in \mathcal{B}^*$ ,  $\frac{1}{C^p} \|\mathbf{x}\|_{\mathcal{H}^*}^p \leq \Psi^*(\mathbf{x}) \leq \|\mathbf{x}\|_{\mathcal{X}}^p$ .

For any  $\mathbf{x}_0, \mathbf{y}_0 \in \mathcal{B}^*$ , by definition of  $\Psi^*(\mathbf{x}_0)$  and  $\Psi^*(\mathbf{y}_0)$ , for any  $\gamma > 0$ , there exist  $\mathcal{B}^*$ -valued trees  $\mathbf{x}$  and



$\mathbf{y}$  of infinite depth s.t. :

$$\Psi^*(\mathbf{x}_0) \leq \left( \frac{1}{C^p} \sup_n \mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] - \sum_{i \geq 1} \mathbb{E} [\| \mathbf{x}_i(\epsilon) \|_{\mathcal{X}}^p] \right) + \gamma$$

and

$$\Psi^*(\mathbf{y}_0^{(j)}) \leq \left( \frac{1}{C^p} \sup_n \mathbb{E} \left[ \left\| \mathbf{y}_0 + \sum_{i=1}^n \epsilon_i \mathbf{y}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] - \sum_{i \geq 1} \mathbb{E} [\| \mathbf{y}_i(\epsilon) \|_{\mathcal{X}}^p] \right) + \gamma$$

In fact in the above two inequalities if the supremum over  $n$  were achieved at some finite  $n_0$ , by replacing the original sequence by one which is identical up to  $n_0$  and for any  $i > n_0$  using  $\mathbf{x}_i(\epsilon) = 0$  (and similarly  $\mathbf{y}_i(\epsilon) = 0$ ), we can in fact conclude that using these  $\mathbf{x}$ 's and  $\mathbf{y}$ 's instead,

$$\Psi^*(\mathbf{x}_0) \leq \left( \frac{1}{C^p} \mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i \geq 1} \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] - \sum_{i \geq 1} \mathbb{E} [\| \mathbf{x}_i(\epsilon) \|_{\mathcal{X}}^p] \right) + \gamma \quad (7.6.1)$$

and

$$\Psi^*(\mathbf{y}_0^{(j)}) \leq \left( \frac{1}{C^p} \mathbb{E} \left[ \left\| \mathbf{y}_0 + \sum_{i \geq 1} \epsilon_i \mathbf{y}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] - \sum_{i \geq 1} \mathbb{E} [\| \mathbf{y}_i(\epsilon) \|_{\mathcal{X}}^p] \right) + \gamma \quad (7.6.2)$$

Now consider a sequence formed by taking  $z_0 = \frac{\mathbf{x}_0 + \mathbf{y}_0}{2}$  and further let

$$z_1 = \left( \frac{1 + \epsilon_0}{2} \right) \frac{\mathbf{x}_0 - \mathbf{y}_0}{2} + \left( \frac{1 - \epsilon_0}{2} \right) \frac{\mathbf{y}_0 - \mathbf{x}_0}{2} = \epsilon_0 (\mathbf{x}_0 - \mathbf{y}_0)$$

and for any  $i \geq 2$ , define

$$z_i = \left( \frac{1 + \epsilon_0}{2} \right) \epsilon_{i-1} \mathbf{x}_{i-1}(\epsilon) + \left( \frac{1 - \epsilon_0}{2} \right) \epsilon_{i-1} \mathbf{y}_{i-1}(\epsilon)$$

where  $\epsilon_0 \in \{\pm 1\}$  is drawn uniformly at random. That is essentially at time  $i = 1$  we flip a coin and decide

to go with tree  $\mathbf{x}$  with probability 1/2 and  $\mathbf{y}$  with probability 1/2. Clearly using the tree  $\mathbf{z}$ , we have that,

$$\begin{aligned}
\Psi^* \left( \frac{\mathbf{x}_0 + \mathbf{y}_0}{2} \right) &= \sup_{\mathbf{z}} \left\{ \left( \frac{1}{C^p} \sup_n \mathbb{E} \left[ \left\| \frac{\mathbf{x}_0 + \mathbf{y}_0}{2} + \sum_{i=1}^n \mathbf{z}_i(\epsilon_0, \epsilon) \right\|_{\mathcal{H}^*}^p \right] - \sum_{i \geq 1} \mathbb{E} [\| \mathbf{z}_i(\epsilon_0, \epsilon) \|_{\mathcal{X}}^p] \right)^{1/p} \right\}^p \\
&\geq \frac{1}{C^p} \mathbb{E} \left[ \left\| \mathbf{z}_0 + \sum_{i \geq 1} \mathbf{z}_i(\epsilon_0, \epsilon) \right\|_{\mathcal{H}^*}^p \right] - \sum_{i \geq 1} \mathbb{E} [\| \mathbf{z}_i(\epsilon_0, \epsilon) \|_{\mathcal{X}}^p] \\
&= \frac{1}{C^p} \frac{\mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i \geq 1} \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] + \mathbb{E} \left[ \left\| \mathbf{y}_0 + \sum_{i \geq 1} \epsilon_i \mathbf{y}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right]}{2} - \sum_{i \geq 1} \mathbb{E} [\| \mathbf{z}_i(\epsilon_0, \epsilon) \|_{\mathcal{X}}^p] \\
&= \frac{1}{C^p} \frac{\mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i \geq 1} \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] + \mathbb{E} \left[ \left\| \mathbf{y}_0 + \sum_{i \geq 1} \epsilon_i \mathbf{y}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right]}{2} - \sum_{i \geq 1} \mathbb{E} [\| \mathbf{z}_i(\epsilon_0, \epsilon) \|_{\mathcal{X}}^p] \\
&= \frac{1}{C^p} \frac{\mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i \geq 1} \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] + \mathbb{E} \left[ \left\| \mathbf{y}_0 + \sum_{i \geq 1} \epsilon_i \mathbf{y}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right]}{2} - \left\| \frac{\mathbf{x}_0 - \mathbf{y}_0}{2} \right\|_{\mathcal{X}}^p - \sum_{i \geq 2} \mathbb{E} [\| \mathbf{z}_i(\epsilon_0, \epsilon) \|_{\mathcal{X}}^p] \\
&= \frac{1}{C^p} \frac{\mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i \geq 1} \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] + \mathbb{E} \left[ \left\| \mathbf{y}_0 + \sum_{i \geq 1} \epsilon_i \mathbf{y}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right]}{2} - \left\| \frac{\mathbf{x}_0 - \mathbf{y}_0}{2} \right\|_{\mathcal{X}}^p \\
&\quad - \sum_{i \geq 1} \frac{\mathbb{E} [\| \mathbf{x}_i(\epsilon) \|_{\mathcal{X}}^p] + \mathbb{E} [\| \mathbf{y}_i(\epsilon) \|_{\mathcal{X}}^p]}{2} \\
&= \frac{\frac{1}{C^p} \mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i \geq 1} \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] - \sum_{i \geq 1} \mathbb{E} [\| \mathbf{x}_i(\epsilon) \|_{\mathcal{X}}^p] + \frac{1}{C^p} \mathbb{E} \left[ \left\| \mathbf{y}_0 + \sum_{i \geq 1} \epsilon_i \mathbf{y}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] - \sum_{i \geq 1} \mathbb{E} [\| \mathbf{y}_i(\epsilon) \|_{\mathcal{X}}^p]}{2} \\
&\quad - \left\| \frac{\mathbf{x}_0 - \mathbf{y}_0}{2} \right\|_{\mathcal{X}}^p} \\
&\geq \frac{\Psi^*(\mathbf{x}_0) + \Psi^*(\mathbf{y}_0)}{2} - \left\| \frac{\mathbf{x}_0 - \mathbf{y}_0}{2} \right\|_{\mathcal{X}}^p - \gamma
\end{aligned}$$

where the last step is obtained by using Equations 7.6.1 and 7.6.2. Since  $\gamma$  was arbitrary taking limit we conclude that for any  $\mathbf{x}_0$  and  $\mathbf{y}_0$ ,

$$\frac{\Psi^*(\mathbf{x}_0) + \Psi^*(\mathbf{y}_0)}{2} \leq \Psi^* \left( \frac{\mathbf{x}_0 + \mathbf{y}_0}{2} \right) + \left\| \frac{\mathbf{x}_0 - \mathbf{y}_0}{2} \right\|_{\mathcal{X}}^p$$

Hence we have shown the existence of a convex function  $\Psi^*$  that is  $p$ -uniformly smooth w.r.t. norm  $\|\cdot\|_{\mathcal{X}}$  such that  $\frac{1}{C^p} \|\cdot\|_{\mathcal{H}^*}^p \leq \Psi^*(\cdot) \leq \|\cdot\|_{\mathcal{X}}^p$ . Using convex duality we can conclude that the convex conjugate  $\Psi$  of function  $\Psi^*$ , is  $q$ -uniformly convex w.r.t. norm  $\|\cdot\|_{\mathcal{X}^*}$  and is such that  $\|\cdot\|_{\mathcal{X}}^q \leq \Psi(\cdot) \leq C^q \|\cdot\|_{\mathcal{H}^*}^q$ . That 2 implies 1 can be easily verified using the smoothness property of  $\Psi^*$ . □

The following sequence of four lemma's give us the essentials towards proving Lemma 67. They use similar

techniques as in [65].

**Lemma 78.** *Let  $1 < r \leq 2$ . If there exists a constant  $D > 0$  such that any  $\mathbf{x}_0 \in \mathcal{B}^*$  and any  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth satisfies :*

$$\forall n \in \mathbb{N}, \quad \mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} \right] \leq D(n+1)^{1/r} \sup_{0 \leq i \leq n} \sup_{\epsilon} \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}$$

then for all  $p < r$  and  $\alpha_p = \frac{20D}{r-p}$  we can conclude that any  $\mathbf{x}_0 \in \mathcal{B}^*$  and any  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth will satisfy :

$$\sup_n \mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} \right] \leq \alpha_p \sup_{\epsilon} \left( \sum_{i \geq 0} \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p \right)^{1/p}$$

*Proof.* To begin with note that in the definition of type, if the supremum over  $n$  were achieved at some finite  $n_0$ , then by replacing the original sequence by one which is identical up to  $n_0$  and then on for any  $i > n_0$  using  $\mathbf{x}_i(\epsilon) = 0$  would only tighten the inequality. Hence it suffices to only consider such sequences. Further to prove the statement we only need to consider finite such sequences (ie. sequences such that there exists some  $n$  so that for any  $i > n$ ,  $\mathbf{x}_i = 0$ ) and show that the inequality holds for every such  $n$  (every such sequence).

Restricting ourselves to such finite sequences, we now use the shorthand,

$S = \sup_{\epsilon} \left( \sum_{i=0}^n \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p \right)^{1/p}$ . Now define

$$\begin{aligned} I_k(\epsilon) &= \left\{ i \geq 0 \mid \frac{S}{2^{(k+1)/p}} < \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}} \leq \frac{S}{2^{k/p}} \right\}, \\ T_0^{(k)}(\epsilon) &= \inf\{i \in I_k(\epsilon)\} \text{ and} \\ \forall m \in \mathbb{N}, T_m^{(k)}(\epsilon) &= \inf\{i > T_{m-1}^{(k)}(\epsilon), i \in I_k(\epsilon)\} \end{aligned}$$

Note that for any  $\epsilon \in \{\pm 1\}^{\mathbb{N}}$ ,

$$S^p \geq \sum_{i \in I_k(\epsilon)} \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p > \frac{S^p |I_k(\epsilon)|}{2^{(k+1)}}$$

and so we get that  $\sup_{\epsilon} |I_k(\epsilon)| < 2^{k+1}$ . From this we conclude that

$$\begin{aligned}
\mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} \right] &\leq \sum_{k \geq 0} \mathbb{E} \left[ \left\| \sum_{i \in I_k(\epsilon)} \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} \right] \\
&= \sum_{k \geq 0} \mathbb{E} \left[ \left\| \sum_{i \geq 0} \epsilon_{T_i^{(k)}(\epsilon)} \mathbf{x}_{T_i^{(k)}(\epsilon)} \right\|_{\mathcal{H}^*} \right] \\
&\leq \sum_{k \geq 0} \left( D \sup_{\epsilon} \{|I_k(\epsilon)|^{1/r}\} \sup_{\epsilon} \{ \sup_{i \in I_k(\epsilon)} \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}} \} \right) \\
&\leq \sum_{k \geq 0} \left( D 2^{(k+1)/r} \sup_{\epsilon} \sup_{i \in I_k(\epsilon)} \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}, \infty} \right) \\
&\leq \sum_{k \geq 0} \left( D 2^{(k+1)/r} 2^{-k/p} S \right) \\
&= D 2^{1/r} \sum_{k \geq 0} 2^{k(\frac{1}{r} - \frac{1}{p})} S \\
&\leq \frac{2D}{1 - 2^{(\frac{1}{r} - \frac{1}{p})}} S \\
&\leq \frac{2D}{1 - 2^{-(r-p)/4}} S \\
&\leq \frac{12D}{r-p} S \\
&= \alpha_p \sup_{\epsilon} \left( \sum_{i=0}^n \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p \right)^{1/p}
\end{aligned}$$

□

**Lemma 79.** *Let  $1 < r \leq 2$ . If there exists a constant  $D > 0$  such that any  $\mathbf{x}_0 \in \mathcal{B}^*$  and any  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth satisfies :*

$$\forall n \in \mathbb{N}, \quad \mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} \right] \leq D(n+1)^{1/r} \sup_{0 \leq i \leq n} \sup_{\epsilon} \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}$$

*then for any  $p < r$ , any  $\mathbf{x}_0 \in \mathcal{B}^*$  and any  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth :*

$$\mathbb{P} \left( \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} > c \right) \leq 2 \left( \frac{\alpha_p}{c} \right)^{p/(p+1)} \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \right)^{1/(p+1)}$$

*Proof.* For any  $\mathbf{x}_0 \in \mathcal{B}^*$  and  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth define

$$V_n(\epsilon) = \sum_{i=0}^n \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p$$

For appropriate choice of  $a > 0$  to be fixed later, define stopping time

$$\tau(\epsilon) = \inf \{n \geq 0 | V_{n+1} > a^p\}$$

Now for any  $c > 0$  we have,

$$\begin{aligned} \mathbb{P} \left( \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} > c \right) &\leq \mathbb{P}(\tau(\epsilon) < \infty) + \mathbb{P} \left( \tau(\epsilon) = \infty, \sup_n \left\| \sum_{i=0}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} > c \right) \\ &\leq \mathbb{P}(\tau(\epsilon) < \infty) + \mathbb{P} \left( \tau(\epsilon) > 0, \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^{n \wedge \tau(\epsilon)} \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} > c \right) \end{aligned} \quad (7.6.3)$$

As for the first term in the above equation note that

$$\mathbb{P}(\tau(\epsilon) < \infty) = \mathbb{P}(\sup_n V_n > a^p) \leq \frac{\|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p]}{a^p} \quad (7.6.4)$$

To consider the second term of Equation 7.6.3 we note that  $(\mathbf{1}_{\{\tau(\epsilon) > 0\}}(\mathbf{x}_0 + \sum_{i=1}^{n \wedge \tau(\epsilon)} \epsilon_i \mathbf{x}_i(\epsilon)))_{n \geq 0}$  is a valid martingale (stopped process) and hence,  $(\|\mathbf{1}_{\{\tau(\epsilon) > 0\}}(\mathbf{x}_0 + \sum_{i=1}^{n \wedge \tau(\epsilon)} \epsilon_i \mathbf{x}_i(\epsilon))\|_{\mathcal{H}^*})_{n \geq 0}$  is a sub-martingale. Hence by Doob's inequality we conclude that,

$$\mathbb{P} \left( T > 0, \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^{n \wedge \tau(\epsilon)} \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} > c \right) \leq \frac{1}{c} \sup_n \mathbb{E} \left[ \left\| \mathbf{1}_{\{\tau(\epsilon) > 0\}} \left( \mathbf{x}_0 + \sum_{i=1}^{n \wedge \tau(\epsilon)} \epsilon_i \mathbf{x}_i(\epsilon) \right) \right\|_{\mathcal{H}^*} \right]$$

Applying conclusion of the previous lemma we get that

$$\begin{aligned} \mathbb{P} \left( T > 0, \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^{n \wedge \tau(\epsilon)} \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} > c \right) &\leq \frac{\alpha_p}{c} \sup_{\epsilon} \left( \mathbf{1}_{\{\tau(\epsilon) > 0\}} \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i=1}^{\tau(\epsilon)} \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p \right) \right)^{1/p} \\ &\leq \frac{\alpha_p}{c} (a^p)^{1/p} = \frac{\alpha_p a}{c} \end{aligned}$$

Plugging the above and Equation 7.6.4 into Equation 7.6.3 we conclude that:

$$\mathbb{P} \left( \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} > c \right) \leq \frac{\|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p]}{a^p} + \frac{\alpha_p a}{c}$$

Using  $a = \left( \frac{c}{\alpha_p} \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \right) \right)^{1/(p+1)}$  we conclude that

$$\mathbb{P} \left( \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} > c \right) \leq 2 \left( \frac{\alpha_p}{c} \right)^{p/(p+1)} \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \right)^{1/(p+1)}$$

This conclude the proof.  $\square$

**Lemma 80.** Let  $1 < r \leq 2$ . If there exists a constant  $D > 0$  such that any  $\mathbf{x}_0 \in \mathcal{B}^*$  and any  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth satisfies :

$$\forall n \in \mathbb{N}, \quad \mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} \right] \leq D(n+1)^{1/r} \sup_{0 \leq i \leq n} \sup_{\epsilon} \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}$$

then for any  $p < r$ , any  $\mathbf{x}_0 \in \mathcal{B}^*$  and  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth will satisfy :

$$\begin{aligned} & \sup_{\lambda > 0} \lambda^p \mathbb{P} \left( \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} > \lambda \right) \\ & \leq \max \left\{ 4^{\frac{p+1}{p}} \alpha_p \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \right)^{\frac{1}{p}}, 2^{2p+3} \log(2) \alpha_p^p \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \right) \right\} \end{aligned}$$

*Proof.* We shall use Proposition 8.53 of Pisier's notes which is restated below to prove this lemma. To this end consider any  $\mathbf{x}_0 \in \mathcal{B}^*$  and any  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth. Given an  $\epsilon \in \{\pm 1\}^{\mathbb{N}}$ , for any  $j \in [M]$  and  $i \in \mathbb{N}$  let  $\epsilon_i^{(j)} = \epsilon_{(i-1)M+j}$ . Let  $z_0 = \mathbf{x}_0 M^{-1/p}$  and define the sequence  $(z_i)_{i \geq 1}$  as follows, for any  $k \in \mathbb{N}$  given by  $k = j + (i-1)M$  where  $j \in [M]$  and  $i \in \mathbb{N}$ ,

$$z_k(\epsilon) = \mathbf{x}_i(\epsilon^{(j)}) M^{-1/p}$$

Clearly,

$$\begin{aligned} \|z_0\|_{\mathcal{X}}^p + \sum_{k \geq 1} \mathbb{E} [\|z_k(\epsilon)\|_{\mathcal{X}}^p] &= \|\mathbf{x}_0\|_{\mathcal{X}}^p + \frac{1}{M} \sum_{j=1}^M \sum_{k \geq 1} \mathbb{E} [\|\mathbf{x}_k(\epsilon^{(j)})\|_{\mathcal{X}}^p] \\ &= \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \end{aligned}$$

By previous lemma we get that for any  $c > 0$ ,

$$\begin{aligned} \mathbb{P} \left( \sup_n \left\| z_0 + \sum_{i=1}^n \epsilon_i z_i(\epsilon) \right\|_{\mathcal{H}^*} > c \right) &\leq 2 \left( \frac{\alpha_p}{c} \right)^{p/(p+1)} \left( \|z_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|z_i(\epsilon)\|_{\mathcal{X}}^p] \right)^{1/(p+1)} \\ &= 2 \left( \frac{\alpha_p}{c} \right)^{p/(p+1)} \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \right)^{1/(p+1)} \end{aligned}$$

Note that

$$\sup_n \left\| z_0 + \sum_{i=1}^n \epsilon_i z_i(\epsilon) \right\|_{\mathcal{H}^*} = M^{-1/p} \sup_{j \in [M]} \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i^{(j)} \mathbf{x}_i(\epsilon^{(j)}) \right\|_{\mathcal{H}^*}$$

Hence we conclude that

$$\mathbb{P} \left( \sup_{j \in [M]} M^{-1/p} \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i^{(j)} \mathbf{x}_i(\epsilon^{(j)}) \right\|_{\mathcal{H}^*} > c \right) \leq 2 \left( \frac{\alpha_p}{c} \right)^{\frac{p}{(p+1)}} \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \right)^{\frac{1}{(p+1)}}$$

For any  $j \in [M]$ , defining  $Z^{(j)} = \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i^{(j)} \mathbf{x}_i(\epsilon^{(j)}) \right\|_{\mathcal{H}^*}$  and using Proposition 82 we conclude that for any  $c > 0$ ,

$$\begin{aligned} & \sup_{\lambda > 0} \lambda^p \mathbb{P} \left( \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} > \lambda \right) \\ & \leq \max \left\{ c, 2c^p \log \left( \frac{1}{1 - 2 \left( \frac{\alpha_p}{c} \right)^{\frac{p}{(p+1)}} \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \right)^{\frac{1}{(p+1)}}} \right) \right\} \end{aligned}$$

Picking

$$c = 4^{\frac{p+1}{p}} \alpha_p \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \right)^{1/p}$$

we conclude that

$$\begin{aligned} & \sup_{\lambda > 0} \lambda^p \mathbb{P} \left( \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} > \lambda \right) \\ & \leq \max \left\{ 4^{\frac{p+1}{p}} \alpha_p \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \right)^{\frac{1}{p}}, 2^{2p+3} \log(2) \alpha_p^p \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \right) \right\} \end{aligned}$$

□

**Lemma 81.** *Let  $1 < r \leq 2$ . If there exists a constant  $D > 0$  such that any  $\mathbf{x}_0 \in \mathcal{B}^*$  and any  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth satisfies :*

$$\forall n \in \mathbb{N}, \quad \mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} \right] \leq D(n+1)^{1/r} \sup_{0 \leq i \leq n} \sup_{\epsilon} \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}$$

*then for all  $p < r$ , we can conclude that any  $\mathbf{x}_0 \in \mathcal{B}^*$  and any  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth will satisfy :*

$$\sup_n \mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] \leq \left( \frac{1104 D}{(r-p)^2} \right)^p \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \right)$$

*That is the pair  $(\mathcal{H}, \mathcal{X})$  is of martingale type  $p$ .*

*Proof.* Given any  $p < r$  pick  $r > p' > p$ , due to the homogeneity of the statement we need to prove, w.l.o.g.

we can assume that

$$\|\mathbf{x}_0\|_{\mathcal{X}}^{p'} + \sum_{i \geq 1} \mathbb{E} \left[ \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^{p'} \right] = 1$$

Hence by previous lemma, we can conclude that

$$\sup_{\lambda > 0} \lambda^{p'} \mathbb{P} \left( \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} > \lambda \right) \leq p' 2^{2p'+3} \log(2) \alpha_{p'}^{p'} \leq (32 \alpha_{p'})^{p'} \quad (7.6.5)$$

Hence,

$$\begin{aligned} \mathbb{E} \left[ \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] &\leq \inf_{a > 0} \left\{ a^{p'} + p \int_a^\infty \lambda^{p-1} \mathbb{P} \left( \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} > \lambda \right) d\lambda \right\} \\ &\leq \inf_{a > 0} \left\{ a^p + p(32 \alpha_{p'})^{p'} \int_a^\infty \lambda^{p-1-p'} d\lambda \right\} \\ &\leq \inf_{a > 0} \left\{ a^p + p(32 \alpha_{p'})^{p'} \left[ \frac{\lambda^{p-p'}}{p-p'} \right]_a^\infty \right\} \\ &\leq \inf_{a > 0} \left\{ a^p + (46 \alpha_{p'})^{p'} \frac{a^{p-p'}}{p-p} \right\} \\ &= 2 \frac{(46 \alpha_{p'})^p}{(p'-p)^{p/p'}} \leq 2 \frac{(46 \alpha_p)^p}{(p'-p)^{p/p'}} \end{aligned}$$

Since  $\|\mathbf{x}_0\|_{\mathcal{X}}^{p'} + \sum_{i \geq 1} \mathbb{E} \left[ \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^{p'} \right] = 1$  and  $p' > p$ , we can conclude that  $\|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} \left[ \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p \right] \geq 1$  and so

$$\begin{aligned} \mathbb{E} \left[ \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] &\leq 2 \frac{(46 \alpha_p)^p}{(p'-p)^{p/p'}} \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} \left[ \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p \right] \right) \\ &\leq 2 \frac{(46 \alpha_p)^p}{(p'-p)} \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} \left[ \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p \right] \right) \end{aligned}$$

Since  $p'$  can be chosen arbitrarily close to  $r$ , taking the limit we can conclude that

$$\mathbb{E} \left[ \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] \leq 2 \frac{(46 \alpha_p)^p}{(r-p)} \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} \left[ \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p \right] \right)$$



Recalling that  $\alpha_p = \frac{12D}{r-p}$  we conclude that

$$\begin{aligned} \mathbb{E} \left[ \sup_n \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] &\leq \left( \frac{1104 D}{(r-p)^{(p+1)/p}} \right)^p \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \right) \\ &\leq \left( \frac{1104 D}{(r-p)^2} \right)^p \left( \|\mathbf{x}_0\|_{\mathcal{X}}^p + \sum_{i \geq 1} \mathbb{E} [\|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^p] \right) \end{aligned}$$

This concludes the proof.  $\square$

We restate below a proposition from Pisier's note (in [66])

**Proposition 82** (Proposition 8.53 of [66]). *Consider a random variable  $Z \geq 0$  and a sequence  $Z^{(1)}, Z^{(2)}, \dots$  drawn iid from some distribution. For some  $0 < p < \infty$ ,  $0 < \delta < 1$  and  $R > 0$ ,*

$$\sup_{M \geq 1} \mathbb{P} \left( \sup_{m \leq M} M^{-1/p} Z^{(m)} > R \right) \leq \delta \implies \sup_{\lambda > 0} \lambda^p \mathbb{P}(Z > \lambda) \leq \max \left\{ R, 2R^p \log \left( \frac{1}{1-\delta} \right) \right\}$$

**Proof of Lemma 67.** By Theorem 66 and our assumption that  $\mathcal{V}_n(\mathcal{H}, \mathcal{X}) \leq Dn^{-(1-1/r)}$ , we have that for any  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth and any  $n \geq 1$ ,

$$\mathbb{E} \left[ \frac{1}{n} \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} \right] \leq Dn^{-(1-\frac{1}{r})}$$

Hence we can conclude for any  $\mathcal{B}^*$ -valued tree  $\mathbf{x}$  of infinite depth and any  $n \geq 1$ ,

$$\mathbb{E} \left[ \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} \right] \leq Dn^{\frac{1}{r}} \sup_{1 \leq i \leq n} \sup_{\epsilon} \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}$$

Hence for any  $\mathbf{x}_0 \in \mathcal{B}^*$ , we have that

$$\begin{aligned} \mathbb{E} \left[ \left\| \mathbf{x}_0 + \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} \right] &\leq \mathbb{E} \left[ \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} \right] + \|\mathbf{x}_0\|_{\mathcal{H}^*} \\ &\leq \mathbb{E} \left[ \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} \right] + D \|\mathbf{x}_0\|_{\mathcal{X}} \\ &\leq Dn^{\frac{1}{r}} \sup_{1 \leq i \leq n} \sup_{\epsilon} \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}} + D \|\mathbf{x}_0\|_{\mathcal{X}} \\ &\leq 2D(n+1)^{\frac{1}{r}} \sup_{0 \leq i \leq n} \sup_{\epsilon} \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}} \end{aligned}$$

Now applying Lemma 81 with  $s = p$  completes the proof.  $\square$

*Proof of Lemma 72.* We now show that the bound got by the Mirror Descent algorithm is tight. First assume without loss of generality that  $\sup_{\mathbf{x} \in \mathcal{X}, \mathbf{h} \in \mathcal{H}} \langle \mathbf{h}, \mathbf{x} \rangle = 1$ . Consider the 1-smooth convex loss function.

$$\phi(z, y) = \begin{cases} |z - y| - \frac{1}{4} & \text{if } |z - y| > \frac{1}{2} \\ (z - y)^2 & \text{otherwise} \end{cases}$$

This is basically a smoothed version of the absolute loss. Consider the smooth convex objective

$$\ell(\mathbf{h}, (\mathbf{x}, y)) = \phi(\langle \mathbf{h}, \mathbf{x} \rangle, y)$$

given by instances  $\mathbf{x} \in \mathcal{X}$  and  $y \in [-1, 1]$ . Now before we proceed we recall from [5] that the value of the online learning game is equal to :

$$\begin{aligned} & \mathcal{V}_n(\mathcal{H}, \mathcal{X}) \\ &= \sup_{p_1} \mathbb{E}_{(\mathbf{x}_1, y_1) \sim p_1} \dots \sup_{p_n} \mathbb{E}_{(\mathbf{x}_n, y_n) \sim p_n} \left[ \frac{1}{n} \sum_{t=1}^n \inf_{\mathbf{h}_t \in \mathcal{H}} \mathbb{E}_{(\mathbf{x}_t, y_t) \sim p_t} [\phi(\langle \mathbf{h}_t, \mathbf{x}_t \rangle, y_t)] - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \phi(\langle \mathbf{h}, \mathbf{x}_t \rangle, y_t) \right] \\ &\geq \mathbb{E}_{(\mathbf{x}_1, y_1) \sim p_1^*} \dots \mathbb{E}_{(\mathbf{x}_n, y_n) \sim p_n^*} \left[ \frac{1}{n} \sum_{t=1}^n \inf_{\mathbf{h}_t \in \mathcal{H}} \mathbb{E}_{(\mathbf{x}_t, y_t) \sim p_t^*} [\phi(\langle \mathbf{h}_t, \mathbf{x}_t \rangle, y_t)] - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \phi(\langle \mathbf{h}, \mathbf{x}_t \rangle, y_t) \right] \quad (7.6.6) \end{aligned}$$

Where  $p_1^*, \dots, p_n^*$  is the distribution on  $\mathcal{X} \times [-1, 1]$  described as follows :

- For the first  $n - m$  rounds,  $p_t^*$  deterministically sets  $\mathbf{x}_t = 0$  and  $y_t = 0$ .
- For the remaining  $m$  rounds, that is for  $t > n - m$ , we first consider a  $\mathcal{B}^*$ -valued tree  $\mathbf{u}$  of infinite depth. The distribution  $p_t^*$  picks  $y_t = \epsilon_{t-n+m}$  where each  $\epsilon_i \sim \text{Unif}\{\pm 1\}$  are Rademacher random variables.  $\mathbf{x}_t$  is chosen by setting  $\mathbf{x}_t = \mathbf{u}_{t-n-m}(\epsilon_1, \dots, \epsilon_{t-n-m-1})$ .

Now notice that on the first  $n - m$  round any algorithm suffers no regret. Using this particular distribution in the lower bound in Equation 7.6.6 we get that :

$$\mathcal{V}_n(\mathcal{H}, \mathcal{X}) \geq \frac{1}{n} \mathbb{E}_{(\mathbf{x}_1, y_1) \sim p_1^*} \dots \mathbb{E}_{(\mathbf{x}_n, y_n) \sim p_n^*} \left[ \sum_{t=n-m+1}^n \inf_{\mathbf{h}_t \in \mathcal{H}} \mathbb{E}_{(\mathbf{x}_t, y_t) \sim p_t^*} [\phi(\langle \mathbf{h}_t, \mathbf{x}_t \rangle, y_t)] - \inf_{\mathbf{h} \in \mathcal{H}} \sum_{t=n-m+1}^n \phi(\langle \mathbf{h}, \mathbf{x}_t \rangle, y_t) \right]$$

Now we shall select the tree  $\mathbf{u}$  such that for any  $n$  and any  $\epsilon \in \{\pm 1\}^{n-1}$ ,  $\|\mathbf{u}_n(\epsilon)\|_{\mathcal{X}} \leq \frac{1}{2}$ . Hence we see that for any choice of  $\mathbf{h} \in \mathcal{H}$ , and any  $t > n - m$ ,  $|\langle \mathbf{h}, \mathbf{x}_t \rangle| \leq \frac{1}{2}$ . Hence we can conclude that for any  $t < n - m$ , and any  $\mathbf{h} \in \mathcal{H}$ ,

$$\phi(\langle \mathbf{h}, \mathbf{x}_t \rangle, y_t) = |\langle \mathbf{h}, \mathbf{x}_t \rangle - y_t| - \frac{1}{4} = \frac{3}{4} - y_t \langle \mathbf{h}, \mathbf{x}_t \rangle$$

Further notice that for any  $t > n - m$ , since  $y_t \sim \text{Unif}\{\pm 1\}$  we have that,

$$\mathbb{E}_{(\mathbf{x}_t, y_t)} [\phi(\langle \mathbf{h}, \mathbf{x}_t \rangle, y_t)] = \frac{3}{4}$$

Plugging this in the lower bound lower bound to the value, we get that

$$\begin{aligned}
\mathcal{V}_n(\mathcal{H}, \mathcal{X}) &\geq \frac{1}{n} \mathbb{E}_{(\mathbf{x}_1, y_1) \sim p_1^*} \dots \mathbb{E}_{(\mathbf{x}_n, y_n) \sim p_n^*} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \sum_{t=n-m+1}^n \left( \frac{3}{4} - \phi(\langle \mathbf{h}, \mathbf{x}_t \rangle, y_t) \right) \right] \\
&= \frac{1}{n} \mathbb{E}_{(\mathbf{x}_1, y_1) \sim p_1^*} \dots \mathbb{E}_{(\mathbf{x}_n, y_n) \sim p_n^*} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \sum_{t=n-m+1}^n \left( \frac{3}{4} - \left( \frac{3}{4} - y_t \langle \mathbf{h}, \mathbf{x}_t \rangle \right) \right) \right] \\
&= \frac{1}{n} \mathbb{E}_{(\mathbf{x}_1, y_1) \sim p_1^*} \dots \mathbb{E}_{(\mathbf{x}_n, y_n) \sim p_n^*} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \sum_{t=n-m+1}^n (\epsilon_{t-n+m} \langle \mathbf{h}, \mathbf{x}_t \rangle) \right] \\
&= \frac{1}{n} \mathbb{E}_\epsilon \left[ \sup_{\mathbf{h} \in \mathcal{H}} \left\langle \mathbf{h}, \sum_{i=1}^m \epsilon_i \mathbf{u}_i(\epsilon) \right\rangle \right] \\
&= \frac{1}{n} \mathbb{E}_\epsilon \left[ \left\| \sum_{i=1}^m \epsilon_i \mathbf{u}_i(\epsilon) \right\|_{\mathcal{H}} \right]
\end{aligned}$$

Specifically picking  $m = \overline{L}^* n$  concludes the proof.  $\square$

*Proof of Lemma 74.* Consider the functions of form  $\ell(\mathbf{h}, z_t) = \langle \mathbf{x}_t, \mathbf{h} \rangle + \sigma\psi(\mathbf{h})$ . Recall from [5] that the value of the online learning game is equal to :

$$\begin{aligned}
\mathcal{V}(\mathcal{H}, \mathcal{Z}_{\text{ucvx}(\sigma, q')})(\mathcal{X}) &= \sup_{p_1} \mathbb{E}_{\mathbf{x}_1 \sim p_1} \dots \sup_{p_n} \mathbb{E}_{\mathbf{x}_n \sim p_n} \left[ \frac{1}{n} \sum_{t=1}^n \inf_{\mathbf{h}_t \in \mathcal{H}} \mathbb{E}_{\mathbf{x}_t \sim p_t} [\ell(\mathbf{h}_t, z_t)] - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n f_t(\mathbf{h}) \right] \\
&\geq \mathbb{E}_{\mathbf{x}_1 \sim p_1^*} \dots \mathbb{E}_{\mathbf{x}_n \sim p_n^*} \left[ \frac{1}{n} \sum_{t=1}^n \inf_{\mathbf{h}_t \in \mathcal{H}} \mathbb{E}_{\mathbf{x}_t \sim p_t^*} [\ell(\mathbf{h}_t, z_t)] - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \right] \\
&= \mathbb{E}_{\mathbf{x}_1 \sim p_1^*} \dots \mathbb{E}_{\mathbf{x}_n \sim p_n^*} \left[ \frac{1}{n} \sum_{t=1}^n \inf_{\mathbf{h}_t \in \mathcal{H}} \mathbb{E}_{\mathbf{x}_t \sim p_t^*} [\langle \mathbf{h}_t, \mathbf{x}_t \rangle] + \sigma\psi(\mathbf{h}_t) - \inf_{\mathbf{h} \in \mathcal{H}} \left\{ \frac{1}{n} \sum_{t=1}^n \langle \mathbf{h}, \mathbf{x}_t \rangle + \sigma\psi(\mathbf{h}) \right\} \right]
\end{aligned}$$

Where  $p_1^*, \dots, p_n^*$  is a distribution on  $\mathcal{X}$  specified as follows : We consider a  $\mathcal{X}$ -valued tree  $\mathbf{u}$  of infinite depth and the distribution  $p_t^*$  picks  $\mathbf{x}_t = \epsilon_t \mathbf{u}_t(\epsilon)$  where each  $\epsilon_t \sim \text{Unif}\{\pm 1\}$  are Rademacher random variables.

Hence we conclude that

$$\begin{aligned}
\mathcal{V}(\mathcal{H}, \mathcal{Z}_{\text{ucvx}(\sigma, q')}(\mathcal{X})) &\geq \mathbb{E}_\epsilon \left[ \frac{1}{n} \sum_{t=1}^n \inf_{\mathbf{h}_t \in \mathcal{H}} \mathbb{E}_{\epsilon_t} [\epsilon_t \langle \mathbf{h}_t, \epsilon_t \mathbf{u}_t(\epsilon) \rangle] + \sigma \psi(\mathbf{h}_t) - \inf_{\mathbf{h} \in \mathcal{H}} \left\{ \frac{1}{n} \sum_{t=1}^n \langle \mathbf{h}, \epsilon_t \mathbf{u}_t(\epsilon) \rangle + \sigma \psi(\mathbf{h}) \right\} \right] \\
&\geq \mathbb{E}_\epsilon \left[ \sup_{\mathbf{h} \in \mathcal{H}} \left\{ \left\langle \mathbf{h}, -\frac{1}{n} \sum_{t=1}^n \epsilon_t \mathbf{u}_t(\epsilon) \right\rangle - \sigma \psi(\mathbf{h}) \right\} \right] \\
&\geq \mathbb{E}_\epsilon \left[ \sup_{\mathbf{h} \in \mathcal{H}} \left\{ \left\langle \mathbf{h}, -\frac{1}{n} \sum_{t=1}^n \epsilon_t \mathbf{u}_t(\epsilon) \right\rangle - \frac{D_{p'}^{q'} \sigma}{q'} \|\mathbf{h}\|_{\mathcal{H}}^{q'} \right\} \right] \\
&= \frac{1}{p' \sigma^{p'-1} D_{p'}^{p'}} \mathbb{E}_\epsilon \left[ \left\| \frac{1}{n} \sum_{t=1}^n \epsilon_t \mathbf{u}_t(\epsilon) \right\|_{\mathcal{H}^*}^{p'} \right] \\
&\geq \frac{1}{p' \sigma^{p'-1} D_{p'}^{p'}} \left( \mathbb{E}_\epsilon \left[ \left\| \frac{1}{n} \sum_{t=1}^n \epsilon_t \mathbf{u}_t(\epsilon) \right\|_{\mathcal{H}^*} \right] \right)^{p'}
\end{aligned}$$

Since choice of  $\mathcal{X}$ -valued the tree  $\mathbf{u}$  is arbitrary we can take supremum over all such trees which concludes the proof.  $\square$

## 7.7 Discussion

We first note that using mirror descent with uniformly convex function (as opposed to strongly convex) is not new and has been used in optimization setting in [1]. The key result of this chapter is to establish universality and near optimality of mirror descent for online learning problems. As we showed it is even optimal for smooth learning problems and some uniformly convex learning problems. While the classic definition of martingale type and the associated results are for dual pairs, in this chapter we extended results by [65] to handle non-dual scenario. Also the proofs have been slightly modified to obtain right dependence on the constants since these constants could be dimension dependent and hence its important to keep track of them.

## Chapter 8

# Optimality of Mirror Descent for Statistical Convex Learning Problems

In the previous chapter we considered convex learning problems in the online learning framework and showed near optimality of mirror descent algorithm for learning rates for online convex learning problems. In this chapter we consider these convex learning problems in statistical learning framework. We show that in fact for most commonly occurring convex learning problems, stochastic mirror descent is in fact near optimal in terms of both sample complexity and efficiency.

To mirror the notation and analysis in previous chapter, given any target hypothesis  $\mathcal{H}$  and any instance class  $\mathcal{Z}$  let us define

$$\mathcal{V}_n^{\text{iid}}(\mathcal{H}, \mathcal{Z}) := \inf_{\mathbf{A}} \sup_{\mathcal{D} \in \Delta(\mathcal{Z})} \mathbb{E}_{S \sim \mathcal{D}^n} \left[ L_{\mathcal{D}}(\mathbf{A}(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right].$$

However unlike the online learning setting for the infimum over the learning algorithm above we shall only consider proper learning algorithms (and as mentioned earlier it suffices to only consider deterministic algorithms), that is  $\bar{\mathcal{H}} = \mathcal{H}$ . Notice that  $\mathcal{V}_n^{\text{iid}}(\mathcal{H}, \mathcal{Z})$  is essentially the same as term  $\epsilon_{\text{cons}}(n)$  we introduced in chapter 3. We introduce this notation to mirror the results in previous chapter and because in this chapter the sets  $\mathcal{H}$  and  $\mathcal{Z}$  we are referring to is rather important and so we prefer to explicitly show this. Also as in the previous chapter the linear learning problem is central to obtaining many of the lower bounds and throughout this chapter we shall use the notation

$$\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X}) = \{\mathbf{x} \mapsto \langle \mathbf{h}, \mathbf{x} \rangle : \mathbf{h} \in \mathcal{H}\}$$

to represent the linear loss class  $\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X}) \subset \mathbb{R}^{\mathcal{X}}$ .

In Section 8.1 we provide lower bounds on learning rates for various convex learning problems (including for smooth losses) using the statistical Rademacher complexity of the linear class  $\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})$ . While all the lower bounds are provided based on statistical Rademacher complexity of the linear function class, in

Section 8.2, analogous to previous chapter we show that the concept of Rademacher type and the Rademacher Complexity of the linear function class are closely related. In chapter 6 we described the stochastic mirror descent algorithm and provided guarantees for it for various statistical convex learning problems. As we saw in the previous chapter these rates are characterized by Martingale type. However the lower bounds in Section 8.2 are in terms of statistical Rademacher complexity or equivalently characterized by Rademacher type. In general, it turns out that Rademacher type and Martingale types do not match and so in general the upper bound on performance of stochastic mirror descent do not match lower bound. However in Section 8.3 we show that for most reasonable cases, the concepts of Rademacher type and Martingale types do in fact match and so we use this result to argue that stochastic mirror descent is near optimal even for statistical learning problems. In the same section we also show that stochastic mirror descent algorithm is optimal also in terms of number of gradient access and for certain supervised learning problems, optimal in terms of computational complexity. In the succeeding section 8.4 we review the examples we saw in the previous chapter to see how mirror descent is indeed optimal. Section 8.5 provides detailed proof of the results in this chapter and we finally conclude the chapter with some discussions in Section 8.6.

## 8.1 Lower Bounds for Statistical Learning Rates

We would like to point out here that linear instance class  $\mathcal{Z}_{\text{lin}}$  is a subclass of the Lipschitz class  $\mathcal{Z}_{\text{Lip}}$ . Hence any lower bound on learning rates for linear class is also a lower bound for the convex Lipschitz class. Similarly, if we consider the supervised learning class with label  $y$  always being  $b = -\sup_{\mathbf{h} \in \mathcal{H}, \mathbf{x} \in \mathcal{X}} \langle \mathbf{h}, \mathbf{x} \rangle$  then we see that  $\ell(\mathbf{h}; (\mathbf{x}, y)) = \langle \mathbf{h}, \mathbf{x} \rangle - b$ . Hence we see that in excess risk for any distribution  $\mathcal{D}$  on  $\mathbf{x}$ 's with  $y$  being deterministically set to  $b$  is same as excess risk of linear class with distribution  $\mathcal{D}$  on  $\mathbf{x}$ 's. Hence we see that lower bound on linear class is also a lower bound on supervised learning class. The following proposition formalizes this.

**Proposition 83.** *For any hypothesis set  $\mathcal{H} \subset \mathcal{B}$ , any  $\mathcal{X} \subset \mathcal{B}^*$  and  $\epsilon > 0$*

$$\mathcal{V}_n^{\text{iid}}(\mathcal{H}, \mathcal{Z}_{\text{Lip}}(\mathcal{X})) \geq \mathcal{V}_n^{\text{iid}}(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X})) \quad \text{and} \quad \mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{sup}}(\mathcal{X})) \geq \mathcal{V}_n(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X})) .$$

Owing to the fact that optimal learning rates for the linear class provide lower bounds for convex lipschitz and supervised convex learning problems, we define for each  $p \in [1, 2]$  the constant  $V_p^{\text{iid}}$  analogous to the definition of  $V_p$  in the previous chapter.

$$V_p^{\text{iid}} := \inf \left\{ V \mid \forall n \in \mathbb{N}, \mathcal{V}_n^{\text{iid}}(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X})) \leq V n^{-(1-\frac{1}{p})} \right\} \quad (8.1.1)$$

In the previous chapter we saw that the value of the linear game was closely related to the sequential Rademacher complexity of the linear class,  $\mathfrak{R}_n^{\text{seq}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X}))$ . We will see an analogous relationship in the statistical setting with the worst case statistical Rademacher complexity. Recall that for the linear class

specified by sets  $\mathcal{H} \subset \mathcal{B}$  and  $\mathcal{X} \subset \mathcal{B}^*$ , the worst case statistical Rademacher complexity is given by

$$\mathcal{R}_n^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) = \sup_{\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathcal{X}} \mathbb{E}_\epsilon \left[ \sup_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n \epsilon_i \langle \mathbf{h}, \mathbf{x}_i \rangle \right] = \sup_{\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathcal{X}} \mathbb{E}_\epsilon \left[ \left\| \frac{1}{n} \sum_{i=1}^n \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right]$$

where  $\epsilon \in \{\pm 1\}^n$  are iid Rademacher random variables. In the following lemma we show that the learning rate of the linear class can in turn be bounded by the worst case statistical Rademacher complexity of the linear class.

**Lemma 84.** *For any  $\mathcal{H} \subset \mathcal{B}$  and  $\mathcal{X} \subset \mathcal{B}^*$ ,*

$$\mathcal{V}_n^{\text{iid}}(\mathcal{H}, \mathcal{Z}_{\text{lin}}(\mathcal{X})) \geq \mathcal{R}_{2n}^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) - \frac{1}{2} \mathcal{R}_n^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X}))$$

where  $n = |S|$ .

While the above lemma along Proposition 83 lower bounds learning rates for convex lipschitz, supervised and linear problems by Rademacher complexity of linear class, for linear and supervised learning problems we can also show that Rademacher complexity of linear class can be used to upper bound the optimal learning rates. The following proposition which is a direct consequence of results from [38] show that the learning rates for linear and supervised convex learning problems are upper bounded by the statistical Rademacher complexity.

**Proposition 85.** [38] *For any set  $\mathcal{X} \in \mathcal{B}^*$  if  $\mathcal{Z}(\mathcal{X})$  is one of either  $\mathcal{Z}_{\text{lin}}(\mathcal{X})$  or  $\mathcal{Z}_{\text{sup}}(\mathcal{X})$ , then for any  $n \in \mathbb{N}$ ,*

$$\sup_{\mathcal{D} \in \Delta(\mathcal{Z}(\mathcal{X}))} \mathbb{E}_{S \sim \mathcal{D}^n} \left[ L_{\mathcal{D}}(\mathbf{A}_{\text{ERM}}(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] \leq 2 \mathcal{R}_n^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{X}))$$

and hence :

$$\mathcal{V}_n^{\text{iid}}(\mathcal{H}, \mathcal{Z}(\mathcal{X})) \leq 2 \mathcal{R}_n^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) .$$

*Proof.* The inequality of linear class is a direct consequence of symmetrization (see [38]). The inequality for the supervised learning class  $\mathcal{Z}_{\text{sup}}(\mathcal{X})$  additionally uses the Lipschitz contraction property (Theorem 10 (4) of [38] ) since the absolute loss is 1-Lipschitz.  $\square$

While the above proposition bounds the learning rates of linear and supervised convex learning classes using Rademacher complexity of the linear class we don't know if such a result is true for the convex Lipschitz class (unlike the online case).

While Lemma 84 provides lower bound on learning rates of linear learning problems in terms of statistical Rademacher complexity of linear class, the RHS there involves the difference of two Rademacher complexity terms. The following theorem shows that learning rates for these convex learning problems provide direct upper bounds on statistical Rademacher complexity of the linear class that captures all polynomial dependences right.

**Theorem 86.** Given any target hypothesis set  $\mathcal{H} \subset \mathcal{B}$  and instance space  $\mathcal{X} \subset \mathcal{B}^*$ , let  $\mathcal{Z}(\mathcal{X})$  be one of  $\mathcal{Z}_{\text{Lip}}(\mathcal{X})$ ,  $\mathcal{Z}_{\text{sup}}(\mathcal{X})$  or  $\mathcal{Z}_{\text{lin}}(\mathcal{X})$ . If for some  $q \in [2, \infty)$  and  $V > 0$ ,

$$\mathcal{V}_n^{\text{iid}}(\mathcal{H}, \mathcal{Z}(\mathcal{X})) \leq \frac{V}{n^{1/q}}$$

then,

$$\forall n \in \mathbb{N}, \quad \mathcal{R}_{2n}^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) \leq \frac{5V}{(2n)^{\frac{1}{q}}} .$$

### 8.1.1 Lower Bounds for Smooth Losses

The following lemma lower bounds learning rate for non-negative smooth convex learning problems and also captures dependence on expected loss of the optimal hypothesis in target class  $\mathcal{H}$ . This result is an analog to Lemma 72 of previous chapter.

**Lemma 87.** Given  $\mathcal{H} \subset \mathcal{B}$ ,  $\mathcal{X} \subset \mathcal{B}^*$  and  $L^* \in (0, 3/4]$  for any learning algorithm  $\mathbf{A}$ , there exists a distribution  $\mathcal{D}$  over instances in  $\mathcal{Z}_{\text{smt}(1)}(\mathcal{X})$  s.t.  $\inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \leq L^*$  and

$$\mathbb{E}_S \left[ L_{\mathcal{D}}(\mathbf{A}(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] \geq \frac{L^*}{4} \left( \mathcal{R}_{nL^*}^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) - \frac{1}{2} \mathcal{R}_{nL^*}^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) \right)$$

## 8.2 Optimal Rates and Rademacher Type

In this section we extend the classic notion of Rademacher type of a Banach space (see for instance [74]) to one that accounts for the pair  $(\mathcal{H}^*, \mathcal{X})$ . The results of this section are analogous to the results in Section 7.2.

**Definition 33.** A pair  $(\mathcal{H}^*, \mathcal{X})$  of subsets of a vector space  $\mathcal{B}^*$  is said to be of Rademacher type  $p$  if there exists a constant  $C \geq 1$  such that for any  $n \in \mathbb{N}$  and any  $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n \in \mathcal{B}^*$  :

$$\mathbb{E} \left[ \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*}^p \right] \leq C^p \left( \sum_{i=1}^n \|\mathbf{x}_i\|_{\mathcal{X}}^p \right) \quad (8.2.1)$$

It can be shown that rate of convergence of i.i.d. random variables in Banach spaces is governed by the above notion of Rademacher type of the associated Banach space. We point the reader to [66] for more details. Further, for any  $p \in [1, 2]$  we also define constant  $C_p^{\text{iid}}$ , analogous to the definition of  $C_p$  in previous chapter.

$$C_p^{\text{iid}} := \inf \left\{ C \mid \forall n \in \mathbb{N}, \forall \mathbf{x}_1, \dots, \mathbf{x}_n \in \mathcal{B}^*, \mathbb{E} \left[ \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*}^p \right] \leq C^p \left( \sum_{i=1}^n \|\mathbf{x}_i\|_{\mathcal{X}}^p \right) \right\}$$

$C_p^{\text{iid}}$  is useful in determining if the pair  $(\mathcal{H}^*, \mathcal{X})$  has Rademacher type  $p$ .

The following lemma is an analog to Lemma 67 of previous section.



**Lemma 88.** *If for some  $r \in (1, 2]$  there exists a constant  $D > 0$  such that for any  $n$ ,*

$$\mathcal{R}_n^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) = \sup_{\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathcal{X}} \mathbb{E} \left[ \left\| \frac{1}{n} \sum_{i=1}^n \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] \leq D n^{-(1-\frac{1}{r})}$$

*then for all  $p < r$ , we can conclude that for any  $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathcal{B}^*$  :*

$$\mathbb{E} \left[ \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*}^p \right] \leq \left( \frac{12\sqrt{2} D}{(r-p)} \right)^p \left( \sum_{i=1}^n \|\mathbf{x}_i\|_{\mathcal{X}}^p \right)$$

*That is, the pair  $(\mathcal{H}^*, \mathcal{X})$  is of Rademacher type  $p$ .*

**Corollary 89.** *Given any pair  $(\mathcal{H}^*, \mathcal{X})$  and any  $p$  and  $p' < p$ , we have that*

$$\frac{(p-p')}{60\sqrt{2}} C_{p'}^{\text{iid}} \leq V_p^{\text{iid}} \leq 2C_p^{\text{iid}}$$

*Proof.* Owing to Proposition 85 and definition of  $V_p^{\text{iid}}$  using Jensen's inequality, we get that  $V_p^{\text{iid}} \leq 2C_p^{\text{iid}}$ . The second inequality is a consequence of using the above Lemma 88 with Theorem 86 and the definitions of  $C_p^{\text{iid}}$  and  $V_p^{\text{iid}}$ .  $\square$

The following figure summarizes the relationship between  $V_p^{\text{iid}}$  and  $C_p^{\text{iid}}$ . The arrow mark from  $C_{p'}^{\text{iid}}$  to  $C_p^{\text{iid}}$  indicates that for any  $n$ , the quantities are within  $\log n$  factor of each other.




Figure 8.1: Relationship between the various constants in statistical setting

### 8.3 Main Result : Optimality of Stochastic Mirror Descent

In Section 8.1 we provided lower bounds on learning rates of convex learning problems in the statistical learning framework. In previous chapter we provided upper bounds on learning rate of stochastic mirror descent algorithm for statistical convex learning problems. A natural question that arises is whether the stochastic mirror descent algorithm is optimal in terms of learning rates or efficiency or both and whether they match the lower bounds obtained.

It turns out that in general even for convex learning problems, this is not true. A problem could be statistically learnable but not online learnable. However in this section we will see that for large class of convex learning problems (in fact most of the commonly occurring ones) it is true that online methods are near optimal for statistical convex learning problems in terms of both rates and efficiency. In the previous we saw that mirror descent algorithm was near optimal and the key to showing this result is depicted in Figure 7.1. In this

chapter again in Section 8.2 we obtained an analogous result depicted in Figure 8.1. The attractive feature in the online learning scenario was that optimal learning rates and hence concept of martingale type was closely connected to existence of appropriate uniformly convex function which then we could use with mirror descent to get near optimal guarantees. In the statistical case we only have the connection between optimal rates (or at least lower bounds) and the concept of Rademacher type. By the definitions of  $C_p$  and  $C_p^{\text{iid}}$ , it is easy to see that  $C_p^{\text{iid}} \leq C_p$  and due to online to batch conversion we can also infer that  $V_p^{\text{iid}} \leq V_p$ . However assume that we could get a reverse bound, this would then imply that online learning methods are optimal at least up to logarithmic factors. Such optimality of learning rate is essentially captured in the following theorem.

**Theorem 90.** *Let  $\mathcal{Z}(\mathcal{X})$  stand for one of  $\mathcal{Z}_{\text{lin}}(\mathcal{X})$ ,  $\mathcal{Z}_{\text{sup}}(\mathcal{X})$  or  $\mathcal{Z}_{\text{Lip}}(\mathcal{X})$ . If there exists some constant  $G \geq 1$  such that for any  $1 < p' < r \leq 2$*

$$C_{p'} \leq G C_r^{\text{iid}}$$

*and it is true that for some  $V > 0$  and  $p \in (1, \infty)$ ,*

$$\forall n \in \mathbb{N}, \mathcal{V}_n^{\text{iid}}(\mathcal{H}, \mathcal{Z}(\mathcal{X})) \leq \frac{V}{n^{1/q}}$$

*then there exists function  $\Psi$  and step size  $\eta$  using which the stochastic mirror descent algorithm enjoys the learning guarantee*

$$\forall n \in \mathbb{N}, \sup_{\mathcal{D} \in \Delta(\mathcal{Z}(\mathcal{X}))} \mathbb{E}_{S \sim \mathcal{D}^n} \left[ L_{\mathcal{D}}(\overline{\mathcal{A}_{\text{MD}}}(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] \leq \frac{30010GV \log^3 n}{n^{1/q}}$$

*Proof.* Using Theorem 86 we first get a bound on  $\mathcal{R}_n^{\text{iid}}(\mathcal{F}_{\text{lin}})$ . Next using Corollary 89 we see that for any  $p' < p$ ,  $C_{p'}^{\text{iid}} \leq \frac{60\sqrt{2}V}{p-p'}$ . Using the assumption that  $C_p \leq G C_p^{\text{iid}}$  we get bound on  $C_{p'}$  and from here on using exactly the same proof as that of theorem 71 we get the bound on the regret. To convert it to bound on excess risk we can use Proposition 58.  $\square$

Even for the smooth loss case one can show that if for any  $1 < p' < r \leq 2$   $C_{p'} \leq G C_r^{\text{iid}}$  then online mirror descent is near optimal as the following theorem shows.

**Theorem 91.** *If there exists some constant  $G \geq 1$  such that for any  $1 < p' < r \leq 2$*

$$C_{p'} \leq G C_r^{\text{iid}}$$

*and it is true that for some  $V > 0$  and  $p \in (1, \infty)$ ,*

$$\forall n \in \mathbb{N}, \mathcal{V}_n^{\text{iid}}(\mathcal{H}, \mathcal{Z}_{\text{smt}(1)}(\mathcal{X})) \leq \frac{V}{n^{1/q}},$$

*then there exists function  $\Psi$  and step size  $\eta$  using which the mirror descent algorithm followed by online to batch conversion technique, that is the algorithm  $\overline{\mathcal{A}_{\text{MD}}}$ , enjoys the guarantee that for any  $L^* \in (0, 3/4]$  and*

any distribution  $\mathcal{D} \in \Delta(\mathcal{Z}_{\text{smt}(H)}(\mathcal{X}))$  s.t.  $\inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \leq L^*$ ,

$$\mathbb{E}_{S \sim \mathcal{D}^n} \left[ L_{\mathcal{D}}(\overline{\mathcal{A}_{\text{MD}}}(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] \leq \frac{240080 G V \sqrt{H L^*} \log^3 n}{n^{1/q}} + \frac{(189800 G V)^2 H \log^6 n}{n^{2/q}}$$

*Proof.* Again the proof is similar to previous two proofs. The main difference is that at the very first step we use Lemma 87 with  $L^* = 3/4$  to bound  $\mathcal{R}_n^{\text{iid}}(\mathcal{F}_{\text{lin}})$ . Also in the step before last (ie. just before the online to batch conversion), instead of using Theorem 71 we instead use Theorem 73. Also after the online to batch step note that

$$\mathbb{E}_{S \sim \mathcal{D}^n} \left[ \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \right] \leq \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \leq L^*$$

□

Also note that the discussion after Theorem 73 in the previous chapter for the online setting also applies here and specifically when  $q = 2$  we also get tightness w.r.t. dependence in  $L^*$

Owing to the above two theorems the main condition we are now looking to satisfy is that there exists some constant  $G \geq 1$  such that for any  $1 < p' < r \leq 2$

$$C_{p'} \leq G C_r^{\text{iid}}$$

**Remark 92.** In some cases we might only be able to prove that there exists some constant  $G \geq 1$  such that for any  $1 < p' < r \leq 2$

$$C_{p'} \leq \frac{G}{r - p'} C_r^{\text{iid}}$$

such an inequality is also fine and basically all the above three theorems still hold with the only modification that the logarithmic factors in the theorems increase by one more power.

In the following subsections we will show that for a large class of spaces such  $G$  exists (or owing to above remark with additional  $(r - p')^{-1}$  factor). From this we can infer that mirror descent is near optimal for convex learning optimization problems in these spaces.

### 8.3.1 Banach Lattices

In this subsection we show that if the Banach space specified by norm  $\|\cdot\|_{\mathcal{H}^*}$  is a Banach lattice, then one can relate martingale type and Rademacher type constants  $C_p$  and  $C_p^{\text{iid}}$  to within constant factor of each other.

**Definition 34** (Banach Lattice [75]). A partially ordered Banach space is called a Banach lattice provided :

1. For any  $\mathbf{x}, \mathbf{x}', \tilde{\mathbf{x}} \in \mathcal{B}^*$ ,  $\mathbf{x} \preceq \mathbf{x}'$  implies that  $\mathbf{x} + \tilde{\mathbf{x}} \preceq \mathbf{x}' + \tilde{\mathbf{x}}$
2.  $a\mathbf{x} \succeq \mathbf{0}$  for every  $\mathbf{x} \succeq \mathbf{0}$  and non-negative scalar  $a$ .

3. For all  $\mathbf{x}, \mathbf{x}' \in \mathcal{B}^*$ , there exists a least upper bound (l.u.b.) represented by  $\mathbf{x} \vee \mathbf{x}'$ .
4. For any  $\mathbf{x}, \mathbf{x}' \in \mathcal{B}^*$  such that  $|\mathbf{x}| \preceq |\mathbf{x}'|$ , we have that  $\|\mathbf{x}\| \leq \|\mathbf{x}'\|$  (where the absolute value  $|\mathbf{x}|$  is defines as  $|\mathbf{x}| = \mathbf{x} \vee (-\mathbf{x})$ ).

Before we proceed we notice that all  $\ell_p$  spaces with partial order given by,  $\mathbf{x} \preceq \mathbf{y}$  if and only if on each co-ordinate  $i$ ,  $\mathbf{x}_i \leq \mathbf{y}_i$  is a Banach lattice. In fact all the examples we saw in previous chapters were Banach lattices under appropriate partial order. In fact one could safely say that most Banach spaces one would encounter in machine learning applications would be Banach lattices. Hence the results in this section are very general from a practical view-point.

Now we introduce the notation that for any  $p \in [1, \infty)$ , we will use  $(\sum_{i=1}^n |\mathbf{x}_i|^p)^{1/p}$  to represent the vector,

$$\left( \sum_{i=1}^n |\mathbf{x}_i|^p \right)^{1/p} := \text{l.u.b.} \left\{ \sum_{i=1}^n a_i \mathbf{x}_i \mid (a_1, \dots, a_n) \in \mathbb{R}, \sum_{i=1}^n |a_i|^q \leq 1 \right\}$$

where  $q = \frac{p}{p-1}$ . Notice that if  $\mathbf{x}_i$ 's were reals this would be simply the  $\ell_p$  norm. Here it is a vector though. The main technology behind proving results about Banach lattices arises from the so called functional calculus over Banach lattices introduced by Krivine [76] (See [75]). The basic idea is a theorem (theorem 1.d.1 in [75]) which roughly states that if we prove any inequality involving continuous degree 1, homogenous equations involving finite number of real valued variables, then the same inequality is true with of course appropriate changes like  $\leq$  replaced by  $\preceq$  and absolute value replaced by the lattice version and so on.

At first glance the statements we would like to prove would involve expectation over Rademacher random variables and trees (of depth  $n$ ). However simple observation that expectation over Rademacher are finite averages and that the tree of depth  $n$  (involving the  $n$  mappings) can be expanded to the  $2^n - 1$  variables involved we see automatically that this general technology can be used to prove results that involve trees and expectations w.r.t.  $\epsilon$ 's too. Detailed proof and associated definitions are provided in Section 8.5. We delineate the main results below.

Before we proceed we first define below the notion of co-type of a Banach space (again extended to the non-dual case, see ?? for details about classical definition).

**Definition 35.** A pair  $(\mathcal{H}^*, \mathcal{X})$  of subsets of a vector space  $\mathcal{B}^*$  is said to be of Rademacher co-type  $q$  if there exists a constant  $C \geq 1$  such that for any  $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n \in \mathcal{B}^*$  :

$$\left( \sum_{i=1}^n \|\mathbf{x}_i\|_{\mathcal{X}}^q \right) \leq C^q \mathbb{E} \left[ \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*}^q \right] \quad (8.3.1)$$

The following lemma shows that if  $\mathcal{H}^*$  formed a Banach lattice and were of some finite co-type  $r$  then for any  $p$ ,  $C_p$  can be bounded by a constant factor times  $C_p^{\text{iid}}$ .

**Lemma 93.** *If  $(\mathcal{B}^*, \|\cdot\|_{\mathcal{H}^*}^*)$  is a Banach lattice and the pair  $(\mathcal{H}^*, \mathcal{H}^*)$  is of co-type  $r$  for some  $r \in [2, \infty)$  with*

some constant  $\tilde{C}_r$ , then for any  $p \in (1, 2]$  we have that

$$C_p \leq \frac{72^3 r^2 \tilde{C}_r}{\sqrt{p-1}} C_p^{\text{iid}}$$

The above Lemma shows that as long as  $\mathcal{H}^*$  has a lattice structure and is of finite co-type,  $G$  is bounded by  $\frac{72^3 r^2 \tilde{C}_r}{\sqrt{p-1}}$  and so for any such case owing to the theorems in the beginning of this section we automatically can conclude optimality of mirror descent.

### Dual Learning Problem :

In Lemma 93 we needed that  $\mathcal{H}^*$  had finite co-type. Though this is a fairly weak condition we still needed to verify that this was satisfied before we used the result. However if we consider the dual learning problem, that is the case when  $\mathcal{X}$  is the dual ball of  $\mathcal{H}$  then it turns out that this condition is automatically satisfied as long as the problem is statistically learnable (if it is not we anyway are not in a position to give any meaningful bounds). This is due to the celebrated result of Maurey and Pisier [77] which (in the dual case) assures that any space with non-trivial type also has finite co-type. The following corollary which uses a result by König and Tzafriri [?] shows the exact relationship between  $C_p$  and  $C_p^{\text{iid}}$  for Banach lattices.

**Corollary 94.** *For the dual learning problem with  $\mathcal{X} = \mathcal{H}^*$ , if  $(\mathcal{B}^*, \|\cdot\|_{\mathcal{H}^*})$  is a Banach lattice, then for any  $p \in (1, 2]$  we have that*

$$C_p \leq \frac{72^3}{\sqrt{p-1}} (3C_p^{\text{iid}})^{2q+1}$$

where  $q = \frac{p}{p-1}$ .

*Proof.* Note that for the dual problem  $\mathcal{X} = \mathcal{H}^*$  and by definition,  $(\mathcal{H}^*, \mathcal{H}^*)$  has Rademacher type  $p$  with constant  $C_p^{\text{iid}}$ . However by Theorem 3 of [?] we have that the pair  $(\mathcal{H}, \mathcal{H})$  is of co-type  $2 + (2C_p^{\text{iid}})^q$  with constant 2 (for co-type constant of 2 refer proof of the theorem). Hence using this in Lemma 93 with  $r = 2 + (2C_p^{\text{iid}})^q$  and  $\tilde{C}_r = 2$  and simplifying yields the required statement.  $\square$

Thus for dual learning problems if the Banach space is a Banach Lattice then online methods (specifically mirror descent) is always near optimal in terms of dependence on  $n$  for learning rate and in terms of dependence on  $\epsilon$  for oracle complexity. More specifically, for a dual learning problem on a Banach lattice, if for instance optimal rate is  $V/\sqrt{n}$  then mirror descent will guarantee a rate of order  $V^5 \log^2(n)/\sqrt{n}$ .

### 8.3.2 Decoupling Inequalities

Another way to guarantee that  $C_p \leq GC_p^{\text{iid}}$  for some finite  $G$  is by using the so called decoupling inequalities (see [78, 79] for more details).

**Definition 36.** We say that a Banach space with norm  $\|\cdot\|_{\mathcal{H}}^*$  satisfies  $p$ -decoupling inequality with constant  $B > 0$  if :

$$\mathbb{E}_{\epsilon} \left[ \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] \leq B^p \mathbb{E}_{\epsilon, \epsilon'} \left[ \left\| \sum_{i=1}^n \epsilon_i \epsilon'_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^p \right] \quad (8.3.2)$$

where  $\epsilon'_1, \dots, \epsilon'_n$  are Rademacher random variables (drawn independent of  $\epsilon_1, \dots, \epsilon_n$ ).

We would like to point out that the above definition is not the same as the decoupling inequalities in [78, 79] where the above needs to be true for all martingale difference sequences where as above we only consider Walsh-Paley martingales. However since above condition is weaker, all positive results for the stronger definition also hold for the above definition. In any case, the below lemma shows that if a Banach space satisfies decoupling inequality, then there exists  $G < \infty$  through which we can upper bound  $C_p$  using  $C_p^{\text{iid}}$ .

**Lemma 95.** *If the Banach space equipped with norm  $\|\cdot\|_{\mathcal{H}^*}$  satisfies 1-decoupling inequality with some constant  $B > 0$ , then it is true that for all  $1 < p' < r \leq 2$ ,*

$$C_{p'} \leq \frac{92B}{(r - p')} C_r^{\text{iid}}$$

First we would like to point out that if a Banach space satisfies  $p$ -decoupling inequality for some  $p \in [1, \infty)$  with constant  $D$ , then it satisfies 1-decoupling inequality with constant  $K_p D$  where  $K_p$  only depends on  $p$  (see Theorem 4.1 of [79]). Hence we can talk of space satisfying decoupling inequality without referring to exponent (of course the exponent plays a role in the constant). Several spaces we commonly encounter satisfy the decoupling inequality. It can be shown that called Unconditional Martingale difference (UMD) Banach spaces always satisfy decoupling inequality. These spaces include all  $L_p, \ell_p$  Schatten  $p$  norms for  $p \in (1, \infty)$ . However while every UMD space satisfies decoupling for instance the  $\ell_1$  space while satisfying the decoupling inequality is not UMD. Using some of the results in [79, 80, 81, 82, 83] we can conclude that for many interesting spaces we commonly encounter, there in fact even exists a universal 1-decoupling constant, call it  $B_{\mathbb{R}}$ . (This constant  $B_{\mathbb{R}}$  is the one referred to as  $D_{\mathbb{R}}$  in [79]). The following proposition proved in [79] (see also [80]) is particularly useful especially to provide decoupling inequalities for group norms and interpolation norms.

**Proposition 96** (Corollaries 4.6 [79]). *If  $Y$  is a Banach space that satisfies  $p$ -decoupling inequality with constant  $B$ , then for any  $\sigma$ -finite measure space  $(S, \Sigma, \mu)$  and any  $p \in [1, \infty)$ , the space  $X = L^p(S; Y)$  satisfies  $p$ -decoupling inequality with same constant  $B$ .*

From the above we immediately see that  $\ell_p$  spaces satisfy  $p$ -decoupling inequality with universal constant  $B_{\mathbb{R}}$  (the space  $Y$  in this case is the reals). We can also use the above result to estimate decoupling constants for group norm by taking the the above Proposition space  $Y$  to be the  $\ell_p$  space corresponding to the inner index of the group norm (mixed norm). In Corollary 4.11 of [79] it has been shown that for  $p \geq \log_2(d)$ , the  $p$ -decoupling constant of  $\ell_{\infty}^d$  spaces is bounded by  $2B_{\mathbb{R}}$ . In [81] it has been shown that a large class of Orlicz and Rearrangement invariant function spaces satisfy 1-decoupling inequality with universal constant

$B_{\mathbb{R}}$  combining these result with the above proposition one can obtain decoupling inequalities for even more examples.

Hence overall, the following figure captures the scenario for the statistical convex learning problems.




Figure 8.2: Relationship between the various constants in both online and statistical settings. Red inequality is not always true but true for most commonly encountered problems.

### 8.3.3 Optimality of Mirror Descent in Terms of Efficiency

In this section we argue that mirror descent is not only optimal in terms of learning rate but also in terms of efficiency. Specifically we are interested in achieving excess risk less than some target sub-optimality  $\epsilon > 0$ , that is, we are interested in coming up with learning algorithm  $\mathbf{A}$  such that

$$\sup_{\mathcal{D} \in \Delta(\mathcal{Z})} \mathbb{E}_{S \sim \mathcal{D}^n} \left[ L_{\mathcal{D}}(\mathbf{A}(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] \leq \epsilon .$$

Now there are two main question one can ask, the first is the minimum sample size required by any algorithm to be able to ensure the above guarantee on sub-optimality. This question is essentially the question about optimal learning rates and we already showed optimality of stochastic mirror descent. The second question one can ask is the computational time required by any algorithm given any amount of requested samples to ensure the above sub-optimality guarantee. Note that in general giving exact computational complexity is cumbersome and even impossible at full generality. A good proxy for computational complexity is the number of gradient calculations or calculation or some kind of other local information like Hessians or higher order derivative. Note that mirror descent is a sub-gradient based method which in each iteration only calculates one gradient. Since we already established that mirror descent is optimal for many statistical convex learning problem, we can easily conclude that mirror descent is also optimal in terms of number of gradient calculations needed to achieve a target sub-optimality.

As mentioned, in general getting a handle on exact computational complexity is hard. However for a few important class of supervised learning problems with absolute loss (can also be extended to hinge loss, logistic loss etc.) one can show that mirror descent is optimal even in terms of exact computational complexity. To see this consider a supervised learning problems where instances given are  $S = (\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)$ . Let  $d$  be the intrinsic dimensionality of the sample  $S$  (in vector space  $\mathcal{B}$ ). Then the total time required just to read the sample is  $dn$ . However note that for many cases like when  $\mathcal{H}/\mathcal{X}$  are various unit balls of  $\ell_p$  norms

or  $L_p$  norms or group norms then as we will see in the next section, mirror descent update step has triple complexity of order  $d$  and so since mirror descent is a single pass algorithm that goes over samples one by one, its time complexity when run once over the sample is again of order  $nd$ . However we already argued that the learning rate of mirror descent is near optimal and so sample size  $n$  required for getting sub-optimality of at most  $\epsilon$  is also near optimal. Hence we can conclude that mirror descent algorithm for these cases have time complexity of same order as time complexity needed just to read the minimum required sample. Hence mirror descent for these cases is near optimal even in terms of exact computational time.

## 8.4 Examples

First of all, we start by noticing that all the examples we considered in the previous chapter were Banach Lattices with finite co-type and so for all these problems by Theorems 90, 104 and 91 stochastic mirror descent technique is near optimal in terms of both rates and number of gradient access, for convex Lipschitz problems, supervised learning problems and non-negative smooth problems. We will further see below that for these examples we don't just get near optimality but that in fact  $C_p$  and  $C_p^{\text{iid}}$  are within fixed numeric constant factor of each other. Hence we see that beyond the logarithmic factor there aren't any space dependent hidden factors even.

### 8.4.1 Example : $\ell_p$ non-dual pairs

In the previous chapter we gave characterization of constant  $D_2$  (corresponding to  $1/\sqrt{n}$  rates in online setting) for the  $\ell_p$  pairs. Recall the setting, that is  $\mathcal{H}$  is the unit ball of  $\ell_{p_1}^d$  ball and  $\mathcal{X}$  is the dual ball of  $\ell_{p_2}^d$  norm. It turns out that for  $\ell_p$  norms, when  $p \in [1, 2]$  we have a universal constant for 1-decoupling by Proposition 96. On the other hand, when  $p \in (2, \infty)$ , the type constant of  $\ell_p$  spaces is of order  $\sqrt{p}$  and so we can conclude that the table is essentially tight even for the statistical learning setting because martingale type constant ( $C_p$ ) and type constant ( $C_p^{\text{iid}}$ ) are within fixed numeric constant factor of each other.

### 8.4.2 Example : Non-dual Schatten norm pairs in finite dimensions

As we showed in the previous chapter, the constants  $D_p$  ( $C_p, V_p$  etc.) for Schatten norms were same as those for  $\ell_p$  norms in the online learning setting. However as we saw in the previous section, constants for  $\ell_p$  norms match in statistical and online frameworks. But since lower bounds for  $\ell_p$  case can be converted to atleast same lower bound on Schatten norm case (by diagonalizing). We can again conclude that rates of Mirror descent in statistical case are near optimal and moreover that  $C_p$  and  $C_p^{\text{iid}}$  are within fixed numeric constant factor of each other.



### 8.4.3 Example : Non-dual group norm pairs in finite dimensions

For the group norm case using Decoupling inequalities for  $\ell_p$  along with Proposition 95 we can show again that  $C_p$  and  $C_p^{\text{iid}}$  are within constant factor of each other which guarantees tightness of mirror descent for for these problems in the statistical setting.

### 8.4.4 Computational Efficiency Issues

Up to now we used number of gradient calculations to argue that mirror descent is optimal in terms of efficiency. Notice that we showed optimality using mirror descent algorithm which has a simple update step that at every round uses only previous hypothesis and gradient of loss at the hypothesis. Hence the time complexity of the update at each round is of the order of effective dimension (as an example in the  $\ell_p$  case it is linear in  $d$ ). Thus once can translate in these cases oracle complexity to time complexity of the algorithm. An for  $\ell_p$  cases one can even show that time complexity is near optimal. However this is not always the case, it might be that the complexity of update step of mirror descent (which depends on  $\Psi$ ) is large that time complexity blows up. This can be seen for instance in the max norm example. In the previous chapter when we considered the max norm example, the function  $\Psi$  we considered in Equation 7.5.1 had a summation over  $2^{M+N}$  elements for matrix of size  $M \times N$ . This is of course prohibitive to use in practice. However notice that while oracle complexity is still tight, for matrix completion problem in the statistical learning framework, one can use a constrained minimization approach (minimize average loss subject to constraint on max norm or equivalently max norm regularization) and using an SDP approach this can be done in time polynomial in the matrix size. This shows a disparity between oracle complexity and actual time complexity. It is an interesting open question whether there is a polynomial time algorithm for max norm based learning problem in the online framework. In [84], in the transductive online setting for max norm it is shown that once can again use the SDP to obtain a poly-time algorithm. An even stronger flavor of questioning is whether one can provide a mirror descent (or variant) algorithm for learning with max-norm that works in time polynomial in matrix size.

## 8.5 Detailed Proofs

*Proof of Lemma 84.* First we pick elements  $\mathbf{x}_1, \dots, \mathbf{x}_{2n} \in \mathcal{X}$  and then draw Rademacher random variables  $\epsilon \in \{\pm 1\}^{2n}$  uniformly at random. We shall now use this to construct a distribution  $\mathcal{D}_\epsilon$  over instances which is the one we shall use for the lower bound. Specifically, given a draw  $\epsilon \in \{\pm 1\}^{2n}$  consider the distribution  $\mathcal{D}_\epsilon$  to be the uniform distribution over the set  $\{\epsilon_1 \mathbf{x}_1, \dots, \epsilon_{2n} \mathbf{x}_{2n}\}$ . Now consider a  $S$  of size  $n$  drawn iid from the distribution  $\mathcal{D}_\epsilon$ . Now note that for any learning algorithm  $\mathbf{A}$ ,

$$\sup_{\mathcal{D}} \mathbb{E}_S \left[ L_{\mathcal{D}}(\mathbf{A}(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] \geq \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2n}} \mathbb{E}_\epsilon \mathbb{E}_{S \sim \mathcal{D}_\epsilon^n} \left[ L_{\mathcal{D}_\epsilon}(\mathbf{A}(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}_\epsilon}(\mathbf{h}) \right]$$

Since  $\mathcal{D}_\epsilon$  is the uniform distribution over set  $\{\epsilon_1 \mathbf{x}_1, \dots, \epsilon_{2n} \mathbf{x}_{2n}\}$  one can rewrite sampling  $S$  from the distribution as follows : First we sample  $n$  numbers,  $t_1, \dots, t_n$  uniformly at random from the set  $[2n]$ . Next, the sample  $S$  is given by  $S = \epsilon_{t_1} \mathbf{x}_{t_1}, \dots, \epsilon_{t_n} \mathbf{x}_{t_n}$ . Hence we can rewrite the above inequality as

$$\sup_{\mathcal{D}} \mathbb{E}_S \left[ L_{\mathcal{D}}(A(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] \geq \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2n}} \mathbb{E}_\epsilon \mathbb{E}_{t_1, \dots, t_n \sim \text{unif}([2n])} \left[ L_{\mathcal{D}_\epsilon}(A(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}_\epsilon}(\mathbf{h}) \right]$$

Now let the set  $J \subset [2n]$  be the set

$$J = \{i \in [2n] : i \in \{t_1, \dots, t_n\}\}$$

that is the set of indices  $i$  such that  $\mathbf{x}_i$  appeared at least once in the sample  $S$  provided to the learner. Also let  $J^c \in [2n]$  stand for the complement of the set  $J$ . For linear class, for the distribution  $\mathcal{D}_\epsilon$  note that for any  $\mathbf{h} \in \mathcal{H}$ ,

$$L_{\mathcal{D}_\epsilon}(\mathbf{h}) = \frac{1}{2n} \sum_{i=1}^{2n} \langle \mathbf{h}, \epsilon_i \mathbf{x}_i \rangle$$

Hence we see that,

$$\begin{aligned} \sup_{\mathcal{D}} \mathbb{E}_S \left[ L_{\mathcal{D}}(A(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] &\geq \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2n} \in \mathcal{X}} \mathbb{E}_\epsilon \mathbb{E}_{t_1, \dots, t_n \sim \text{unif}([2n])} \left[ L_{\mathcal{D}_\epsilon}(A(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}_\epsilon}(\mathbf{h}) \right] \\ &= \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2n} \in \mathcal{X}} \mathbb{E}_\epsilon \mathbb{E}_{t_1, \dots, t_n \sim \text{unif}([2n])} \left[ \sup_{\mathbf{h} \in \mathcal{H}} \frac{1}{2n} \sum_{i=1}^{2n} \langle \mathbf{h}, -\epsilon_i \mathbf{x}_i \rangle - \frac{1}{2n} \sum_{i=1}^{2n} \langle A(S), -\epsilon_i \mathbf{x}_i \rangle \right] \\ &= \frac{1}{2n} \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2n} \in \mathcal{X}} \mathbb{E}_\epsilon \mathbb{E}_{t_1, \dots, t_n \sim \text{unif}([2n])} \left[ \left\| \sum_{i=1}^{2n} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} - \sum_{i \in J} \langle A(S), -\epsilon_i \mathbf{x}_i \rangle - \sum_{i \in J^c} \langle A(S), -\epsilon_i \mathbf{x}_i \rangle \right] \\ &= \frac{1}{2n} \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2n} \in \mathcal{X}} \mathbb{E}_{t_1, \dots, t_n \sim \text{unif}([2n])} \left[ \mathbb{E}_\epsilon \left[ \left\| \sum_{i=1}^{2n} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] - \mathbb{E}_\epsilon \left[ \sum_{i \in J} \langle A(S), -\epsilon_i \mathbf{x}_i \rangle \right] - \mathbb{E}_\epsilon \left[ \sum_{i \in J^c} \langle A(S), -\epsilon_i \mathbf{x}_i \rangle \right] \right] \end{aligned}$$

Now we notice that given  $t_1, \dots, t_n \in [2n]$ , the sets  $J$  and  $J^c$  are fixed and  $S$  only consists of  $\epsilon_i \mathbf{x}_i$  s.t.  $i \in J$ .

Thus,  $S$  is only a function of  $\epsilon_i$  where  $i \in J$ . Hence,

$$\mathbb{E}_\epsilon \left[ \left\langle A(S), - \sum_{i \in J^c} \epsilon_i \mathbf{x}_i \right\rangle \right] = \mathbb{E}_{\epsilon_J} \left[ \mathbb{E}_{\epsilon_{J^c}} \left[ \left\langle A(S), - \sum_{i \in J^c} \epsilon_i \mathbf{x}_i \right\rangle \right] \right] = \mathbb{E}_{\epsilon_J} \left[ \left\langle A(S), - \mathbb{E}_{\epsilon_{J^c}} \left[ \sum_{i \in J^c} \epsilon_i \mathbf{x}_i \right] \right\rangle \right] = 0$$

where we use the shorthand  $\epsilon_J$  to refer to Rademacher random variables  $\epsilon_i$  where  $i \in J$  and  $\epsilon_{J^c}$  refers to the remaining  $\epsilon_i$ 's not in  $J$ . Hence we see that

$$\begin{aligned}
& \sup_{\mathcal{D}} \mathbb{E}_S \left[ L_{\mathcal{D}}(A(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] \\
& \geq \frac{1}{2n} \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2n} \in \mathcal{X}} \mathbb{E}_{t_1, \dots, t_n \sim \text{unif}([2n])} \left[ \mathbb{E}_{\epsilon} \left[ \left\| \sum_{i=1}^{2n} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] - \mathbb{E}_{\epsilon} \left[ \sum_{i \in J} \langle A(S), -\epsilon_i \mathbf{x}_i \rangle \right] \right] \\
& \geq \frac{1}{2n} \sup_{\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathcal{X}} \mathbb{E}_{t_1, \dots, t_n \sim \text{unif}([2n])} \left[ \mathbb{E}_{\epsilon} \left[ \left\| \sum_{i=1}^{2n} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] - \mathbb{E}_{\epsilon} \left[ \left\| \sum_{i \in J} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] \right] \\
& = \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2n} \in \mathcal{X}} \left\{ \mathbb{E}_{\epsilon} \left[ \left\| \frac{1}{2n} \sum_{i=1}^{2n} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] - \frac{1}{2n} \mathbb{E}_{t_1, \dots, t_n \sim \text{unif}([2n])} \mathbb{E}_{\epsilon} \left[ \left\| \sum_{i \in J} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] \right\} \\
& \geq \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2n} \in \mathcal{X}} \mathbb{E}_{\epsilon} \left[ \left\| \sum_{i=1}^{2n} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] - \frac{1}{2n} \mathbb{E}_{t_1, \dots, t_n \sim \text{unif}([2n])} \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{|J|} \in \mathcal{X}} \mathbb{E}_{\epsilon} \left[ \left\| \sum_{i=1}^{|J|} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right]
\end{aligned}$$

Note that since sample size  $|S| = n$ , we have that  $|J| \leq n$ . Therefore,

$$\sup_{\mathbf{x}_1, \dots, \mathbf{x}_{|J|} \in \mathcal{X}} \mathbb{E}_{\epsilon} \left[ \left\| \sum_{i \in J} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] \leq \sup_{\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathcal{X}} \mathbb{E}_{\epsilon} \left[ \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right]$$

(the above is obvious because worst case we can always put  $\mathbf{x}_{|J|+1}, \dots, \mathbf{x}_n = 0$  so the two are equal). Thus we conclude that :

$$\begin{aligned}
\sup_{\mathcal{D}} \mathbb{E}_S \left[ L_{\mathcal{D}}(A(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] & \geq \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2n} \in \mathcal{X}} \mathbb{E}_{\epsilon} \left[ \left\| \sum_{i=1}^{2n} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] - \frac{1}{2n} \mathbb{E}_{t_1, \dots, t_n \sim \text{unif}([2n])} \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{|J|} \in \mathcal{X}} \mathbb{E}_{\epsilon} \left[ \left\| \sum_{i=1}^{|J|} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] \\
& \geq \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2n} \in \mathcal{X}} \mathbb{E}_{\epsilon} \left[ \left\| \sum_{i=1}^{2n} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] - \frac{1}{2n} \sup_{\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathcal{X}} \mathbb{E}_{\epsilon} \left[ \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] \\
& = \mathcal{R}_{2n}^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) - \frac{1}{2} \mathcal{R}_n^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X}))
\end{aligned}$$

This conclude the lemma.  $\square$

*Proof of Theorem 86.* Applying the upper bound guaranteed by the assumption of the theorem and using the lower bound from lemma 84 we get that for any  $n \in \mathbb{N}$ :

$$\mathcal{R}_{2n}^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) \leq \frac{1}{2} \mathcal{R}_n^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) + \frac{V}{n^{1/q}}$$

Expanding the recursive inequality on the right above we get that

$$\mathcal{R}_{2n}^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) \leq \frac{V}{n^{1/q}} \left( 1 + \frac{1}{2^{1/p}} + \frac{1}{2^{2/p}} + \dots \right) \leq \frac{2^{1/p}}{2^{1/p} - 1} \frac{V}{n^{1/q}}$$

Thus we conclude the first inequality in the theorem statement. As for the upper bound on the fat shattering dimension, we have that whenever  $\beta \leq \mathcal{R}_n^{\text{id}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X}))$ ,

$$\text{fat}_\beta^{\text{id}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) \leq \frac{n (\mathcal{R}_n^{\text{id}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})))^2}{\beta^2} \leq n$$

Hence we conclude that for all  $n$  such that  $\mathcal{R}_n^{\text{id}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) \geq \beta$ ,  $\text{fat}_\beta^{\text{id}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) \leq n$ . In other words,

$$\text{fat}_\beta^{\text{id}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) \leq \inf\{n : \mathcal{R}_n^{\text{id}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) \geq \beta\} \inf\{2n : \mathcal{R}_{2n}^{\text{id}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) \geq \beta\}$$

However since we already proved that for all  $n \in \mathbb{N}$ ,  $\mathcal{R}_{2n}^{\text{id}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) \leq \frac{5V}{n^{1/q}}$  we can conclude that for any  $\beta > 0$ ,

$$\text{fat}_\beta^{\text{id}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) \leq \left(\frac{5V}{\beta}\right)^q.$$

This concludes the theorem. □

*Proof of Lemma 87.* Consider instance space  $\mathcal{Z} = \mathcal{X} \times [-1, 1]$  and the non-negative function  $\phi : \mathbb{R} \times [-1, 1] \mapsto \mathbb{R}^+$  that is 1-smooth (in its first argument) given by :

$$\phi(z, y) = \begin{cases} (z - y)^2 & \text{if } |z - y| \leq \frac{1}{2} \\ |z - y| - \frac{1}{4} & \text{otherwise} \end{cases}$$

Now the loss function we consider is given by  $\ell(\mathbf{h}, (\mathbf{x}, y)) = \phi(\langle \mathbf{h}, \mathbf{x} \rangle, y)$ . Since  $\phi$  is 1-smooth in its first argument, we see that

$$\begin{aligned} \|\nabla \ell(\mathbf{h}, (\mathbf{x}, y)) - \nabla \ell(\mathbf{h}', (\mathbf{x}, y))\|_{\mathcal{X}} &= |\partial \phi(\langle \mathbf{h}, \mathbf{x} \rangle, y) - \partial \phi(\langle \mathbf{h}', \mathbf{x} \rangle, y)| \|\mathbf{x}\|_{\mathcal{X}} \leq |\langle \mathbf{h} - \mathbf{h}', \mathbf{x} \rangle| \|\mathbf{x}\|_{\mathcal{X}} \\ &\leq \|\mathbf{h} - \mathbf{h}'\|_{\mathcal{X}^*} \|\mathbf{x}\|_{\mathcal{X}}^2 \leq \|\mathbf{h} - \mathbf{h}'\|_{\mathcal{X}^*} \end{aligned}$$

Thus we conclude that the loss is 1-smooth and so belongs to class  $\mathcal{Z}_{\text{smooth}}$ . The distribution we pick for showing the lower bound is a slight modification of the one used in Lemma 84. We start by picking elements  $\mathbf{x}_1, \dots, \mathbf{x}_{2nL^*} \in \mathcal{X}$  and then draw Rademacher random variables  $\epsilon \in \{\pm 1\}^{2nL^*}$  uniformly at random. We shall now use this to construct a distribution  $\mathcal{D}'_\epsilon$  over instances which is the one we shall use for the lower bound. Specifically, given a draw  $\epsilon \in \{\pm 1\}^{2nL^*}$  consider the distribution  $\mathcal{D}_\epsilon$  on  $\mathcal{X}$  to be the uniform distribution over the set  $\{\epsilon_1 \mathbf{x}_1, \dots, \epsilon_{2nL^*} \mathbf{x}_{2nL^*}\}$ . Now the distribution  $\mathcal{D}'_\epsilon$  is the distribution on  $\mathcal{Z} = \mathcal{X} \times [-1, 1]$  to be the one that picks  $(\mathbf{0}, 0)$  with probability  $1 - L^*$  and with probability  $L^*$ , draws input instance  $\mathbf{x} \in \mathcal{X}$  i.i.d. from distribution  $\mathcal{D}_\epsilon$  and deterministically picks label  $y = -1$  (or whatever  $\sup_{\mathbf{h} \in \mathcal{H}, \mathbf{x} \in \mathcal{X}} \langle \mathbf{h}, \mathbf{x} \rangle$  is with appropriate scaling changes but for simplicity let us assume its bounded by 1).

Note that for any  $\mathbf{h} \in \mathcal{H}$ ,

$$L_{\mathcal{D}'_\epsilon}(\mathbf{h}) = L^* \mathbb{E}_{\mathcal{D}'_\epsilon} \left[ \left| \langle \mathbf{h}, \mathbf{x} \rangle - y \right| - \frac{1}{4} \right] = \frac{1}{2n} \sum_{i=1}^{2nL^*} \left( \langle \mathbf{h}, \epsilon_i \mathbf{x}_i \rangle - 1 - \frac{1}{4} \right).$$

Hence we see that for any  $\epsilon \in \{\pm 1\}^{2n}$ ,

$$L_{\mathcal{D}'_\epsilon}(A(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}'_\epsilon}(\mathbf{h}) = \frac{1}{2n} \sum_{i=1}^{2nL^*} \langle A(S), \epsilon_i \mathbf{x}_i \rangle - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{2n} \sum_{i=1}^{2nL^*} \langle \mathbf{h}, \epsilon_i \mathbf{x}_i \rangle$$

Now consider sample  $S$  of size  $n$  drawn iid from the distribution  $\mathcal{D}'_\epsilon$ . Note that for any learning algorithm  $\mathbf{A}$ ,

$$\begin{aligned} \sup_{\mathcal{D}} \mathbb{E}_S \left[ L_{\mathcal{D}}(A(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] &\geq \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2nL^*}} \mathbb{E}_\epsilon \mathbb{E}_{S \sim \mathcal{D}'_\epsilon} \left[ L_{\mathcal{D}'_\epsilon}(A(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}'_\epsilon}(\mathbf{h}) \right] \\ &= \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2nL^*}} \mathbb{E}_\epsilon \mathbb{E}_{S \sim \mathcal{D}'_\epsilon} \left[ \frac{1}{2n} \sum_{i=1}^{2nL^*} \langle A(S), \epsilon_i \mathbf{x}_i \rangle - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{2n} \sum_{i=1}^{2nL^*} \langle \mathbf{h}, \epsilon_i \mathbf{x}_i \rangle \right] \\ &= \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2nL^*}} \mathbb{E}_\epsilon \mathbb{E}_{S \sim \mathcal{D}'_\epsilon} \left[ \frac{1}{2n} \sum_{i=1}^{2nL^*} \langle A(S), \epsilon_i \mathbf{x}_i \rangle + \left\| \frac{1}{2n} \sum_{i=1}^{2nL^*} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] \end{aligned}$$

Proceeding in similar fashion as in the proof of Lemma 84, we see that one can rewrite sampling  $S$  from the distribution as follows : First we sample  $m$  from binomial distribution  $\text{Binomial}(L^*, n)$  (represents the samples for which  $y \neq 0$ ). Next, numbers  $t_1, \dots, t_m$  is drawn uniformly at random from the set  $[2nL^*]$ . Hence we get the inequality :

$$\begin{aligned} \sup_{\mathcal{D}} \mathbb{E}_S \left[ L_{\mathcal{D}}(A(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] &\geq \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2nL^*}} \mathbb{E}_\epsilon \mathbb{E}_{m \sim \text{Binomial}(n, L^*)} \mathbb{E}_{t_1, \dots, t_m \sim \text{Unif}([2nL^*])} \left[ \frac{1}{2n} \sum_{i=1}^{2nL^*} \langle A(S), \epsilon_i \mathbf{x}_i \rangle + \left\| \frac{1}{2n} \sum_{i=1}^{2nL^*} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] \\ &= \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2nL^*}} \mathbb{E}_{m \sim \text{Binomial}(n, L^*)} \mathbb{E}_{t_1, \dots, t_m \sim \text{Unif}([2nL^*])} \left[ \mathbb{E}_\epsilon \left[ \left\langle A(S), \frac{1}{2n} \sum_{i=1}^{2nL^*} \epsilon_i \mathbf{x}_i \right\rangle \right] + \mathbb{E}_\epsilon \left[ \left\| \frac{1}{2n} \sum_{i=1}^{2nL^*} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] \right] \end{aligned}$$

Now let the set  $J \subset [2nL^*]$  be the set  $J = \{i \in [2nL^*] : i \in \{t_1, \dots, t_m\}\}$  that is the set of indices  $i$  such that  $(\mathbf{x}_i, -1)$  appeared at least once in the sample  $S$  provided to the learner. Also let  $J^c \in [2nL^*]$  stand for the complement of the set  $J$ . Following the same line of proof as in Lemma 84 noting that  $S$  only depends on the  $\epsilon$ 's that occur in the sample  $S$ , we can see that

$$\mathbb{E}_\epsilon \left[ \left\langle A(S), \frac{1}{2n} \sum_{i=1}^{2nL^*} \epsilon_i \mathbf{x}_i \right\rangle \right] \geq -\mathbb{E}_\epsilon \left[ \left\| \frac{1}{2n} \sum_{i \in J} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right]$$

and so we conclude that :

$$\begin{aligned}
& \sup_{\mathcal{D}} \mathbb{E}_S \left[ L_{\mathcal{D}}(A(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] \\
& \geq \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2nL^*}} \mathbb{E}_{m \sim \text{Binomial}(n, L^*)} \mathbb{E}_{t_1, \dots, t_m \sim \text{Unif}(\{[2nL^*]\})} \left[ \mathbb{E}_{\epsilon} \left[ \left\| \frac{1}{2n} \sum_{i=1}^{2nL^*} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] - \mathbb{E}_{\epsilon} \left[ \left\| \frac{1}{2n} \sum_{i \in J} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] \right] \\
& \geq \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2nL^*}} \mathbb{E}_{m \sim \text{Binomial}(n, L^*)} \left[ \mathbb{E}_{\epsilon} \left[ \left\| \frac{1}{2n} \sum_{i=1}^{2nL^*} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] - \sup_{\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{X}} \mathbb{E}_{\epsilon} \left[ \left\| \frac{1}{2n} \sum_{i=1}^{\min\{m, 2L^*n\}} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] \right]
\end{aligned}$$

However note that with probability  $1/2$ ,  $m \leq nL^*$  and so we have that

$$\begin{aligned}
\sup_{\mathcal{D}} \mathbb{E}_S \left[ L_{\mathcal{D}}(A(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] & \geq \frac{1}{2} \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{2nL^*}} \mathbb{E}_{\epsilon} \left[ \left\| \frac{1}{2n} \sum_{i=1}^{2nL^*} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] - \sup_{\mathbf{x}_1, \dots, \mathbf{x}_{nL^*} \in \mathcal{X}} \mathbb{E}_{\epsilon} \left[ \left\| \frac{1}{2n} \sum_{i=1}^{L^*n} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] \\
& = \frac{L^*}{2} \left( \mathcal{R}_{2L^*n}^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) - \frac{1}{2} \mathcal{R}_{L^*n}^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) \right)
\end{aligned}$$

This conclude the lemma.  $\square$

*Proof of Lemma 88.* First, since both sides below are homogenous, the premise of the lemma can be rewritten as, for all  $n \in \mathbb{N}$  and all  $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathcal{B}^*$ ,

$$\mathbb{E}_{\epsilon} \left[ \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] \leq D n^{\frac{1}{r}} \max_{i \in [n]} \|\mathbf{x}_i\|_{\mathcal{X}} \quad (8.5.1)$$

Let  $S = (\sum_{i=1}^n \|\mathbf{x}_i\|_{\mathcal{X}}^p)^{1/p}$ , define

$$\begin{aligned}
I_k & := \left\{ i \geq 1 \mid \frac{S}{2^{(k+1)/p}} < \|\mathbf{x}_i\|_{\mathcal{X}} \leq \frac{S}{2^{k/p}} \right\}, \\
T_0^{(k)} & := \inf\{i \in I_k\} \text{ and} \\
\forall m \in \mathbb{N}, T_m^{(k)} & := \inf\{i > T_{m-1}^{(k)}, i \in I_k\}
\end{aligned}$$

Note that,  $S^p \geq \sum_{i \in I_k} \|\mathbf{x}_i\|_{\mathcal{X}}^p > \frac{S^p |I_k|}{2^{(k+1)}}$  and so we get that  $|I_k| < 2^{k+1}$ . From this, using the premise in

Equation 8.5.1 we conclude that

$$\begin{aligned}
\mathbb{E}_\epsilon \left[ \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] &\leq \sum_{k \geq 0} \mathbb{E}_\epsilon \left[ \left\| \sum_{i \in I_k} \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right] = \sum_{k \geq 0} \mathbb{E}_\epsilon \left[ \left\| \sum_{i \geq 0} \epsilon_{T_i^{(k)}} \mathbf{x}_{T_i^{(k)}} \right\|_{\mathcal{H}^*} \right] \\
&\leq \sum_{k \geq 0} \left( D \{|I_k|^{1/r}\} \{ \sup_{i \in I_k} \|\mathbf{x}_i\|_{\mathcal{X}} \} \right) \\
&\leq \sum_{k \geq 0} \left( D 2^{(k+1)/r} \sup_{i \in I_k} \|\mathbf{x}_i\|_{\mathcal{X}, \infty} \right) \\
&\leq \sum_{k \geq 0} \left( D 2^{(k+1)/r} 2^{-k/p} S \right) \\
&= D 2^{1/r} \sum_{k \geq 0} 2^{k(\frac{1}{r} - \frac{1}{p})} S \leq \frac{2D}{1 - 2^{(\frac{1}{r} - \frac{1}{p})}} S \leq \frac{12D}{r - p} S \\
&= \frac{12D}{r - p} \left( \sum_{i=0}^n \|\mathbf{x}_i\|_{\mathcal{X}}^p \right)^{1/p}
\end{aligned}$$

We conclude the proof by using Kahane Inequality (see [85]) which asserts that for any  $p \in [1, 2]$ ,

$$\left( \mathbb{E}_\epsilon \left[ \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*}^p \right] \right)^{1/p} \leq \sqrt{2} \mathbb{E}_\epsilon \left[ \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i \right\|_{\mathcal{H}^*} \right]$$

□

*Proof of Proposition 58.* Since we are dealing with convex loss, using Jensen's inequality we have that

$$L_{\mathcal{D}} \left( \frac{1}{n} \sum_{t=1}^n \mathbf{A}(z_{1:t}) \right) \leq \frac{1}{n} \sum_{t=1}^n L_{\mathcal{D}} (\mathbf{A}(z_{1:t})) .$$

Hence we see that

$$\begin{aligned}
\mathbb{E}_{S \sim \mathcal{D}^n} \left[ \left( \frac{1}{n} \sum_{t=1}^n \mathbf{A}(z_{1:t-1}) \right) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] &\leq \frac{1}{n} \sum_{t=1}^n \mathbb{E}_S [L_{\mathcal{D}}(\mathbf{A}(z_{1:t-1}))] - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n L_{\mathcal{D}}(\mathbf{h}) \\
&= \frac{1}{n} \sum_{t=1}^n \mathbb{E}_{S \sim \mathcal{D}^n} [\ell(\mathbf{A}(z_{1:t-1}), z_t)] - \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \mathbb{E}_{S \sim \mathcal{D}^n} [\ell(\mathbf{h}, z_t)] \\
&= \mathbb{E}_{S \sim \mathcal{D}^n} \left[ \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{A}(z_{1:t-1}), z_t) \right] - \inf_{\mathbf{h} \in \mathcal{H}} \mathbb{E}_S \left[ \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \right] \\
&\leq \mathbb{E}_{S \sim \mathcal{D}^n} \left[ \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{A}(z_{1:t-1}), z_t) \right] - \mathbb{E}_S \left[ \inf_{\mathbf{h} \in \mathcal{H}} \frac{1}{n} \sum_{t=1}^n \ell(\mathbf{h}, z_t) \right] \\
&= \mathbb{E}_{S \sim \mathcal{D}^n} [\mathbf{R}_n(\mathbf{A}, z_1, \dots, z_n)] \\
&\leq \sup_{z_1, \dots, z_n \in \mathcal{Z}} \mathbf{R}_n(\mathbf{A}, z_1, \dots, z_n)
\end{aligned}$$

where the second step is because  $\mathbf{A}(z_1, \dots, z_{t-1})$  only depends on  $z_1, \dots, z_{t-1}$  and so  $\mathbb{E}_S [\ell(\mathbf{A}(z_1, \dots, z_{t-1}), z_t)] = \mathbb{E}_S [L_{\mathcal{D}}(\mathbf{A}(z_1, \dots, z_{t-1}))]$ . The second part of the proposition is from the fact that the above holds for any online learning algorithm.  $\square$

**Proposition 97.** *If  $(\mathcal{H}^*, \mathcal{X})$  has type  $p$  with some constant  $C_p^{\text{iid}}$  then for  $q = \frac{p}{p-1}$ , the pair  $(\mathcal{H}, \mathcal{X}^*)$  has co-type  $q$  with constant  $\frac{1}{C_p^{\text{iid}}}$*

*Proof.* Given  $\mathbf{h}_1, \dots, \mathbf{h}_n \in \mathcal{B}^*$  for any  $\epsilon > 0$  we can pick  $\mathbf{x}_1, \dots, \mathbf{x}_n$  such that

$$\sum_{t=1}^n \langle \mathbf{h}_t, \mathbf{x}_t \rangle \geq (1 - \epsilon) \left( \sum_{t=1}^n \|\mathbf{h}_t\|_{\mathcal{X}^*}^q \right)^{1/q} \left( \sum_{t=1}^n \|\mathbf{x}_t\|_{\mathcal{X}}^p \right)^{1/p} \quad (8.5.2)$$

On the other hand we have that,

$$\begin{aligned}
\sum_{t=1}^n \langle \mathbf{h}_t, \mathbf{x}_t \rangle &= \mathbb{E}_{\epsilon} \left[ \left\langle \sum_{t=1}^n \epsilon_t \mathbf{h}_t, \sum_{t=1}^n \epsilon_t \mathbf{x}_t \right\rangle \right] \\
&\leq \left( \mathbb{E}_{\epsilon} \left[ \left\| \sum_{t=1}^n \epsilon_t \mathbf{h}_t \right\|_{\mathcal{H}}^q \right] \right)^{1/q} \left( \mathbb{E}_{\epsilon} \left[ \left\| \sum_{t=1}^n \epsilon_t \mathbf{x}_t \right\|_{\mathcal{H}^*}^p \right] \right)^{1/p} \\
&\leq C_p^{\text{iid}} \left( \mathbb{E}_{\epsilon} \left[ \left\| \sum_{t=1}^n \epsilon_t \mathbf{h}_t \right\|_{\mathcal{H}}^q \right] \right)^{1/q} \left( \sum_{t=1}^n \|\mathbf{x}_t\|_{\mathcal{X}}^p \right)^{1/p}
\end{aligned}$$

where the last step is by the type inequality for  $(\mathcal{H}, \mathcal{X})$ . Combining the above and Equation 8.5.2 and taking limit of  $\epsilon \rightarrow 0$  proves the statement.  $\square$

**Definition 37.** For  $p \in [1, 2]$  the pair  $(\mathcal{H}^*, \mathcal{X})$  is said to be  $p$ -convex with constant  $K_p$  if for any  $\mathbf{x}_1, \dots, \mathbf{x}_N \in$



$\mathcal{X}$ ,

$$\left\| \left( \sum_{t=1}^N |\mathbf{x}_t|^p \right)^{1/p} \right\|_{\mathcal{H}^*} \leq K_p \left( \sum_{t=1}^N \|\mathbf{x}_t\|_{\mathcal{X}}^p \right)^{1/p}$$

**Definition 38.** For  $q \in [2, \infty)$  the pair  $(\mathcal{H}, \mathcal{X}^*)$  is said to be  $q$ -concave with constant  $\bar{K}_q$  if for any  $K \leq \bar{K}_q(\mathcal{H}, \mathcal{X}^*)$  and any  $\mathbf{x}_1, \dots, \mathbf{x}_N \in \mathcal{X}$ ,

$$\left\| \left( \sum_{t=1}^N |\mathbf{x}_t|^q \right)^{1/q} \right\|_{\mathcal{H}^*} \geq \frac{1}{\bar{K}_q} \left( \sum_{t=1}^N \|\mathbf{x}_t\|_{\mathcal{X}}^q \right)^{1/q}$$

**Lemma 98.** If the pair  $(\mathcal{H}^*, \mathcal{X})$  is of type  $p$  with constant  $C_p^{\text{iid}}$  then, for any  $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathcal{X}$ ,

$$\sqrt{2} C_p^{\text{iid}} \left( \sum_{t=1}^n \|\mathbf{x}_t\|_{\mathcal{X}^*}^p \right)^{1/p} \geq \left\| \left( \sum_{t=1}^n |\mathbf{x}_t|^p \right)^{\frac{1}{p}} \right\|_{\mathcal{H}^*}$$

That is it is  $p$ -convex with constant  $\sqrt{2} C_p^{\text{iid}}$ .

*Proof.* By type  $p$  with constant  $C_p^{\text{iid}}$ ,

$$\begin{aligned} C_p^{\text{iid}} \left( \sum_{t=1}^n \|\mathbf{x}_t\|_{\mathcal{X}^*}^p \right)^{1/p} &\geq \left( \mathbb{E}_\epsilon \left[ \left\| \sum_{t=1}^n \epsilon_t \mathbf{x}_t \right\|_{\mathcal{H}^*}^p \right] \right)^{\frac{1}{p}} \geq \mathbb{E}_\epsilon \left[ \left\| \sum_{t=1}^n \epsilon_t \mathbf{x}_t \right\|_{\mathcal{H}^*} \right] \geq \left\| \mathbb{E}_\epsilon \left[ \sum_{t=1}^n \epsilon_t \mathbf{x}_t \right] \right\|_{\mathcal{H}^*} \\ &\geq \frac{1}{\sqrt{2}} \left\| \left( \sum_{t=1}^n |\mathbf{x}_t|^p \right)^{\frac{1}{p}} \right\|_{\mathcal{H}^*} \end{aligned}$$

where the last inequality is got by using scalar version of Kintchine's inequality with functional calculus for Banach lattice (specifically Theorem 1.d.1 of [75] by noting that expectation w.r.t.  $\epsilon$  can be written as finite average). This concludes the proof.  $\square$

**Proposition 99.** If  $(\mathcal{H}^*, \mathcal{X})$  is  $p$ -convex with some constant  $K_p$  then for  $q = \frac{p}{p-1}$ , the pair  $(\mathcal{H}, \mathcal{X}^*)$  is  $q$ -concave with constant  $K_p$ . That is for any  $\mathbf{h}_1, \dots, \mathbf{h}_n \in \mathcal{H}$

$$\left\| \left( \sum_{t=1}^n |\mathbf{h}_t|^q \right)^{1/q} \right\|_{\mathcal{H}} \geq \frac{1}{K_p} \left( \sum_{t=1}^n \|\mathbf{h}_t\|_{\mathcal{X}^*}^q \right)^{1/q}$$

*Proof.* Given  $\mathbf{h}_1, \dots, \mathbf{h}_n \in \mathcal{B}^*$  for any  $\epsilon > 0$  we can pick  $\mathbf{x}_1, \dots, \mathbf{x}_n$  such that

$$\sum_{t=1}^n \langle \mathbf{h}_t, \mathbf{x}_t \rangle \geq (1 - \epsilon) \left( \sum_{t=1}^n \|\mathbf{h}_t\|_{\mathcal{X}^*}^q \right)^{1/q} \left( \sum_{t=1}^n \|\mathbf{x}_t\|_{\mathcal{X}}^p \right)^{1/p} \quad (8.5.3)$$

On the other hand, using 1.d.2 (iii) [75] we can conclude that

$$\begin{aligned}
\sum_{t=1}^n \langle \mathbf{h}_t, \mathbf{x}_t \rangle &\leq \left\langle \left( \sum_{t=1}^n |\mathbf{h}_t|^q \right)^{1/q}, \left( \sum_{t=1}^n |\mathbf{x}_t|^p \right)^{1/p} \right\rangle \\
&\leq \left\| \left( \sum_{t=1}^n |\mathbf{h}_t|^q \right)^{1/q} \right\|_{\mathcal{H}} \left\| \left( \sum_{t=1}^n |\mathbf{x}_t|^p \right)^{1/p} \right\|_{\mathcal{H}^*} \\
&\leq K_p \left\| \left( \sum_{t=1}^n |\mathbf{h}_t|^q \right)^{1/q} \right\|_{\mathcal{H}} \left( \sum_{t=1}^n \|\mathbf{x}_t\|_{\mathcal{X}}^p \right)^{1/p}
\end{aligned}$$

where the last step is by the  $p$ -convex inequality of  $(\mathcal{H}^*, \mathcal{X})$ . Combining the above and Equation ?? and taking limit of  $\epsilon \rightarrow 0$  proves the statement.  $\square$

**Proposition 100.** For any  $n \in \mathbb{N}$ , any sequence real valued tree  $\mathbf{a}$  of depth  $n$  and for any  $1 \leq p \leq 2 \leq r < \infty$ ,

$$\frac{18^3 \sqrt{2} r^3 p^{\frac{3}{2}}}{\sqrt{p-1} (r-1)} \left( \mathbb{E}_\epsilon \left[ \left| \sum_{t=1}^n \epsilon_t \mathbf{a}_t(\epsilon) \right|^p \right] \right)^{1/p} \geq \left( \mathbb{E}_\epsilon \left[ \left| \sum_{t=1}^n \epsilon_t \mathbf{a}_t(\epsilon) \right|^r \right] \right)^{1/r}$$

*Proof.*

$$\begin{aligned}
\left( \mathbb{E}_\epsilon \left[ \sum_{t=1}^n |\mathbf{a}_t(\epsilon)|^p \right] \right)^{\frac{1}{p}} &\geq \left( \mathbb{E}_\epsilon \left[ \left( \sum_{t=1}^n |\mathbf{a}_t(\epsilon)|^2 \right)^{\frac{p}{2}} \right] \right)^{\frac{1}{p}} && \forall p \leq 2, \|\cdot\|_p^p \geq \|\cdot\|_2^p \\
&\geq \frac{\sqrt{p-1}}{18 p^{\frac{3}{2}}} \left( \mathbb{E}_{\epsilon, \epsilon'} \left[ \left| \sum_{t=1}^n \epsilon'_t \epsilon_t \mathbf{a}_t(\epsilon) \right|^p \right] \right)^{1/p} && \text{Burkholder's Inequality [86]} \\
&\geq \frac{\sqrt{p-1}}{18 \sqrt{2(r-1)} p^{\frac{3}{2}}} \left( \mathbb{E}_{\epsilon, \epsilon'} \left[ \left| \sum_{t=1}^n \epsilon'_t \epsilon_t \mathbf{a}_t(\epsilon) \right|^r \right] \right)^{1/r} \\
&\geq \frac{\sqrt{p-1} (r-1)}{18^3 \sqrt{2} r^3 p^{\frac{3}{2}}} \left( \mathbb{E}_\epsilon \left[ \left| \sum_{t=1}^n \epsilon_t \mathbf{a}_t(\epsilon) \right|^r \right] \right)^{1/r}
\end{aligned}$$

This concludes the proof.  $\square$

*Proof of Lemma 93.* Consider any  $\mathcal{X}$  valued tree  $\mathbf{x}$  of infinite depth. We start by noting that type  $p$  with

constant  $C_p^{\text{iid}}$  implies  $p$ -convexity with constant  $\sqrt{2}C_p^{\text{iid}}$  and so

$$\begin{aligned}
\frac{18^3 \sqrt{8} p^{\frac{3}{2}} r^3 C_p^{\text{iid}} \tilde{C}_r}{\sqrt{p-1}(r-1)} \left( \mathbb{E}_\epsilon \left[ \sum_{t=1}^n \|\mathbf{x}_t(\epsilon)\|_{\mathcal{X}}^p \right] \right)^{\frac{1}{p}} &= \frac{18^3 \sqrt{8} p^{\frac{3}{2}} r^3 C_p^{\text{iid}} \tilde{C}_r}{\sqrt{p-1}(r-1)} \left( \frac{1}{2^n} \sum_{\epsilon \in \{\pm 1\}^n} \sum_{t=1}^n \|\mathbf{x}_t(\epsilon)\|_{\mathcal{X}}^p \right)^{\frac{1}{p}} \\
&\geq \frac{18^3 \sqrt{4} p^{\frac{3}{2}} r^3 \tilde{C}_r}{\sqrt{p-1}(r-1)} \left\| \left( \frac{1}{2^n} \sum_{\epsilon \in \{\pm 1\}^n} \sum_{t=1}^n |\mathbf{x}_t(\epsilon)|^p \right)^{\frac{1}{p}} \right\|_{\mathcal{H}^*} \\
&= \frac{18^3 \sqrt{4} p^{\frac{3}{2}} r^3 \tilde{C}_r}{\sqrt{p-1}(r-1)} \left\| \left( \mathbb{E}_\epsilon \left[ \sum_{t=1}^n |\mathbf{x}_t(\epsilon)|^p \right] \right)^{\frac{1}{p}} \right\|_{\mathcal{H}^*} \quad (8.5.4)
\end{aligned}$$

Now note that for the real-valued case by Proposition 100 we have that, for any real valued tree  $\mathbf{a}$  of depth  $n$  and for any  $1 \leq p \leq 2 \leq r < \infty$ ,

$$\frac{18^3 \sqrt{2} r^3 p^{\frac{3}{2}}}{\sqrt{p-1}(r-1)} \left( \mathbb{E}_\epsilon \left[ \left| \sum_{t=1}^n \epsilon_t \mathbf{a}_t(\epsilon) \right|^p \right] \right)^{1/p} \geq \left( \mathbb{E}_\epsilon \left[ \left| \sum_{t=1}^n \epsilon_t \mathbf{a}_t(\epsilon) \right|^r \right] \right)^{1/r}$$

Since both sides the expressions are degree one homogenous, applying Theorem 1.d.1 [75] (by expanding out the tree of depth  $n$  to its  $2^n - 1$  elements and noting that expectation w.r.t.  $n$  Rademacher variables is in fact a finite average of  $2^n$  signs) we conclude that

$$\frac{18^3 \sqrt{2} r^3 p^{\frac{3}{2}}}{\sqrt{p-1}(r-1)} \left( \mathbb{E}_\epsilon \left[ \sum_{t=1}^n |\mathbf{x}_t(\epsilon)|^p \right] \right)^{\frac{1}{p}} \geq \left( \mathbb{E}_\epsilon \left[ \left| \sum_{t=1}^n \epsilon_t \mathbf{x}_t(\epsilon) \right|^r \right] \right)^{1/r}$$

Plugging this back in Equation 8.5.4 and noting that co-type  $r$  with constant  $\tilde{C}_r$  of the pair  $(\mathcal{H}^*, \mathcal{H}^*)$  implies its  $r$ -concavity with constant  $\sqrt{2}\tilde{C}_r$ , we conclude that :

$$\begin{aligned}
\frac{18^3 \sqrt{8} p^{\frac{3}{2}} r^3 C_p^{\text{iid}} \tilde{C}_r}{\sqrt{p-1}(r-1)} \left( \mathbb{E}_\epsilon \left[ \sum_{t=1}^n \|\mathbf{x}_t(\epsilon)\|_{\mathcal{X}}^p \right] \right)^{\frac{1}{p}} &\geq \sqrt{2} \tilde{C}_r \left\| \left( \mathbb{E}_\epsilon \left[ \left| \sum_{t=1}^n \epsilon_t \mathbf{x}_t(\epsilon) \right|^r \right] \right)^{\frac{1}{r}} \right\|_{\mathcal{H}^*} \\
&\geq \left( \mathbb{E}_\epsilon \left[ \left\| \sum_{t=1}^n \epsilon_t \mathbf{x}_t(\epsilon) \right\|_{\mathcal{H}^*}^r \right] \right)^{\frac{1}{r}} \\
&\geq \left( \mathbb{E}_\epsilon \left[ \left\| \sum_{t=1}^n \epsilon_t \mathbf{x}_t(\epsilon) \right\|_{\mathcal{H}^*}^p \right] \right)^{1/p}
\end{aligned}$$

Noting that  $r^3/(r-1) \leq 2r^2$  and that  $p^{3/2} \leq \sqrt{8}$  and by definition of  $C_p$  we conclude the proof.  $\square$

*Proof of Lemma 95.* Consider any  $\mathcal{X}$  valued infinite depth tree  $\mathbf{x}$ . We have,

$$\begin{aligned}
 \mathbb{E}_\epsilon \left[ \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} \right] &\leq D \mathbb{E}_{\epsilon, \epsilon'} \left[ \left\| \sum_{i=1}^n \epsilon_i \epsilon'_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} \right] \\
 &= D \mathbb{E}_\epsilon \left[ \mathbb{E}_{\epsilon'} \left[ \left\| \sum_{i=1}^n \epsilon_i \epsilon'_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*} \right] \right] \\
 &\leq DC_r^{\text{iid}} \mathbb{E}_\epsilon \left[ \left( \sum_{i=1}^n \|\epsilon_i \mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^r \right)^{1/r} \right] \\
 &\leq DC_r^{\text{iid}} \sup_{\epsilon \in \{\pm 1\}^n} \left( \sum_{i=1}^n \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^r \right)^{1/r}
 \end{aligned}$$

Now this is effectively what we had in the proof of Lemma 67 and exactly as we did there applying Lemma's 79, 80 and 81 and repeating the steps in proof of Lemma 67 we can conclude that

$$\mathbb{E}_\epsilon \left[ \left\| \sum_{i=1}^n \epsilon_i \mathbf{x}_i(\epsilon) \right\|_{\mathcal{H}^*}^{p'} \right] \leq \frac{92DC_r^{\text{iid}}}{(r-p')} \sum_{i=1}^n \mathbb{E}_\epsilon \left[ \|\mathbf{x}_i(\epsilon)\|_{\mathcal{X}}^{p'} \right]$$

This concludes the proof. □

## 8.6 Discussion

The highlight of this chapter was that we showed that for most commonly occurring learning problems constant  $C_p$  can be almost bounded by  $C_p^{\text{iid}}$  which we could use to conclude that mirror descent was near optimal in terms of both learning rate and oracle complexity for statistical convex problems too. In fact for the common examples we saw that the constant factor relating  $C_p$  and  $C_p^{\text{iid}}$  can be a fixed universal constant. We also saw that while many a times oracle complexity could be directly associated with time complexity of the algorithms this is not always true and for the very practical problem of matrix completion with max norm this issue arose leading to the open question of whether we can efficiently solve in the online setting, matrix completion with max-norm with optimal rates.

We also saw strong connection between oracle complexity of offline optimization of convex lipschitz class and statistical complexity of linear class and using this we also informally argued that for large dimensional problems in “reasonable space”, mirror descent algorithm is near optimal even for offline optimization. Overall the aim of this section is to show that mirror descent is near optimal ubiquitously for most practical convex statistical learning problems and for high dimensional offline optimization problems not just in the online learning setting as was shown in the previous chapter.

## Chapter 9

# Optimality of Mirror Descent for Offline Convex Optimization

In this chapter we consider the problem of offline convex optimization. To address the issue of efficiency of optimization methods for the convex optimization problems, we use the notion of oracle complexity introduced by Nemirovski and Yudin in [1]. We show interesting connections between convex optimization and statistical convex learning. Furthermore for several commonly encountered high dimensional convex optimization problems we also show that mirror descent algorithm is near optimal, in terms of oracle complexity, even for offline convex optimization. Based on the results we also show that for several statistical convex learning problems, mirror descent is optimal even when one has access to more powerful oracles that can account for parallel computational methods for optimization. In fact we show that for these problems, parallelization does not help (in improving efficiency of the learning procedure) and simply using single pass mirror descent algorithm is optimal both in terms of rates and efficiency.

Section 9.1 introduces the oracle based offline optimization model of Nemirovski and Yudin [1]. Section 9.2 provides lower bound on oracle complexity of offline convex optimization problems in terms of fat shattering dimension of associated linear class and shows connections between convex optimization and statistical learning and specifically also shows that whenever a the convex optimization problem over function class associated with  $\mathcal{Z}_{\text{Lip}}(\mathcal{X})$  is efficiently optimizable, then the supervised learning problem is also efficiently learnable. Section 9.3 shows that for several high dimensional problems, mirror descent is also optimal for offline convex optimization. Section 9.4 deals with statistical learning with distributed oracles and shows that for several cases, mirror descent is near optimal even when compared to learning algorithms that have access to distributed oracles and thus show that in these cases parallelization does not help. Section 9.5 provides the detailed proofs of this chapter and we conclude with some discussions in Section 9.6.

## 9.1 Oracle-based Offline Convex Optimization

A typical convex optimization algorithm initially picks some point in the feasible set  $\mathcal{H}$  and iteratively updates these points based on some local information about the function it calculates around this point. Examples of these type of procedures are gradient descent methods that uses first order gradient information, newton's method that uses second over hessian information, interior point methods and so on. In fact most procedures one can think of for optimization are based on iteratively updating based on some local information about the function at the point. In general computing the exact computational complexity of these methods is cumbersome and may not even be possible in full generality. To capture efficiency of optimization procedures in a general way, we consider the oracle based optimization problem and associated oracle complexity of the problem. To this end we first formally define an Oracle. Such models have been introduced and analyzed in [1]. In general, given an instance set  $\mathcal{Z}$  indexing convex functions on hypothesis set  $\bar{\mathcal{H}}$ , we use the term oracle to refer to any mapping  $\mathcal{O} : \bar{\mathcal{H}} \times \mathcal{Z} \mapsto \mathcal{I}$  from hypothesis instance pairs to an answer  $I \in \mathcal{I}$ , where set  $\mathcal{I}$  is some arbitrary information set  $\mathcal{I}$ . However at this generality note that the information set  $\mathcal{I}$  could be all of  $\mathcal{Z}$  and Oracle  $\mathcal{O}$  could be the identity mapping. This would defeat the purpose of introducing oracle models and associated oracle complexity. To address this issue we use the notion of local oracle defined by Nemirovski and Yudin [1] and whenever we use the term oracle we mean local oracle. Before we formally define a (local) oracle we first define the neighborhood set of a point  $\mathbf{h} \in \bar{\mathcal{H}}$ . Given  $\delta > 0$  and a point  $\mathbf{h} \in \mathcal{H}$ , the  $\delta$ -neighborhood of the point  $\mathbf{h}$  is the set

$$B_\delta(\mathbf{h}) = \{\mathbf{h}' \in \bar{\mathcal{H}} : \|\mathbf{h} - \mathbf{h}'\|_{\mathcal{H}} \leq \delta\} .$$

**Definition 39.** A (local) oracle  $\mathcal{O} : \bar{\mathcal{H}} \times \mathcal{Z} \mapsto \mathcal{I}$  is a mapping, which, given any point  $\mathbf{h} \in \bar{\mathcal{H}}$  and instances  $z, z' \in \mathcal{Z}$  such that

$$\lim_{\delta \rightarrow 0} \sup_{\mathbf{h}', \mathbf{h}'' \in B_\delta(\mathbf{h})} |\ell(\mathbf{h}', z) - \ell(\mathbf{h}'', z')| = 0 ,$$

outputs answers that satisfy the equality  $\mathcal{O}(z, \mathbf{h}) = \mathcal{O}(z', \mathbf{h})$ .

The definition basically says if two instances  $z, z' \in \mathcal{Z}$  correspond to convex losses  $\ell(\cdot, z)$  and  $\ell(\cdot, z')$  respectively that are indistinguishable about some neighborhood of a point  $\mathbf{h} \in \bar{\mathcal{H}}$ , then querying an oracle about the two instances at this point  $\mathbf{h}$  leads to the same answer. For a query with instance  $\mathbf{h} \in \bar{\mathcal{H}}$  at instance  $z \in \mathcal{Z}$ , the oracle only provides local information about the function  $\ell(\cdot, z)$  at the point of query  $\mathbf{h}$ . To make the concept clearer we provide the following examples of (local) oracles commonly used in practice.

**Example 13 (Zero'th-order Oracle).** This is perhaps the simplest oracle that simply evaluates the given function at the query point and returns the value. That is  $\mathcal{O}(\mathbf{h}, z) = \ell(\mathbf{h}, z)$ . Clearly  $\mathcal{I} \subseteq \mathbb{R}$ .

The zero'th order oracle captures bandit learning problems.

**Example 14 (First-order Oracle).** This is an Oracle that provides the sub-gradient of the function at query point. That is  $\mathcal{O}(\mathbf{h}, z) = \nabla_{\mathbf{h}} \ell(\mathbf{h}, z)$ . Clearly  $\mathcal{I} \subset \mathcal{B}^*$  the dual space of  $\mathcal{B}$  (the banach space containing  $\bar{\mathcal{H}}$ ).

**Example 15** (Second-order Oracle). Consider the example where  $\bar{\mathcal{H}} \subseteq \mathbb{R}^d$  and each  $z \in \mathcal{Z}$  corresponds to twice differentiable convex loss. In this case a second-order oracle is one that returns the hessian of the function at the query point. That is  $\mathcal{O}(\mathbf{h}, z) = \nabla_{\mathbf{h}}^2 \ell(\mathbf{h}, z)$ .

As mentioned earlier since we are restricting ourselves to convex problems, we only need to consider deterministic algorithms. We now generically define Oracle Based Learning/Optimization Algorithm for a given oracle  $\mathcal{O}$ .

**Definition 40.** For a given Oracle  $\mathcal{O}$ , an “Oracle Based Optimization/Learning Algorithm”,  $\mathbf{A}^{\mathcal{O}} : \bigcup_{n \in \mathbb{N}} \mathcal{I}^n \rightarrow \bar{\mathcal{H}}$  is a mapping from a sequences of oracle answers in  $\mathcal{I}$  to an element of hypothesis set  $\bar{\mathcal{H}}$ .

We now describe the oracle-based offline convex optimization protocol. Given  $z \in \mathcal{Z}$  corresponding to convex function  $\ell(\cdot, z)$  (unknown to the learner) the optimization procedure is as follows :

**Oracle-based Offline Optimization Protocol :**

**for**  $t = 1$  **to**  $m$

Pick hypothesis  $\mathbf{h}_t \in \bar{\mathcal{H}}$  for query

Oracle provides answer  $\mathbf{I}_t = \mathcal{O}(\mathbf{h}_t, z)$

**end for**

The goal is to solve the optimization problem

$$\operatorname{argmin}_{\mathbf{h} \in \bar{\mathcal{H}}} \ell(\mathbf{h}, z)$$

and at the end of  $m$  steps the sub-optimality of the *Learner* is given by

$$\ell(\mathbf{h}_m, z) - \inf_{\mathbf{h} \in \bar{\mathcal{H}}} \ell(\mathbf{h}, z) .$$

Given an Oracle-based learning/optimization algorithm,  $\mathbf{A}^{\mathcal{O}}$  we shall use the short-hand,  $\mathbf{I}_1 = \mathcal{O}(\mathbf{h}_1, z)$  and further iteratively, use the notation  $\mathbf{I}_t = \mathcal{O}(\mathbf{h}_t, z)$  where of course each hypothesis  $\mathbf{h}_t = \mathbf{A}^{\mathcal{O}}(\mathbf{I}_1, \dots, \mathbf{I}_{t-1})$  is picked using the Oracle-based learning/optimization algorithm. We now define what it means to be offline optimizable using an oracle  $\mathcal{O}$ .

**Definition 41.** For a given class of convex functions corresponding to instance set  $\mathcal{Z}$  and a given Oracle  $\mathcal{O}$ , the offline oracle complexity of a given “Oracle based optimization/learning algorithm”,  $\mathbf{A}^{\mathcal{O}}$ , is defined as

$$m_{\text{off}}(\epsilon, \mathbf{A}^{\mathcal{O}}, \mathcal{Z}) = \inf \left\{ m \in \mathbb{N} \mid \sup_{z \in \mathcal{Z}} \left\{ \ell(\mathbf{A}^{\mathcal{O}}(\mathbf{I}_1, \dots, \mathbf{I}_m), z) - \inf_{\mathbf{h} \in \bar{\mathcal{H}}} \ell(\mathbf{h}, z) \right\} \leq \epsilon \right\}$$

Further, for the given oracle  $\mathcal{O}$ , the  $\mathcal{O}$ -offline oracle complexity of the given offline optimization problem is defined as  $m_{\text{off}}(\epsilon, \mathcal{Z}, \mathcal{O}) = \inf_{\mathbf{A}^{\mathcal{O}}} m_{\text{off}}(\epsilon, \mathbf{A}^{\mathcal{O}}, \mathcal{Z})$ .

Roughly speaking, given  $\epsilon > 0$ , the oracle complexity of an algorithm is the minimum number of oracle answers needed by the algorithm to guarantee sub-optimality smaller than  $\epsilon$  against any instance  $z \in \mathcal{Z}$ .

Further, the above definition basically implies that for any  $\epsilon$ , there exists an Oracle Based Optimization Algorithm,  $\mathbf{A}^{\mathcal{O}}$  that provides an  $\epsilon$  sub-optimality for any  $z \in \mathcal{Z}$  within  $m(\epsilon, \mathcal{Z}, \mathcal{O})$  steps. While the above definition of oracle complexity tells us what is the best (oracle) efficiency achievable for a problem for a given problem using oracle  $\mathcal{O}$ , one might still wonder if using some other oracle one can improve efficiency. To address this issue we now define oracle complexity of a problem (independent of any particular oracle).

**Definition 42.** For a given class of convex functions corresponding to instance set  $\mathcal{Z}$ , the offline oracle complexity of the given offline optimization problem is defined as

$$m_{\text{off}}(\epsilon, \mathcal{Z}) = \inf_{\mathcal{O}} m_{\text{off}}(\epsilon, \mathcal{Z}, \mathcal{O}) .$$

Further we say that a given problem is oracle-based offline optimizable if there exists some Oracle  $\mathcal{O}$  such that  $\forall \epsilon > 0, m_{\text{off}}(\epsilon, \mathcal{Z}, \mathcal{O}) < \infty$ .

## 9.2 Lower Bounding Oracle Complexity: Connections to Statistical Convex Learning

In this section we establish some connections between offline optimization of convex Lipschitz instance class  $\mathcal{Z}_{\text{Lip}}(\mathcal{X})$  and fat-shattering dimension and Rademacher complexities of the linear function class. We also establish interesting connection between convex optimization and statistical convex learning.

The following lemma which lower bounds oracle complexity by fat-shattering dimension of the corresponding linear function class is based on the proof technique for lower bounds on oracle complexity for offline optimization of convex Lipschitz function classes in [1].

**Lemma 101.** *For any  $\mathcal{H} \subset \mathcal{B}$  and  $\mathcal{X} \subset \mathcal{B}^*$  that are convex and centrally symmetric,,*

$$m_{\text{off}}(\beta, \mathcal{Z}_{\text{Lip}}(\mathcal{X})) \geq \text{fat}_{2\beta}^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) .$$

Based on the above lemma and relationship between fat-shattering dimension and statistical Rademacher complexity, we prove the following theorem which bounds Rademacher complexity of the linear function class in terms using any polynomial upper bound on oracle complexity of offline optimization of the convex Lipschitz function class.

**Theorem 102.** *If there exists some  $V > 0$  and  $q \in (0, \infty)$  such that for any  $\beta > 0$ ,*

$$m_{\text{off}}(\beta, \mathcal{Z}_{\text{Lip}}(\mathcal{X})) \leq \left( \frac{V}{\beta} \right)^q$$



then using the shorthand  $r = \max\{2, q\}$ , we have that :

$$\forall n \in \mathbb{N}, \mathcal{R}_n^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) \leq \frac{9V \log^{1+\frac{1}{r}}(n)}{n^{1/r}} .$$

Notice that in the above theorem if  $q \geq 2$  then  $r = q$  and so like in Theorem ?? we get a tight relationship between Rademacher complexity and oracle complexity of offline optimization. However this is not always true. In  $d$  dimensional space where  $d$  is small, using the centroid methods once can guarantee oracle complexity upper bound of  $d \log(1/\epsilon)$ . However it turns out that at least for the dual learning problems (when  $\mathcal{H} = \mathcal{X}^*$ ) if the dimension is large enough, by the celebrated Dvoretzky-Roger's theorem (see for instance [87]), we can infer that for all  $\beta < \epsilon$ ,  $\text{fat}_\beta^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X}))$  is larger than order  $1/\beta^2$ . Later on in Section 8.3 as a side interest we can use this to conclude that for dual learning problems in high dimensional problems in appropriate spaces (most common ones), mirror descent is almost optimal even for offline optimization for the convex Lipschitz class!

The above theorem establishes connections between offline convex optimization and statistical convex learning especially for supervised learning problems. The following corollary implies that for any  $\mathcal{H}$  and  $\mathcal{X}$ , if we can find an efficient optimization algorithm for optimizing functions in class  $\mathcal{Z}_{\text{Lip}}(\mathcal{X})$ , then we can find an efficient algorithm for statistical learning of supervised convex learning problem.

**Corollary 103.** *If there exists some  $V > 0$  and  $q \in (0, \infty)$  such that for any  $\beta > 0$ ,*

$$m_{\text{off}}(\beta, \mathcal{Z}_{\text{Lip}}(\mathcal{X})) \leq \left(\frac{V}{\beta}\right)^q$$

*then using the shorthand  $r = \max\{2, q\}$ , we have that there exists a learning algorithm  $\mathbf{A}_{\text{aERM}}$  (that solves the empirical risk minimization problem approximately) which enjoys learning guarantee for the supervised convex learning problem in the statistical framework:*

$$\sup_{\mathcal{D} \in \Delta(\mathcal{Z}_{\text{sup}}(\mathcal{X}))} \mathbb{E}_{S \sim \mathcal{D}^n} \left[ L_{\mathcal{D}}(\mathbf{A}_{\text{aERM}}(S)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right] \leq \frac{18V \log^{1+\frac{1}{r}}(n)}{n^{1/r}} .$$

*Furthermore, the number of oracle calls needed by this algorithm is bounded by  $\frac{n^{1+q/r}}{18^q \log^{q+\frac{q}{r}}(n)}$ , and so number of oracle calls is at most order  $n^2$ , the sample size.*

*Proof.* We provide a proof sketch of this simple corollary. We start by noting that by Proposition 85, learning rate for the empirical risk minimizer is bounded by Rademacher complexity of the associated linear class and by previous theorem we see that upper bound on oracle complexity for optimization of  $\mathcal{Z}_{\text{Lip}}(\mathcal{X})$  implies upper bound on Rademacher complexity. Thus combining these results we conclude the statement of the corollary.  $\square$

First of all, while the above corollary we give for supervised convex learning problem with absolute loss, it can also be extended to logistic loss, hinge losses and basically any convex Lipschitz loss.

### 9.3 Main Result : Optimality of Mirror Descent for Offline Convex Optimization

The following theorem shows that for a large class of convex problems, mirror descent is near optimal even for offline convex optimization in terms of oracle complexity.

**Theorem 104.** *If there exists some constant  $G \geq 1$  such that for any  $1 < p' < r \leq 2$*

$$C_{p'} \leq G C_r^{\text{iid}}$$

*and it is true that for some  $V > 0$  and  $p \in (1, \infty)$ ,*

$$m_{\text{off}}(\epsilon, \mathcal{Z}_{\text{Lip}}(\mathcal{X})) \leq \left(\frac{V}{\epsilon}\right)^q$$

*where  $q = p/(p-1)$ . Then there exists function  $\Psi$  and step size  $\eta$  using which the stochastic mirror descent algorithm enjoys the guarantee*

$$m_{\text{off}}(\epsilon, \overline{\mathcal{A}_{\text{MD}}}, \mathcal{Z}(\mathcal{X})) \leq \left(\frac{54018 G V \log^5\left(\frac{1}{\epsilon}\right)}{\epsilon}\right)^r$$

*Proof.* First step we use Theorem 102 to bound  $\mathcal{R}_n^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{X}))$ . Next we use proceed as in the proof of 86. Next using Corollary 89 we see that for any  $p' < p$ ,  $C_{p'}^{\text{iid}} \leq \frac{60\sqrt{2}V}{p-p'}$ . Using the assumption that  $C_p \leq G C_p^{\text{iid}}$  we get bound on  $C_{p'}$ . However  $C_p$  is an upper bound on  $D_p$  and so using Corollary 62 we conclude the proof.  $\square$

The first thing we notice is that if  $q \geq 2$  then  $r = q$ . Hence in this case the guarantee on the oracle complexity of offline optimization using mirror descent algorithm matches the upper bound on oracle complexity for the problem. Hence we see that for such a case mirror descent algorithm is near optimal even for offline convex optimization. In particular, if for a problem no algorithm can guarantee a better oracle complexity than order  $1/\epsilon^2$  then from the above theorem we can conclude that mirror descent is near optimal (up to polynomials).

Now let us consider the dual problem, that is the case when  $\mathcal{H} = \mathcal{X}^*$ . In this case, the celebrated Dvoretzky Roger's theorem (see [87] for geometric interpretation we use here) implies that for any Banach space of dimension large enough (larger than  $2^n$ ), there exists set of  $n$  points,  $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathcal{X}$  such that for all  $\epsilon_1, \dots, \epsilon_n \in \{\pm 1\}^n$ ,

$$\frac{1}{\sqrt{n}} \leq \left\| \frac{1}{n} \sum_{t=1}^n \epsilon_t \mathbf{x}_t \right\|.$$

Hence we can conclude that for any  $\beta > 0$ , if  $\dim(\mathcal{H}) \rightarrow \infty$ , then

$$\text{fat}_\beta(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{H}^*)) \geq \frac{1}{\beta^2}.$$

Hence by Lemma 101 we can conclude that oracle complexity of offline optimization for the convex Lipschitz class in the dual case when the dimensionality is large enough is lower bounded by  $1/2\beta^2$ . Hence we can conclude that for large dimensional problems, mirror descent is near optimal even for offline optimization problem. Also as per the discussion in sub-section ?? we can also conclude that for statistical learning problem with any powerful oracle  $\overline{\mathcal{O}}$ , mirror descent is near optimal for convex Lipschitz problems in the dual case when dimension is large enough.

The Dvoretzky Roger's theorem tells us very generally that for any large dimensional dual problem, mirror descent is always near optimal. However the dimension of the problem needs to be very large (of order exponential in  $1/\epsilon$  or more) for this claim to be true. While at complete generality we can only ensure this much, for many commonly encountered problems one can provide lower bounds for fat-shattering dimension of order  $1/\epsilon^2$  or worse even when dimensionality is only as large as some polynomial in  $1/\epsilon$ . More specifically, for the case when the pair  $\mathcal{H}/\mathcal{X}$  are either  $\ell_2^d/\ell_2^d$  or  $\ell_1^d/\ell_\infty^d$ , fat-shattering dimension can be lower bound by  $1/\epsilon^2$  as long as dimension is larger than order  $1/\epsilon^2$ . Hence we can conclude that for these cases, when we are dealing with high dimensional convex optimization problems when dimension is as large as order  $1/\epsilon^2$  where  $\epsilon$  is desired sub-optimality, then mirror descent is near optimal even for offline optimization problem. This result can also be extended to  $\ell_p/\ell_q$  pair case when  $p \in [2, \infty)$  and Schatten norm counterparts of all the  $\ell_p$  cases. We would also like to point out that such high dimensional (relative to desired sub-optimality) problems are common in machine learning and high-dimensional statistics application for cases when dimension  $d$  is large than the sample size  $n$ . In all these cases mirror descent is also near optimal for offline convex optimization.

## 9.4 Statistical Learning With Distributed Oracles

In the previous chapter, in section ?? we showed that for most reasonable statistical convex learning problems, stochastic mirror descent is near optimal not only in terms of learning rate but also in terms of number of gradient access (or even any other local oracle information) needed to guarantee excess risk smaller than some target value. Of course in the result we counted accessing gradient at a point of one point sample point as one oracle query. While we abstract off oracle as a black box, in reality the answer that the oracle computes, example gradient, is generally needs to be calculated by the learning algorithm itself. With the growth of cluster and grid computing and its influence in the design of machine learning and optimization algorithms one may wonder if having access to several machines can help speed up learning. For instance one could think of performing gradient descent or mirror descent on training sample but then calculate gradient of empirical average loss in parallel so that computation of in fact several gradients is done more or less in the time of computing one. In this section we try to formalize a more powerful oracle based statistical learning model that can capture such distributed computing scenarios. We show that for most reasonable cases in high dimension for statistical convex learning (of  $\mathcal{Z}_{\text{Lip}}(\mathcal{X})$ ) even having such more powerful oracles do not help and the single pass gradient based mirror descent is near optimal in terms of efficiency even when compared against learning algorithms that have access to these more powerful distributed oracles. In short we show that parallelization does not help for the case when instance set is  $\mathcal{Z}_{\text{Lip}}(\mathcal{X})$ .

More formally, in the oracle-based statistical learning model we consider, learner has access to a local distributed oracle  $\overline{\mathcal{O}} : \mathcal{H} \times \bigcup_{n \in \mathbb{N}} \mathcal{Z}^n \mapsto \mathcal{I}$  that can provide answers when queried at a point about multiple instances simultaneously. Of course here local means the oracle is local w.r.t. the sequence of instances  $z_1, \dots, z_n$  it is queried on at any times. The corresponding learning protocol is given below.

**Powerful Oracle-based Statistical Learning Protocol :**  
 Sample  $S = (z_1, \dots, z_n) \sim \mathcal{D}^n$   
**for**  $t = 1$  **to**  $m$   
     Learner picks hypothesis  $\mathbf{h}_t \in \overline{\mathcal{H}}$   
     Oracle provides answer  $I_t = \overline{\mathcal{O}}(\mathbf{h}_t, S)$   
**end for**

Now to study efficiency of statistical learning algorithms that have access to such distributed oracles, similar to  $m_{\text{off}}$ , the oracle complexity of offline optimization procedures, we introduce oracle complexity  $m_{\text{stat}}$  of oracle-based statistical learning algorithms that have access to distributed oracles.

**Definition 43.** For a given class of convex functions corresponding to instance set  $\mathcal{Z}$  and a distributed Oracle  $\overline{\mathcal{O}}$ , the distributed statistical oracle complexity of a given "Oracle based optimization/learning algorithm",  $\mathbf{A}^{\overline{\mathcal{O}}}$ , is defined as

$$m_{\text{stat}}(\epsilon, \mathbf{A}^{\overline{\mathcal{O}}}, \mathcal{Z}) = \inf \left\{ m \in \mathbb{N} \mid \sup_{\mathcal{D} \in \Delta(\mathcal{Z})} \left\{ L_{\mathcal{D}}(\mathbf{A}^{\overline{\mathcal{O}}}(I_1, \dots, I_m)) - \inf_{\mathbf{h} \in \mathcal{H}} L_{\mathcal{D}}(\mathbf{h}) \right\} \leq \epsilon \right\}$$

Further, for the given distributed oracle  $\overline{\mathcal{O}}$ , the  $\overline{\mathcal{O}}$ -distributed statistical oracle complexity of the given statistical convex learning problem is defined as  $m_{\text{stat}}(\epsilon, \mathcal{Z}, \overline{\mathcal{O}}) = \inf_{\mathbf{A}^{\overline{\mathcal{O}}}} m_{\text{stat}}(\epsilon, \mathbf{A}^{\overline{\mathcal{O}}}, \mathcal{Z})$ .

Further we can define oracle complexity of a distributed oracle-based statistical convex learning problem (independent of any particular oracle) as follows :

**Definition 44.** For a given class of convex functions corresponding to instance set  $\mathcal{Z}$ , the distributed statistical oracle complexity of the given distributed oracle-based statistical convex learning problem is defined as

$$m_{\text{stat}}(\epsilon, \mathcal{Z}) = \inf_{\overline{\mathcal{O}}} m_{\text{stat}}(\epsilon, \mathcal{Z}, \overline{\mathcal{O}}).$$

Now to provide lower bounds on oracle complexity under this more powerful oracle based statistical learning scenario, we start by noticing that lower bounds on oracle complexity of offline optimization problems also provide lower bounds on oracle complexity for learning under the powerful oracle based statistical learning protocol. The reasoning for this is simple : If distribution  $\mathcal{D}$  deterministically picked a single function then the problem is identical to offline optimization in the function class and querying oracle on a sequence or just the one function is exactly the same. The following proposition captures exactly this.

**Proposition 105.** For any  $\mathcal{H}$  and convex learning instance space  $\mathcal{Z}$ , we have that

$$m_{\text{stat}}(\epsilon, \mathcal{Z}) \geq m_{\text{off}}(\epsilon, \mathcal{Z})$$

In the previous section specifically Theorem 104 we showed that for several reasonable, high dimensional offline convex optimization problem over instance space  $\mathcal{Z}_{\text{Lip}}(\mathcal{X})$  mirror descent is in fact near optimal even when we consider distributed oracle-based statistical convex learning algorithm. That is we have the following corollary which is trivial given previous proposition and Theorem 104.

**Corollary 106.** If there exists some constant  $G \geq 1$  such that for any  $1 < p' < r \leq 2$

$$C_{p'} \leq G C_r^{\text{iid}}$$

and it is true that for some  $V > 0$  and  $p \in (1, \infty)$ ,

$$m_{\text{stat}}(\epsilon, \mathcal{Z}_{\text{Lip}}(\mathcal{X})) \leq \left(\frac{V}{\epsilon}\right)^q$$

where  $q = p/(p - 1)$ . Then there exists function  $\Psi$  and step size  $\eta$  using which the stochastic mirror descent algorithm enjoys the guarantee

$$m_{\text{stat}}(\epsilon, \overline{\mathcal{A}_{\text{MD}}}, \mathcal{Z}(\mathcal{X})) \leq \left(\frac{54018 G V \log^5(\frac{1}{\epsilon})}{\epsilon}\right)^r$$

However note that mirror descent is a gradient-based single pass algorithm that at each iteration only queries gradient at a single sample point. Hence we see that for these convex learning problems over instance space  $\mathcal{Z}_{\text{Lip}}(\mathcal{X})$ , stochastic mirror descent is near optimal and that having access to any kind of distributed local oracle does not help. That is in the worst case parallelization does not help.

## 9.5 Detailed Proofs

*Proof of Lemma 101.* The proof is essentially a vderly simple modification of the one provided by Nemirovski and Yudin in Section 4.4.2 of [1]. We provide an abridged version here with the appropriate modifications needed to deal with the non-dual case with a few minor alterations to relate to fat-shattering dimension. To prove the lower bound, we first start by picking  $\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{X}$  add  $s_1, \dots, s_m \in \mathbb{R}$ . Now the functions we shall consider are of form

$$z_\epsilon(\mathbf{h}; (\mathbf{x}_1, s_1), \dots, (\mathbf{x}_m, s_m)) = \max_{i \in [m]} \epsilon_i(\langle \mathbf{h}, -\mathbf{x}_i \rangle + s_i)$$

where  $\epsilon \in \{\pm 1\}^m$ . Notice that each  $z_\epsilon \in \mathcal{Z}_{\text{Lip}}(\mathcal{X})$ . Note also that for any  $\epsilon \in \{\pm 1\}^m$ ,

$$-\inf_{\mathbf{h} \in \mathcal{H}} z_\epsilon(\mathbf{h}; (\mathbf{x}_1, s_1), \dots, (\mathbf{x}_m, s_m)) = -\inf_{\mathbf{h} \in \mathcal{H}} \max_{i \in [m]} \epsilon_i (\langle \mathbf{h}, -\mathbf{x}_i \rangle + s_i) = \sup_{\mathbf{h} \in \mathcal{H}} \min_{i \in [m]} \epsilon_i (\langle \mathbf{h}, \mathbf{x}_i \rangle - s_i) . \quad (9.5.1)$$

Remember that we want to show that for any  $\mathbf{A}^\mathcal{O}$  there exists a function that requires at least  $m$  calls to some Oracle  $\mathcal{O}$  to ensure sub-optimality less than  $\epsilon > 0$ . The first thing we notice is that the family of functions we consider are piece wise linear and so any local oracle can give no more information that function value and gradient at point of query. Now given an Optimization algorithm  $\mathbf{A}^\mathcal{O}$  the exact function we shall use for the lower bound will be constructed in  $m$  steps based on the algorithm and the choosen  $\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{X}$  add  $s_1, \dots, s_m \in \mathbb{R}$ . The procedure for constructing the function is given below :

```

Initialize  $I_1 = [m]$ 
for  $t = 1$  to  $m$ 
     $\mathbf{A}^\mathcal{O}$  picks  $\mathbf{h}_t \in \mathcal{H}$  for query
     $i(t) = \operatorname{argmax}_{i \in I_t} \{|\langle \mathbf{h}_t, -\mathbf{x}_i \rangle + s_i|\}$ 
     $\epsilon_t = \begin{cases} +1 & \text{if } (\langle \mathbf{h}_t, -\mathbf{x}_{i(t)} \rangle + s_{i(t)}) \geq 0 \\ -1 & \text{otherwise} \end{cases}$ 
     $I_{t+1} = I_t \setminus \{i(t)\}$ 
     $z^t(\mathbf{h}) = \max_{j \in [t]} \{\epsilon_j (\langle \mathbf{h}, -\mathbf{x}_{i(j)} \rangle + s_{i(j)})\}$ 
    Return answer to query as  $I_t = \mathcal{O}(\mathbf{h}_t, z^t)$ .
end for

```

The first thing we notice about  $z^m$  is that it is of the form :

$$z^m(\cdot) = z_\epsilon(\cdot; (\mathbf{x}_{i(1)}, s_{i(1)}), \dots, (\mathbf{x}_{i(m)}, s_{i(m)})) \quad (9.5.2)$$

where  $\epsilon_1, \dots, \epsilon_m$  are given by the procedure above. Next,  $z^m$  is such that, for any  $i \in [m]$  and any local oracle  $\mathcal{O}$ ,

$$\mathcal{O}(\mathbf{h}_i, z^m) = \mathcal{O}(\mathbf{h}_i, z^i)$$

hence the  $\mathbf{h}_1, \dots, \mathbf{h}_m$  returned by the algorithm when it is presented with function  $f^m$  is the same as the corresponding ones in the above procedure. Finally, by the way the functions are constructed (specifically the way  $\epsilon_t$  is picked),

$$z^m(\mathbf{h}_m) \geq 0$$

Hence we conclude that

$$\begin{aligned}
z^m(\mathbf{h}_m) - \inf_{\mathbf{h} \in \mathcal{H}} z^m(\mathbf{h}) &\geq - \inf_{\mathbf{h} \in \mathcal{H}} z^m(\mathbf{h}) \\
&= - \inf_{\mathbf{h} \in \mathcal{H}} z_\epsilon(\mathbf{h}; (\mathbf{x}_{i(1)}, s_{i(1)}), \dots, (\mathbf{x}_{i(m)}, s_{i(m)})) \quad (\text{by Eq. 9.5.2}) \\
&= \sup_{\mathbf{h} \in \mathcal{H}} \min_{i \in [m]} \epsilon_i(\langle \mathbf{h}, \mathbf{x}_i \rangle - s_i) \quad (\text{by Eq. 9.5.1}) \\
&\geq \inf_{\epsilon \in \{\pm 1\}^m} \sup_{\mathbf{h} \in \mathcal{H}} \min_{i \in [m]} \epsilon_i(\langle \mathbf{h}, \mathbf{x}_i \rangle - s_i) .
\end{aligned}$$

Furthermore note that the choice of  $\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{X}$  and  $s_1, \dots, s_m$  are arbitrary. Hence we can conclude that for any  $\beta > 0$ , if for any  $m$  there exists  $\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{X}$  and  $s_1, \dots, s_m \in \mathbb{R}$  such that

$$\inf_{\epsilon \in \{\pm 1\}^m} \sup_{\mathbf{h} \in \mathcal{H}} \min_{i \in [m]} \epsilon_i(\langle \mathbf{h}, \mathbf{x}_i \rangle - s_i) > \beta ,$$

then no oracle based algorithm can achieve sub-optimality smaller than  $\beta$  in  $m$  or less steps. However note that this is exactly the definition of fat-shattering dimension at scale  $2\beta$  (Definition 13) for the linear class given by

$$\mathcal{F} = \{\mathbf{x} \mapsto \langle \mathbf{h}, \mathbf{x} \rangle : \mathbf{h} \in \mathcal{H}\} .$$

Hence we conclude the lemma statement. □

*Proof of Theorem 102.* The first inequality is a direct consequence of Lemma 101. For the upper bound on the Rademacher complexity note that, for any  $n \in \mathbb{N}$ , by the refined Dudley integral bound we have that :

$$\mathcal{R}_n^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) \leq \inf_{\alpha > 0} \left\{ 4\alpha + 10 \int_{\alpha}^1 \sqrt{\frac{\text{fat}_{\beta}^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) \log(n)}{n}} d\beta \right\}$$

We now divide the analysis into two cases, first where  $q \in [2, \infty)$  and next where  $q \in (0, 2)$ . We start for the case when  $q \in [2, \infty)$  and see that using the assumption of this theorem and Lemma 101 we see that for any

$q \in [2, \infty)$ :

$$\begin{aligned}
\mathcal{R}_n^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) &\leq \inf_{\alpha > 0} \left\{ 4\alpha + \sqrt{\frac{V^q \log(n)}{n}} 10 \int_{\alpha}^1 \frac{1}{\beta^{\frac{q}{2}}} d\beta \right\} \\
&\leq \inf_{\alpha > 0} \left\{ 4\alpha + \sqrt{\frac{V^q \log(n)}{n}} 10 \int_{\alpha}^1 \frac{(\frac{q}{2} - 1) \log(1/\beta) + 1}{\beta^{\frac{q}{2}}} d\beta \right\} \\
&\leq \inf_{\alpha > 0} \left\{ 4\alpha + 10 \sqrt{\frac{V^q \log(n)}{n}} \left[ \frac{\log(\beta)}{\beta^{\frac{q}{2}-1}} \right]_{\alpha}^1 \right\} \\
&\leq \inf_{\alpha > 0} \left\{ 4\alpha + 10 \sqrt{\frac{V^q \log(n)}{n}} \frac{\log(1/\alpha)}{\alpha^{\frac{q}{2}-1}} \right\} \\
&\leq \frac{9V \log^{1+\frac{1}{q}}(n)}{n^{1/q}}
\end{aligned}$$

where in the last step above we used the value  $\alpha = \frac{V \log^{1/q}(n)}{n^{1/q}}$ .

Now we turn our attention to case when  $q \in (0, 2)$ . As for this case we simply note that  $(\frac{V}{\epsilon})^q \leq (\frac{V}{\epsilon})^2$  and so using the case when  $q = 2$  we conclude that

$$\mathcal{R}_n^{\text{iid}}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})) \leq \frac{9V \log^{3/2}(n)}{n^{1/2}}$$

This concludes the proof. □

## 9.6 Discussion

The key result of this chapter is that for most reasonable cases, if the dimension of the vector space  $\mathcal{B}$  is large enough, then stochastic mirror descent algorithm is near optimal even for offline convex optimization for instance space  $\mathcal{Z}_{\text{Lip}}(\mathcal{X})$ . We further show that for statistical convex learning problems over instance space  $\mathcal{Z}_{\text{Lip}}(\mathcal{X})$  even when we consider learning algorithm that use distributed oracles, (ie. uses distributed computation of local oracle information like gradients etc.) mirror descent is still near optimal and parallelization does not help.



# Chapter 10

## Conclusion and Future Work

In this section we delineate some important questions that are related to the work in this dissertation. We also discuss some further directions of research. Finally we summarize and give some concluding remarks.

### 10.1 Open Problems

#### 10.1.1 Online Optimization and Stability

Another direction yet to be explored is the question of online learnability in the general learning setting. In the statistical paradigm we used the tool of stability and properties of asymptotic empirical minimizer of learning rule to determine learnability for the general setting. We would like to explore the problem of online learnability in the general setting on similar lines.

**Question 1.** Are there properties analogous to stability and AERM property in the online paradigm that guarantee online learnability in the general setting for learning ?

**Question 2.** Can we provide a generic strategy for *Learner* in the online learning framework that guarantees diminishing regret whenever the problem is learnable ?

#### 10.1.2 Upper Bounding Oracle Complexity in Terms of Fat-Shattering Dimension

In the second part of the dissertation especially in Chapter ?? we showed that the oracle complexity of offline convex optimization problem,  $m_{\text{off}}(\epsilon, \mathcal{Z}_{\text{Lip}}(\mathcal{X}))$  is lower bounded by the fat shattering dimension of the associated linear class  $\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X})$ . Using this we showed that at least for supervised learning problems, if one can efficiently optimize convex function corresponding to the class  $\mathcal{Z}_{\text{Lip}}(\mathcal{X})$  then one can also statistically learn and efficiently. In the same chapter we also showed that for most reasonable cases, if dimension is large enough, then mirror descent is near optimal even for offline optimization. Using the results in the thesis one

can also conclude that for these large dimensional cases,  $m_{\text{off}}(\epsilon, \mathcal{Z}_{\text{Lip}}(\mathcal{X}))$  can also be upper bounded by  $\tilde{O}(\text{fat}_\epsilon)$ . Can this result be generalized and can we show the upper bound on oracle complexity in terms of fat shattering dimension always hold? We pose the question.

**Question 3.** Is it always true that

$$m_{\text{off}}(\epsilon, \mathcal{Z}_{\text{Lip}}(\mathcal{X})) \leq \text{fat}_{c\epsilon}(\mathcal{F}_{\text{lin}}(\mathcal{H}, \mathcal{X}))$$

where  $c$  is some universal constants? If it is true, then can one give an optimization algorithm that has oracle complexity bounded by fat-shattering dimension?

## 10.2 Further Directions

The thesis mainly covers the story of learning from the perspective of optimization and answers questions about learnability. However there were a few results that emerged out of results and techniques provided in this thesis and we delineate a few below.

We mainly considered two extreme scenarios while considering statistical and online learning framework. In statistical framework, instances were sampled iid and in the online learning framework instances picked adversarially. It is interesting to consider the intermediate scenarios where learner is not faced with a completely worst case adversary but is also not faced with iid sampling of instances. Maybe adversary might have some constraints on instances that can be chosen or chooses instances in a stochastic way that is more complex than iid sampling. Such a scenario is analyzed in [88] based on techniques in Chapter 4. Another orthogonal way in which the results in the chapter were extended was to games beyond online learning to include games like Blackwell's approachability, calibration etc. in [89, 90].

Results in chapters ?? and ?? influenced and shaped the work in paper [91] where we showed how one can make appropriate changes to stochastic mirror descent and accelerated methods to include mini-batching (breaking sample into blocks and instead of updating in each step with single gradient update with average of the block of gradients). We showed that this helped in guaranteeing better time complexity with parallelization of these methods.

## 10.3 Summary

An important question in the field of theoretical machine learning is that of learnability and learning rates. We have explored this question for various learning problems in both statistical and online learning frameworks. In the statistical learning framework we provide the first general characterization of learnability in the general setting using the notion of stability of learning algorithms. We also provided a generic algorithm for learning in the statistical learning framework. As for the problem of learnability in the online framework while we don't yet have a complete picture we introduced various complexity measures analogous to the

ones in statistical learning framework. We also provide characterization of online learnability for real valued supervised learning problem.

An integral part of machine learning is optimization. While the question of learnability and learning rates are central to machine learning theory, from a practical point of view one would like to consider problems that are efficiently learnable. To address this issue in a general way, we considered convex learning and optimization problems in both statistical and online learning framework. We used the notion of oracle complexity to address issue of efficiency. For the online learning problems, we showed mirror descent is universal and near optimal. That is whenever a convex problem is online learnable, it is learnable with near optimal rates using mirror descent. Since mirror descent is a first order method (sub-gradient based) we could infer that for online learning scenario mirror descent is near optimal in terms of both rates and oracle complexity. We also explored connections between learning in the various frameworks and oracle based optimization. For the statistical convex learning problem, unlike online setting, in general it is not true that mirror descent is universal. However we saw that for problems we would encounter in practical applications though, this was in fact the case. Mirror descent would indeed be near optimal. We also saw that for certain offline optimization problems in high enough dimensions, mirror descent can again shown to be near optimal.

We expect the work to provide a better understanding of learning algorithms especially from the perspective of optimization. While it is common that for machine learning practitioners optimization is often an after thought and is in a sense mainly a computational issue, through this work we would like to stress that learning can be seen as optimization and should in fact be seen as so. On the other hand, we also show some strong connections between optimization and showed how tools from learning theory can be used to prove results on optimization. Hence we would also like to stress overall the strong and inevitable connections between the two.

In this work we also used several concepts from the theory of Banach space geometry. It would certainly be interesting to see if more connections can be made and techniques from Banach space geometry be used to prove more results about learning and optimization.

# Bibliography

- [1] A. Nemirovski and D. Yudin. *Problem complexity and method efficiency in optimization*. Nauka Publishers, Moscow, 1978.
- [2] S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan. Stochastic convex optimization. In *Conference on Learning Theory*, 2009.
- [3] S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan. Learnability and stability in the general learning setting. In *Proceedings of the 22nd Annual Conference on Computational Learning Theory*, 2009. .
- [4] Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. Learnability, stability and uniform convergence. *J. Mach. Learn. Res.*, 11:2635–2670, December 2010. ISSN 1532-4435. .
- [5] A. Rakhlin, K. Sridharan, and A. Tewari. Online learning: Random averages, combinatorial parameters, and learnability. *Neural Information Processing Systems (NIPS)*, 2010. .
- [6] Nathan Srebro, Karthik Sridharan, and Ambuj Tewari. On the universality of online mirror descent. *Arxiv.*, 2011. .
- [7] K. Sridharan and A. Tewari. Convex games in Banach spaces. In *Proceedings of the 23rd Annual Conference on Learning Theory*, 2010.
- [8] V.N. Vapnik. *The Nature of Statistical Learning Theory*. Springer, 1995.
- [9] Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Learnability and the Vapnik-Chervonenkis dimension. *Journal of the Association for Computing Machinery*, 36(4):929–965, October 1989.
- [10] N. Alon, S. Ben-David, N. Cesa-Bianchi, and D. Haussler. Scale-sensitive dimensions, uniform convergence, and learnability. *Journal of the ACM*, 44(4):615–631, 1997.
- [11] M. J. Kearns, R. E. Schapire, and L. M. Sellie. Toward efficient agnostic learning. In *Proceedings of the Fifth Annual ACM Workshop on Computational Learning Theory*, pages 341–352, July 1992. To appear, *Machine Learning*.

- [12] D. Haussler. Decision theoretic generalizations of the PAC model for neural net and other learning applications. *Information and Computation*, 100(1):78–150, 1992.
- [13] V. N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its applications*, XVI(2):264–280, 1971.
- [14] Martin Anthony and Peter Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, 1999.
- [15] V. N. Vapnik. *Statistical Learning Theory*. Wiley, 1998.
- [16] David Pollard. *Convergence of Stochastic Processes*. Springer-Verlag, 1984.
- [17] R. M. Dudley. The sizes of compact subsets of Hilbert space and continuity of Gaussian processes. *Journal of Functional Analysis*, 1(3):290–330, 1967.
- [18] J. Hadamard. Sur les problèmes aux dérivées partielles et leur signification physique. *Princeton University Bulletin*, 13:49–52, 1902.
- [19] A. N. Tikhonov. On the stability of inverse problems. *Dokl. Akad. Nauk SSSR*, 39(5):195–198, 1943.
- [20] D. L. Phillips. A technique for the numerical solution of certain integral equations of the first kind. *Journal of the ACM*, 9(1):84–97, 1962.
- [21] W. Rogers and T. Wagner. A finite sample distribution-free performance bound for local discrimination rules. *Annals of Statistics*, 6(3):506–514, 1978.
- [22] L. Devroye, L. Györfi, and G. Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer, 1996.
- [23] Leo Breiman. Bias, variance, and arcing classifiers. Technical Report 460, Statistics Department, University of California at Berkeley, 1996.
- [24] M. Kearns and D. Ron. Algorithmic stability and sanity-check bounds for leave-one-out cross-validation. *Neural Computation*, 11(6):1427–1453, 1999.
- [25] Olivier Bousquet and André Elisseeff. Stability and generalization. *J. Mach. Learn. Res.*, 2:499–526, 2002. ISSN 1532-4435.
- [26] S. Kutin and P. Niyogi. Almost-everywhere algorithmic stability and generalization error. In *Proceedings of the 18th Conference in Uncertainty in Artificial Intelligence*, pages 275–282, 2002.
- [27] S. Rakhlin, S. Mukherjee, and T. Poggio. Stability results in learning theory. *Analysis and Applications*, 3(4):397–419, 2005.
- [28] S. Mukherjee, P. Niyogi, T. Poggio, and R. Rifkin. Learning theory: stability is sufficient for generalization and necessary and sufficient for consistency of empirical risk minimization. *Advances in Computational Mathematics*, 25(1-3):161–193, 2006.

- [29] K. Sridharan, N. Srebro, and S. Shalev-Shwartz. Fast rates for regularized objectives. In *Advances in Neural Information Processing Systems 22*, 2008.
- [30] M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *ICML*, pages 928–936, 2003.
- [31] S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan. Learnability and stability in the general learning setting. In *Proceedings of the 22nd Annual Conference on Computational Learning Theory*, 2009. .
- [32] D.A. McAllester and R.E. Schapire. On the convergence rate of good-turing estimators. In *Proceedings of the Thirteenth Annual Conference on Computational Learning Theory*, 2000.
- [33] Amit Daniely, Sivan Sabato, Shai Ben-David, and Shai Shalev-Shwartz. Multiclass learnability and the erm principle. *COLT*, 2011.
- [34] N. Littlestone. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine Learning*, 2(4):285–318, 04 1988.
- [35] S. Ben-David, D. Pal, and S. Shalev-Shwartz. Agnostic online learning. In *Proceedings of the 22th Annual Conference on Learning Theory*, 2009.
- [36] J. Abernethy, A. Agarwal, P. Bartlett, and A. Rakhlin. A stochastic view of optimal regret through minimax duality. In *Proceedings of the 22nd Annual Conference on Learning Theory*, 2009.
- [37] V. Koltchinskii and D. Panchenko. Rademacher processes and bounding the risk of function learning. *High Dimensional Probability II*, 47:443–459, 2000.
- [38] P. L. Bartlett and S. Mendelson. Rademacher and gaussian complexities: risk bounds and structural results. *J. Mach. Learn. Res.*, 3:463–482, 2003. ISSN 1532-4435.
- [39] S. Mendelson. A few notes on statistical learning theory. In S. Mendelson and A. J. Smola, editors, *Advanced Lectures in Machine Learning, LNCS 2600, Machine Learning Summer School 2002, Canberra, Australia, February 11-22*, pages 1–40. Springer, 2003.
- [40] R. M. Dudley. *Uniform Central Limit Theorems*. Cambridge University Press, 1999.
- [41] M. Ledoux and M. Talagrand. *Probability in Banach Spaces*. Springer-Verlag, New York, 1991.
- [42] V. N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.
- [43] N. Alon, S. Ben-David, N. Cesa-Bianchi, and D. Haussler. Scale-sensitive dimensions, uniform convergence, and learnability. *Journal of the ACM*, 44:615–631, 1997.
- [44] P. L. Bartlett, P. M. Long, and R. C. Williamson. Fat-shattering and the learnability of real-valued functions. *Journal of Computer and System Sciences*, 52(3):434–452, 1996. (special issue on COLT’94).

- [45] N. Sauer. On the density of families of sets. *J. Combinatorial Theory*, 13:145–147, 1972.
- [46] S. Shelah. A combinatorial problem: Stability and order for models and theories in infinitary languages. *Pac. J. Math*, 4:247–261, 1972.
- [47] N. Alon and J. Spencer. *The Probabilistic Method*. John Wiley & Sons, 2nd edition, 2000.
- [48] S. Mendelson and R. Vershynin. Entropy and the combinatorial dimension. *Inventiones mathematicae*, 152:37–55, 2003.
- [49] E. Giné and J. Zinn. Some limit theorems for empirical processes. *Annals of Probability*, 12(4):929–989, 1984.
- [50] S.A. van de Geer. *Empirical Processes in M-Estimation*. Cambridge University Press, 2000.
- [51] P. Massart. Some applications of concentration inequalities to statistics. *Annales de la Faculté des Sciences de Toulouse*, IX(2):245–303, 2000.
- [52] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [53] R.E. Schapire, Y. Freund, P. Bartlett, and W.S. Lee. Boosting the margin: A new explanation for the effectiveness of voting methods. *The Annals of Statistics*, pages 322–330, 1997.
- [54] V. Koltchinskii and D. Panchenko. Empirical margin distributions and bounding the generalization error of combined classifiers. *Annals of Statistics*, 30(1):1–50, 2002.
- [55] S. M. Kakade and A. Kalai. From batch to transductive online learning. In *NIPS*, 2005.
- [56] N. Cesa-Bianchi and G. Lugosi. On prediction of individual sequences. *Annals of Statistics*, pages 1865–1895, 1999.
- [57] A. Tauman Kalai and R. Sastry. The isotron algorithm: High-dimensional isotonic regression. In *Proceedings of the 22th Annual Conference on Learning Theory*, 2009.
- [58] A.N. Kolmogorov and V.M. Tikhomirov.  $\varepsilon$ -entropy and  $\varepsilon$ -capacity of sets in function spaces. *Uspekhi Matematicheskikh Nauk*, 14(2):3–86, 1959.
- [59] P.L. Bartlett, O. Bousquet, and S. Mendelson. Local rademacher complexities. *Annals of Statistics*, 33(4):1497–1537, 2005.
- [60] Nathan Srebro, Karthik Sridharan, and Ambuj Tewari. Smoothness, low-noise and fast rates. *NIPS*, 2010.
- [61] S. Shalev-Shwartz. *Online Learning: Theory, Algorithms, and Applications*. PhD thesis, Hebrew University of Jerusalem, 2007.
- [62] N. Cesa-Bianchi, A. Conconi, and C. Gentile. On the generalization ability of on-line learning algorithms. *IEEE Trans. on Information Theory*, 50(9):2050–2057, September 2004.

- [63] E. Hazan, A. Kalai, S. Kale, and A. Agarwal. Logarithmic regret algorithms for online convex optimization. In *Proceedings of the Nineteenth Annual Conference on Computational Learning Theory*, 2006.
- [64] S. Shalev-Shwartz and Y. Singer. Logarithmic regret algorithms for strictly convex repeated games. Technical report, The Hebrew University, 2007. Available at <http://www.cs.huji.ac.il/~shais>.
- [65] G. Pisier. Martingales with values in uniformly convex spaces. *Israel Journal of Mathematics*, 20(3–4): 326–350, 1975.
- [66] G. Pisier. Martingales in banach spaces (in connection with type and cotype). *Winter School/IHP Graduate Course*, 2011.
- [67] Keith Ball, Eric A. Carlen, and Elliott H. Lieb. Sharp uniform convexity and smoothness inequalities for trace norms. *Invent. Math.*, 115:463–482, 1994.
- [68] Alekh Agarwal, Peter L. Bartlett, Pradeep Ravikumar, and Martin J. Wainwright. Information-theoretic lower bounds on the oracle complexity of convex optimization.
- [69] Nathan Srebro, Jason D. M. Rennie, and Tommi S. Jaakola. Maximum-margin matrix factorization. In *Advances in Neural Information Processing Systems 17*, pages 1329–1336. MIT Press, 2005.
- [70] Nathan Srebro and Adi Shraibman. Rank, trace-norm and max-norm. In *Proceedings of the 18th Annual Conference on Learning Theory*, pages 545–560. Springer-Verlag, 2005.
- [71] Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society, Series B*, 67:301–320, 2005.
- [72] V. Chandrasekaran, S. Sanghavi, P. Parrilo, and A. Willsky. Sparse and low-rank matrix decompositions. In *IFAC Symposium on System Identification*, 2009.
- [73] Ali Jalali, Pradeep Ravikumar, Sujay Sanghavi, and Chao Ruan. A Dirty Model for Multi-task Learning. In *NIPS*, December 2010.
- [74] Bernard Maurey. Type, cotype and k-convexity. In *HANDBOOK OF THE GEOMETRY OF BANACH SPACES. VOLUME 2*, pages 1299–1332. North-Holland, 2003.
- [75] Joram Lindenstrauss and Lior Tzafriri. *Classical Banach Spaces I and II*. Classics in Mathematics. Springer, Berlin, Heidelberg, New York, 1996. ISBN 3-540-60628-9.
- [76] J.L. Krivine. Theoremes de factorisation dans les espaces reticules. *Seminaire MAurey-Schwartz 1973-74: Exp. Nos. 22 and 23, Ecole Polytech, Paris*, 1974.
- [77] Bernard Maurey and Gilles Pisier. Series de variables aleatoires vectorielles independantes et proprietes geometriques des espaces de banach. *Studia Math.*, 58:45–90, 1976.
- [78] D. J. H. Garling. Random martingale transform inequalities. *Probability in Banach spaces 6 (Sandbjerg, 1986)*, (20):pp. 101–119, 1986.



- [79] Sonja Cox and Mark Veraar. Vector-valued decoupling and the burkholder-davis gundy inequality. Technical Report arXiv:1107.2218, Jul 2011.
- [80] Sonja Cox and Mark Veraar. Some remarks on tangent martingale difference sequences in  $\ell_1$ -spaces. Technical Report 12, 2007.
- [81] Paweł Hitczenko and Stephen J. Montgomery-Smith. On a domination of sums of random variables by sums of conditionally independent ones. *Mathematical Proceedings of the Cambridge Philosophical Society*, (119):pp. 91–101, 1996.
- [82] Paweł Hitczenko. On a domination of sums of random variables by sums of conditionally independent ones. *The Annals of Probability*, 22(1):pp. 453–468, 1994. ISSN 00911798.
- [83] Paweł Hitczenko. Comparison of moments for tangent sequences of random variables. *Probability Theory and Related Fields*, 78:223–230, 1988. ISSN 0178-8051.
- [84] Nicol Cesa-Bianchi and Ohad Shamir. Efficient online learning via randomized rounding. *Neural Information Processing System*, 2011.
- [85] J.-P. Kahane. *Some random series of functions*. 2nd edition, 1985.
- [86] D. L. Burkholder. Martingale transforms. *The Annals of Mathematical Statistics*, 37(6):pp. 1494–1504, 1966. ISSN 00034851. URL .
- [87] William B. Johnson and Joram Lindenstrauss. Chapter 1 basic concepts in the geometry of banach spaces. volume 1 of *Handbook of the Geometry of Banach Spaces*, pages 1 – 84. Elsevier Science B.V., 2001. doi: 10.1016/S1874-5849(01)80003-6.
- [88] Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Online learning: Stochastic and constrained adversaries. *NIPS*, 2011.
- [89] Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Online learning: Beyond regret. *COLT*, 2011.
- [90] Dean Foster, Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Complexity-based approach to calibration with checking rules. *COLT*, 2011.
- [91] Andrew Cotter, Nathan Srebro, Ohad Shamir, and Karthik Sridharan. Better mini-batch algorithms via accelerated gradient methods. *NIPS*, 2011.
- [92] P. L. Bartlett and S. Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *JMLR*, 3:463–482, 2002.
- [93] P. Massart. Some applications of concentration inequalities to statistics. *Annales de la Faculté des Sciences de Toulouse*, IX(2):245–303, 2000.

- [94] N. Alon, S. Ben-David, N. Cesa-Bianchi, and D. Haussler. Scale-sensitive dimensions, uniform convergence, and learnability. *FOCS*, 0:292–301, 1993. doi: <http://doi.ieeecomputersociety.org/10.1109/SFCS.1993.366858>.
- [95] Shahar Mendelson. Rademacher averages and phase transitions in glivenko-cantelli classes. *IEEE Trans. On Information Theory*, 48(1):251–263, 2002.

## Appendix A

# Relating Various Complexity Measures : Statistical Learning

### A.1 The Refined Dudley Integral: Bounding Rademacher Complexity with $L_2$ Covering Numbers

We shall find it simpler here to use the empirical Rademacher complexity for a given sample  $x_1, \dots, x_n$  [92]:

$$\widehat{\mathcal{R}}_n(\mathcal{H}) = \mathbb{E}_{\sigma \sim \text{Unif}(\{\pm 1\}^n)} \left[ \sup_{h \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n h(x_i) \sigma_i \right] \quad (\text{A.1.1})$$

and the  $L_2$  covering number at scale  $\epsilon > 0$  specific to a sample  $x_1, \dots, x_n$ , denoted by  $N_2(\epsilon, \mathcal{F}, (x_1, \dots, x_n))$  as the size of a minimal cover  $\mathcal{C}_\epsilon$  such that

$$\forall f \in \mathcal{F}, \exists f_\epsilon \in \mathcal{C}_\epsilon \text{ s.t. } \sqrt{\frac{1}{n} \sum_{i=1}^n (f(z_i) - f_\epsilon(z_i))^2} \leq \epsilon.$$

We will also denote  $\widehat{\mathbb{E}}[f^2] = \frac{1}{n} \sum_{i=1}^n f^2(x_i)$ .

We state our bound in terms of the empirical Rademacher complexity and covering numbers. Taking a supremum over samples of size  $n$ , we get the same relationship between the worst-case Rademacher complexity and covering numbers, as is used in Section ??.

**Lemma 107.** *For any function class  $\mathcal{F}$  containing functions  $f : \mathcal{X} \mapsto \mathbb{R}$ , we have that*

$$\widehat{\mathcal{R}}_n(\mathcal{F}) \leq \inf_{\alpha \geq 0} \left\{ 4\alpha + 10 \int_{\alpha}^{\sup_{f \in \mathcal{F}} \sqrt{\widehat{\mathbb{E}}[f^2]}} \sqrt{\frac{\log \mathcal{N}(\epsilon, \mathcal{F}, (x_1, \dots, x_n))}{n}} d\epsilon \right\}.$$

*Proof.* Let  $\beta_0 = \sup_{f \in \mathcal{F}} \sqrt{\mathbb{E}[f^2]}$  and for any  $j \in \mathbb{Z}_+$  let  $\beta_j = 2^{-j} \sup_{f \in \mathcal{F}} \sqrt{\mathbb{E}[f^2]}$ . The basic trick here is the idea of chaining. For each  $j$  let  $T_j$  be a (proper)  $L_2$ -cover at scale  $\beta_j$  of  $\mathcal{F}$  for the given sample. For each  $f \in \mathcal{F}$  and  $j$ , pick an  $\hat{f}_j \in T_j$  such that  $\hat{f}_j$  is a  $\beta_j$  approximation of  $f$ . Now for any  $N$ , we express  $f$  by chaining as

$$f = f - \hat{f}_N + \sum_{i=1}^N (\hat{f}_i - \hat{f}_{i-1})$$

where  $\hat{f}_0 = 0$ . Hence for any  $N$  we have that

$$\begin{aligned} \widehat{\mathcal{R}}_n(\mathcal{F}) &= \frac{1}{n} \mathbb{E}_\sigma \left[ \sup_{f \in \mathcal{F}} \sum_{i=1}^n \sigma_i \left( f(\mathbf{x}_i) - \hat{f}_N(\mathbf{x}_i) + \sum_{j=1}^N (\hat{f}_j(\mathbf{x}_i) - \hat{f}_{j-1}(\mathbf{x}_i)) \right) \right] \\ &\leq \frac{1}{n} \mathbb{E}_\sigma \left[ \sup_{f \in \mathcal{F}} \sum_{i=1}^n \sigma_i (f(\mathbf{x}_i) - \hat{f}_N(\mathbf{x}_i)) \right] + \sum_{j=1}^N \frac{1}{n} \mathbb{E}_\sigma \left[ \sup_{f \in \mathcal{F}} \sum_{i=1}^n \sigma_i (\hat{f}_j(\mathbf{x}_i) - \hat{f}_{j-1}(\mathbf{x}_i)) \right] \\ &\leq \frac{1}{n} \sqrt{\sum_{i=1}^n \sigma_i^2} \sup_{f \in \mathcal{F}} \sqrt{\sum_{i=1}^n (f(x_i) - \hat{f}_N(x_i))^2} + \sum_{j=1}^N \frac{1}{n} \mathbb{E}_\sigma \left[ \sup_{f \in \mathcal{F}} \sum_{i=1}^n \sigma_i (\hat{f}_j(\mathbf{x}_i) - \hat{f}_{j-1}(\mathbf{x}_i)) \right] \\ &\leq \beta_N + \sum_{j=1}^N \frac{1}{n} \mathbb{E}_\sigma \left[ \sup_{f \in \mathcal{F}} \sum_{i=1}^n \sigma_i (\hat{f}_j(\mathbf{x}_i) - \hat{f}_{j-1}(\mathbf{x}_i)) \right] \end{aligned} \quad (\text{A.1.2})$$

where the step before last is due to Cauchy-Schwarz inequality and  $\sigma = [\sigma_1, \dots, \sigma_n]^\top$ . Now note that

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n (\hat{f}_j(x_i) - \hat{f}_{j-1}(x_i))^2 &= \frac{1}{n} \sum_{i=1}^n \left( (\hat{f}_j(x_i)) - f(x_i) + (f(x_i) - \hat{f}_{j-1}(x_i)) \right)^2 \\ &\leq \frac{2}{n} \sum_{i=1}^n (\hat{f}_j(x_i) - f(x_i))^2 + \frac{2}{n} \sum_{i=1}^n (f(x_i) - \hat{f}_{j-1}(x_i))^2 \\ &\leq 2\beta_j^2 + 2\beta_{j-1}^2 = 6\beta_j^2. \end{aligned}$$

Now Massart's finite class lemma [93] states that if for any function class  $\mathcal{G}$ ,  $\sup_{g \in \mathcal{G}} \sqrt{\frac{1}{n} \sum_{i=1}^n g(x_i)^2} \leq R$ , then  $\widehat{\mathcal{R}}_n(\mathcal{G}) \leq \sqrt{\frac{2R^2 \log(|\mathcal{G}|)}{n}}$ . Applying this to function classes  $\{f - f' : f \in T_j, f' \in T_{j-1}\}$  (for each  $j$ )

we get from Eq. (A.1.2) that for any  $N$ ,

$$\begin{aligned}
\widehat{\mathcal{R}}_n(\mathcal{F}) &\leq \beta_N + \sum_{j=1}^N \beta_j \sqrt{\frac{12 \log(|T_j| |T_{j-1}|)}{n}} \\
&\leq \beta_N + \sum_{j=1}^N \beta_j \sqrt{\frac{24 \log |T_j|}{n}} \\
&\leq \beta_N + 10 \sum_{j=1}^N (\beta_j - \beta_{j+1}) \sqrt{\frac{\log |T_j|}{n}} \\
&\leq \beta_N + 10 \sum_{j=1}^N (\beta_j - \beta_{j+1}) \sqrt{\frac{\log \mathcal{N}(\beta_j, \mathcal{F}, (x_1, \dots, x_n))}{n}} \\
&\leq \beta_N + 10 \int_{\beta_{N+1}}^{\beta_0} \sqrt{\frac{\log \mathcal{N}(\epsilon, \mathcal{F}, (x_1, \dots, x_n))}{n}} d\epsilon
\end{aligned}$$

where the third step is because  $2(\beta_j - \beta_{j+1}) = \beta_j$  and we bounded  $\sqrt{24}$  by 5. Now for any  $\alpha > 0$ , pick  $N = \sup\{j : \beta_j > 2\alpha\}$ . In this case we see that by our choice of  $N$ ,  $\beta_{N+1} \leq 2\alpha$  and so  $\beta_N = 2\beta_{N+1} \leq 4\epsilon$ . Also note that since  $\beta_N > 2\alpha$ ,  $\beta_{N+1} = \frac{\beta_N}{2} > \alpha$ . Hence we conclude that

$$\widehat{\mathcal{R}}_n(\mathcal{F}) \leq 4\alpha + 10 \int_{\alpha}^{\sup_{f \in \mathcal{F}} \sqrt{\mathbb{E}[f^2]}} \sqrt{\frac{\log \mathcal{N}(\epsilon, \mathcal{F}, (x_1, \dots, x_n))}{n}} d\epsilon.$$

Since the choice of  $\alpha$  was arbitrary we take an infimum over  $\alpha$ . □

## A.2 Bounding $L_\infty$ covering number by Fat-shattering Dimension

The following proposition and lemma are standard in statistical learning theory and their proof can be found for instance in [94]. We provide the statement and the proof of the Lemma for completeness and so that we can state it in the exact form it is used in, in this work.

**Proposition 108.** *Let  $\mathcal{H} \subseteq \{0, \dots, k\}^{\mathcal{X}}$  be a class of functions with  $\text{fat}_2 = d$ . Then, we have,*

$$\mathcal{N}_\infty(1/2, \mathcal{H}, n) \leq \sum_{i=0}^d \binom{n}{i} k^i$$

and specifically for  $n \geq d$  this gives,

$$\mathcal{N}_\infty(1/2, \mathcal{H}, n) \leq \left(\frac{ekn}{d}\right)^d.$$

**Lemma 109.** For any function class  $\mathcal{H}$  bounded by  $B$  and any  $\alpha > 0$  such that  $\text{fat}_\alpha < n$ , we have,

$$\mathcal{N}_\infty(\alpha, \mathcal{H}, n) \leq \left( \frac{2eBn}{\alpha \text{fat}_\alpha(\mathcal{H})} \right)^{\text{fat}_\alpha(\mathcal{H})}.$$

*Proof.* For any  $\alpha > 0$ , define an  $\alpha$ -discretization of the  $[-B, B]$  interval as  $B_\alpha = \{-B + \alpha/2, -B + 3\alpha/2, \dots, -B + (2k+1)\alpha/2, \dots\}$  for  $0 \leq k$  and  $(2k+1)\alpha \leq 4B$ . Also for any  $a \in [-B, B]$ , define  $[a]_\alpha = \underset{r \in B_\alpha}{\text{argmin}} |r - a|$  with ties being broken by choosing the smaller discretization point. For a function  $h : \mathcal{X} \mapsto [-B, B]$  let the function  $[h]_\alpha$  be defined pointwise as  $[h(x)]_\alpha$ , and let  $[\mathcal{H}]_\alpha = \{[h]_\alpha : h \in \mathcal{H}\}$ . First, we prove that  $\mathcal{N}_\infty(\alpha, \mathcal{H}, \{x_i\}_{i=1}^n) \leq \mathcal{N}_\infty(\alpha/2, [\mathcal{H}]_\alpha, \{x_i\}_{i=1}^n)$ . Indeed, suppose the set  $V$  is a minimal  $\alpha/2$ -cover of  $[\mathcal{H}]_\alpha$  on  $\{x_i\}_{i=1}^n$ . That is,

$$\forall h_\alpha \in [\mathcal{H}]_\alpha, \exists \mathbf{v} \in V \text{ s.t. } |v_i - h_\alpha(x_i)| \leq \alpha/2.$$

Pick any  $h \in \mathcal{H}$  and let  $h_\alpha = [h]_\alpha$ . Then  $\|h - h_\alpha\|_\infty \leq \alpha/2$  and for any  $i \in [n]$

$$|h(x_i) - v_i| \leq |h(x_i) - h_\alpha(x_i)| + |h_\alpha(x_i) - v_i| \leq \alpha,$$

and so  $V$  also provides an  $L_\infty$  cover at scale  $\alpha$ .

We conclude that  $\mathcal{N}_\infty(\alpha, \mathcal{H}, \{x_i\}_{i=1}^n) \leq \mathcal{N}_\infty(\alpha/2, [\mathcal{H}]_\alpha, \{x_i\}_{i=1}^n) = \mathcal{N}_\infty(1/2, \mathcal{G}, \{x_i\}_{i=1}^n)$  where  $\mathcal{G} = \frac{1}{\alpha} [\mathcal{H}]_\alpha$ . The functions of  $\mathcal{G}$  take on a discrete set of at most  $\lfloor 2B/\alpha \rfloor + 1$  values. Obviously, by adding a constant to all the functions in  $\mathcal{G}$ , we can make the set of values to be  $\{0, \dots, \lfloor 2B/\alpha \rfloor\}$ . We now apply Proposition 108 with an upper bound  $\sum_{i=0}^d \binom{n}{i} k^i \leq \left(\frac{ekn}{d}\right)^d$  which holds for any  $n > d$ . This yields  $\mathcal{N}_\infty(1/2, \mathcal{G}, \{x_i\}_{i=1}^n) \leq \left(\frac{2eBn}{\alpha \text{fat}_2(\mathcal{G})}\right)^{\text{fat}_2(\mathcal{G})}$ .

It remains to prove  $\text{fat}_2(\mathcal{G}) \leq \text{fat}_\alpha(\mathcal{H})$ , or, equivalently (by scaling)  $\text{fat}_{2\alpha}([\mathcal{H}]_\alpha) \leq \text{fat}_\alpha(\mathcal{H})$ . To this end, suppose there exists a set  $\{x_{i=1}^n\}$  of size  $d = \text{fat}_{2\alpha}([\mathcal{H}]_\alpha)$  such that there is an witness  $s_1, \dots, s_n$  with

$$\forall \epsilon \in \{\pm 1\}^d, \exists h_\alpha \in [\mathcal{H}]_\alpha \text{ s.t. } \forall i \in [d], \epsilon_i(h_\alpha(x_i) - s_i) \geq \alpha.$$

Using the fact that for any  $h \in \mathcal{H}$  and  $h_\alpha = [h]_\alpha$  we have  $\|h - h_\alpha\|_\infty \leq \alpha/2$ , it follows that

$$\forall \epsilon \in \{\pm 1\}^d, \exists h \in \mathcal{H} \text{ s.t. } \forall i \in [d], \epsilon_i(h(x_i) - s_i) \geq \alpha/2.$$

That is,  $s_1, \dots, s_n$  is a witness to  $\alpha$ -shattering by  $\mathcal{H}$ . Thus for any  $\{x_i\}_{i=1}^n$ , as long as  $n > \text{fat}_\alpha$

$$\mathcal{N}_\infty(\alpha, \mathcal{H}, \{x_i\}_{i=1}^n) \leq \mathcal{N}_\infty(\alpha/2, [\mathcal{H}]_\alpha, \{x_i\}_{i=1}^n) \leq \left(\frac{2eBn}{\alpha \text{fat}_{2\alpha}([\mathcal{H}]_\alpha)}\right)^{\text{fat}_{2\alpha}([\mathcal{H}]_\alpha)} \leq \left(\frac{2eBn}{\alpha \text{fat}_\alpha}\right)^{\text{fat}_\alpha(\mathcal{H})}.$$

□

### A.3 Relating Fat-shattering Dimension and Rademacher complexity

The following lemma upper bounds the fat-shattering dimension at scale  $\epsilon \geq \mathcal{R}_n(\mathcal{H})$  in terms of the Rademacher complexity of the function class. The proof closely follows the arguments of Mendelson [95, discussion after Definition 4.2].

**Lemma 110.** *For any hypothesis class  $\mathcal{H}$ , any sample size  $n$  and any  $\epsilon > \mathcal{R}_n(\mathcal{H})$  we have that*

$$\text{fat}_\epsilon(\mathcal{H}) \leq \frac{4n \mathcal{R}_n(\mathcal{H})^2}{\epsilon^2}.$$

In particular, if  $\mathcal{R}_n(\mathcal{H}) = \sqrt{R/n}$  (the typical case), then  $\text{fat}_\epsilon(\mathcal{H}) \leq R/\epsilon^2$ .

*Proof.* Consider any  $\epsilon \geq \mathcal{R}_n(\mathcal{H})$ . Let  $x_1^*, \dots, x_{\text{fat}_\epsilon}^*$  be the set of  $\text{fat}_\epsilon$  shattered points. This means that there exists  $s_1, \dots, s_{\text{fat}_\epsilon}$  such that for any  $J \subset [\text{fat}_\epsilon]$  there exists  $h_J \in \mathcal{H}$  such that  $\forall i \in J, h_J(x_i) \geq s_i + \epsilon$  and  $\forall i \notin J, h_J(x_i) \leq s_i - \epsilon$ . Now consider a sample  $x_1, \dots, x_{n'}$  of size  $n' = \lceil \frac{n}{\text{fat}_\epsilon} \rceil \text{fat}_\epsilon$ , obtained by taking each  $x_i^*$  and repeating it  $\lceil \frac{n}{\text{fat}_\epsilon} \rceil$  times, i.e.  $x_i = x_{\lfloor \frac{i}{\text{fat}_\epsilon} \rfloor}^*$ . Now, following Mendelson's arguments:

$$\begin{aligned} \mathcal{R}_{n'}(\mathcal{H}) &\geq \mathbb{E}_{\sigma \sim \text{Unif}\{\pm 1\}^{n'}} \left[ \frac{1}{n'} \sup_{h \in \mathcal{H}} \left| \sum_{i=1}^{n'} \sigma_i h(x_i) \right| \right] \\ &\geq \frac{1}{2} \mathbb{E}_{\sigma \sim \text{Unif}\{\pm 1\}^{n'}} \left[ \frac{1}{n'} \sup_{h, h' \in \mathcal{H}} \left| \sum_{i=1}^{n'} \sigma_i (h(x_i) - h'(x_i)) \right| \right] \quad (\text{triangle inequality}) \\ &= \frac{1}{2} \mathbb{E}_{\sigma \sim \text{Unif}\{\pm 1\}^{n'}} \left[ \frac{1}{n'} \sup_{h, h' \in \mathcal{H}} \left| \sum_{i=1}^{\text{fat}_\epsilon} \left( \sum_{j=1}^{\lceil n/\text{fat}_\epsilon \rceil} \sigma_{(i-1)\text{fat}_\epsilon + j} \right) (h(x_i^*) - h'(x_i^*)) \right| \right] \\ &\geq \frac{1}{2} \mathbb{E}_{\sigma \sim \text{Unif}\{\pm 1\}^{n'}} \left[ \frac{1}{n'} \left| \sum_{i=1}^{\text{fat}_\epsilon} \left( \sum_{j=1}^{\lceil n/\text{fat}_\epsilon \rceil} \sigma_{(i-1)\text{fat}_\epsilon + j} \right) (h_R(x_i^*) - h_{\overline{R}}(x_i^*)) \right| \right] \end{aligned}$$

where for each  $\sigma_1, \dots, \sigma_{n'}, R \subseteq [\text{fat}_\epsilon]$  is given by  $R = \left\{ i \in [\text{fat}_\epsilon] \mid \text{sign} \left( \sum_{j=1}^{\lceil n/\text{fat}_\epsilon \rceil} \sigma_{(i-1)\lceil n/\text{fat}_\epsilon \rceil + j} \right) \geq 0 \right\}$ ,  $h_R$  is the function in  $\mathcal{H}$  that  $\epsilon$ -shatters the set  $R$  and  $h_{\bar{R}}$  be the function that shatters the complement of set  $R$ .

$$\begin{aligned}
&\geq \frac{1}{2} \mathbb{E}_{\sigma \sim \text{Unif}\{\pm 1\}^{n'}} \left[ \frac{1}{n'} \sum_{i=1}^{\text{fat}_\epsilon} \left| \sum_{j=1}^{\lceil n/\text{fat}_\epsilon \rceil} \sigma_{(i-1)\text{fat}_\epsilon + j} \right| 2\epsilon \right] \\
&\geq \frac{\epsilon}{n'} \sum_{i=1}^{\text{fat}_\epsilon} \mathbb{E}_{\sigma \sim \text{Unif}\{\pm 1\}^{n'}} \left[ \left| \sum_{j=1}^{\lceil n/\text{fat}_\epsilon \rceil} \sigma_{(i-1)\text{fat}_\epsilon + j} \right| \right] \\
&\geq \frac{\epsilon \text{fat}_\epsilon}{n'} \sqrt{\frac{\lceil n/\text{fat}_\epsilon \rceil}{2}} \quad \text{(Khintchine's inequality)} \\
&= \sqrt{\frac{\epsilon^2 \text{fat}_\epsilon}{2 n'}}.
\end{aligned}$$

We can now conclude that:

$$\text{fat}_\epsilon \leq \frac{2n' \mathcal{R}_{n'}^2(\mathcal{H})}{\epsilon^2} \leq \frac{4n \mathcal{R}_n^2(\mathcal{H})}{\epsilon^2}$$

where last inequality is because Rademacher complexity decreases with increase in number of samples and  $n \leq n' \leq 2n$  (because  $\epsilon \geq \mathcal{R}_n(\mathcal{H})$  which implies that  $\text{fat}_\epsilon < n$ ).  $\square$



# Index

- $\ell_p$  norm, 122, 160
- AERM, 13
  - always AERM, 13
- Banach Lattice, 157
- consistency
  - strict, 23
- cotype, 158
- covering number
  - sequential, 58
  - statistical, 16
- cut norm, 124
- decision trees, 68
- decoupling, 159
- Dudley integrated complexity
  - sequential, 61
- excess risk, 11
- fat-shattering
  - sequential, 57
  - statistical, 16
- generalization, 12
- group norm, 123, 161
- Hilbert space, 17
- interpolation norm, 124
- isotron, 70
- learnability
  - online, 52
  - oracle-based, 109
  - statistical, 11
  - oracle-based, 108
- learning protocol
  - online learning, 50
  - oracle-based
    - offline, 106
    - online, 109
    - statistical, 107
- learning rule, 51
  - online, 51
  - randomized, 51
  - statistical, 11
    - proper, 11
    - randomized, 31
- Littlestone dimension, 57, 67
- margin, 67
- martingale type, 116
- max norm, 124
- Minkowski functional, 103
- mirror descent, 113
  - proxy function, 114
- mixed norm, 123, 161
- neural networks, 67
- online transductive learning, 69
- oracle, 105
- oracle based learning algorithm, 106
- oracle complexity
  - offline, 106
  - online, 109
  - statistical, 108
- oracle-based offline optimizable, 107
- p-convex, 170
- packing number
  - sequential
    - strong, 58
    - weak, 58
- path, 5
- q-concave, 171
- Rademacher complexity
  - sequential, 53
  - statistical, 15
- regret, 50
- sample complexity, 11
  - distribution specific, 12
- Schatten norm, 123, 161
- shattering
  - sequential, 57
- stability, 17, 24–29
  - (leave one out) LOO, 27
    - all-i, 27
    - on average, 27
    - uniform, 26
  - (replace one) RO
    - average, 25
    - strongly uniform, 33
    - uniform, 24
- Stochastic Approximation (SA), 108
- stochastic optimization, 10
- subtree, 5
- tangent sequence, 55
- tree, 5
- type, 152
- uniform convergence
  - statistical, 13
  - universal, 62
- uniform convexity, 113
- value of the online learning game, 51–53, 111
- VC dimension, 59
- Walsh-Paley martingales, 116