

Efficient Online Learning, Deterministic, and Stochastic Optimization

Yoram Singer
Machine Intelligence Center
Google Research




High Dimensional Sparse Data

- To predict the MID of an entity a large number of boolean predicates are built and combined
- Most predicates evaluate to be false for most examples
- Example: $[\omega_t = \text{President-Name}] \ \& \ [\omega_{t+1} = \text{"White-House"}]$

High Dimensional Sparse Data

- To predict the MID of an entity a large number of boolean predicates are built and combined
- Most predicates evaluate to be false for most examples
- Example: $[\omega_t = \text{President-Name}] \ \& \ [\omega_{t+1} = \text{"White-House"}]$



High Dimensional Sparse Data

- To predict the MID of an entity a large number of boolean predicates are built and combined
- Most predicates evaluate to be false for most examples
- Example: $[\omega_t = \text{President-Name}] \ \& \ [\omega_{t+1} = \text{"White-House"}]$

How to use the predicates in order to make accurate predictions ?

$$\mathbf{x} \in \{0, 1\}^n \quad (0, 0, 1, 0, 1, 0, \dots, 0, 0, 1, 0)$$

High Dimensional Sparse Data

- To predict the MID of an entity a large number of boolean predicates are built and combined
- Most predicates evaluate to be false for most examples
- Example: $[\omega_t = \text{President-Name}] \ \& \ [\omega_{t+1} = \text{"White-House"}]$

Which predicates are “important” for performing accurate predictions?

$\mathbf{x} \in \{0, 1\}^n$ (0, 0, 1, 0, 1, 0, . . . , 0, 0, 1, 0)

Setting in a Picture

INSTANTIATED PREDICATES == FEATURES

(e.g. MID="Obama")

Target **Y**

Example X_1

	f_1	f_2	f_3	f_4	f_5	f_6	f_7	f_8
X_1	0	1	0	1	1	0	0	0
X_2	1	0	1	0	0	0	1	0
	1	1	1	1	0	1	0	1
	0	1	0	0	0	0	1	0
	1	0	0	0	0	0	1	1
	1	0	0	1	0	0	1	0
X_7	0	1	1	0	0	1	0	1

X_2

0
0
1
0
0
0
0
1

Setting in a Picture

INSTANTIATED PREDICATES == FEATURES

(e.g. MID="Obama")

Target **Y**

Example X_1

X_2

X_7

	f_1	f_2	f_3	f_4	f_5	f_6	f_7	f_8
X_1	0	1	0	1	1	0	0	0
X_2	1	0	1	0	0	0	1	0
	1	1	1	1	0	1	0	1
	0	1	0	0	0	0	1	0
	1	0	0	0	0	0	1	1
	1	0	0	1	0	0	1	0
X_7	0	1	1	0	0	1	0	1

0
0
1
0
0
0
1

**FREQUENT BUT
NON-INFORMATIVE**



Setting in a Picture

INSTANTIATED PREDICATES == FEATURES

(e.g. MID="Obama")

Target **Y**

Example X_1

X_2

X_7

	f_1	f_2	f_3	f_4	f_5	f_6	f_7	f_8
X_1	0	1	0	1	1	0	0	0
X_2	1	0	1	0	0	0	1	0
	1	1	1	1	0	1	0	1
	0	1	0	0	0	0	1	0
	1	0	0	0	0	0	1	1
	1	0	0	1	0	0	1	0
X_7	0	1	1	0	0	1	0	1

0
0
1
0
0
0
0
1

**FREQUENT BUT
NON-INFORMATIVE**

**INFREQUENT YET
INFORMATIVE**

Setting in a Picture

INSTANTIATED PREDICATES == FEATURES

(e.g. MID="Obama")

Target **Y**

Example X_1

X_2

X_7

	f_1	f_2	f_3	f_4	f_5	f_6	f_7	f_8
X_1	0	1	0	1	1	0	0	0
X_2	1	0	1	0	0	0	1	0
	1	1	1	1	0	1	0	1
	0	1	0	0	0	0	1	0
	1	0	0	0	0	0	1	1
	1	0	0	1	0	0	1	0
X_7	0	1	1	0	0	1	0	1

0
0
1
0
0
0
0
1

**FREQUENT BUT
NON-INFORMATIVE**

**INFREQUENT YET
INFORMATIVE**

**TRUE ALSO FOR REAL
VALUED FEATURES**

Challenges

Challenges

- Large amounts of high dimensional sparse data:
 - Computational time should scale with #“1” features
 - Cannot process entire dataset “all at once”

Challenges

- Large amounts of high dimensional sparse data:
 - Computational time should scale with #“1” features
 - Cannot process entire dataset “all at once”
- Many *frequent* features are *irrelevant*

Challenges

- Large amounts of high dimensional sparse data:
 - Computational time should scale with #“1” features
 - Cannot process entire dataset “all at once”
- Many *frequent* features are *irrelevant*
- Some of the *infrequent* features are *highly relevant*

Challenges

- Large amounts of high dimensional sparse data:
 - Computational time should scale with #“1” features
 - Cannot process entire dataset “all at once”
- Many *frequent* features are *irrelevant*
- Some of the *infrequent* features are *highly relevant*
- Need to learn relatively compact models:
 - Training can use lots of (distributed) memory & CPUs
 - Serving (testing) is performed on many more instances than training and often should be

Outline

- Brief reminder:
linear models, empirical loss, regularization
- Convexity, Smoothness, and L_1 regularization
- Gradients & Subgradients for loss minimization
- Gradient Descent & Stochastic Gradient Methods
- Proximal view of GD & SG
- *Fobos*: dimension efficient proximal method
- *AdaGrad*: feature efficient adaptive proximal method

Elementary Start: Linear Models

Instance **X**

X ₁	X ₂	X ₃	X ₄	X ₅	X ₆	X ₇	X ₈
----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------

Weights **W**

W ₁	W ₂	W ₃	W ₄	W ₅	W ₆	W ₇	W ₈
----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------



Prediction $\hat{y} = \mathbf{w} \cdot \mathbf{x} = \sum_{j=1}^n w_j x_j$

True Target $y \Rightarrow \ell(y, \hat{y})$ (loss function)

Elementary Start: Linear Models

Instance \mathbf{x}

x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8
-------	-------	-------	-------	-------	-------	-------	-------

Weights \mathbf{w}

w_1	w_2	w_3	w_4	w_5	w_6	w_7	w_8
-------	-------	-------	-------	-------	-------	-------	-------




Prediction $\hat{y} = \mathbf{w} \cdot \mathbf{x} = \sum_{j=1}^n w_j x_j$

True Target $y \Rightarrow \ell(y, \hat{y})$ (loss function)


Example of losses

$\ell(y, \hat{y}) = (y - \hat{y})^2$ squared error $\ell(y, \hat{y}) = e^{-y\hat{y}}$ exponential loss


Empirical Loss (Linear Regression)




Empirical Loss (Linear Regression)




Empirical Loss (Linear Regression)




Empirical Loss (Linear Regression)



Empirical Loss (Linear Regression)




Empirical Loss (Linear Regression)



$$\frac{1}{10} (\text{■} + \text{■} + \text{■} + \text{■} + \text{■} + \text{■} + \text{■} + \text{■} + \text{■} + \text{■})$$

Empirical Loss (Linear Regression)





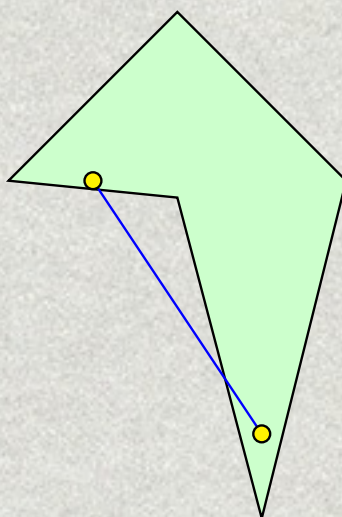
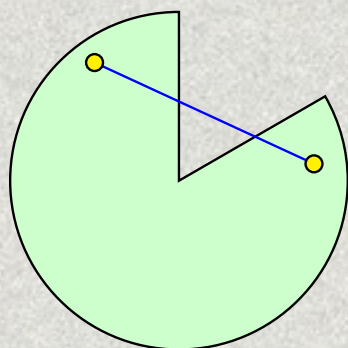
Empirical
Loss

$$L(\mathbf{w}) = \frac{1}{m} \sum_{i=1}^m \ell(y_i, \mathbf{w} \cdot \mathbf{x}_i)$$

Convex Losses

non-convex




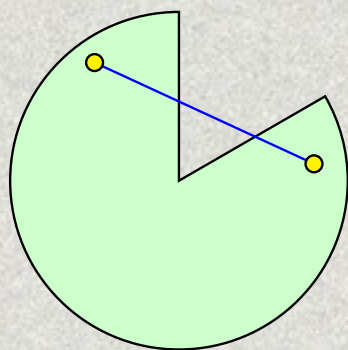
convex




Convex Losses

non-convex


convex



$$f(\alpha \mathbf{u} + (1 - \alpha) \mathbf{v}) \leq \alpha f(\mathbf{u}) + (1 - \alpha) f(\mathbf{v})$$




Overfitting




Overfitting

Suppose we were able to fit a spline function to the data




Overfitting

Then the empirical loss $L(\mathbf{w})$ would be 0




Overfitting



However new data points
are unlikely to reside on
the piecewise linear curve
“*overfitting*”


Overfitting



But, is it possible to overfit with a linear model?

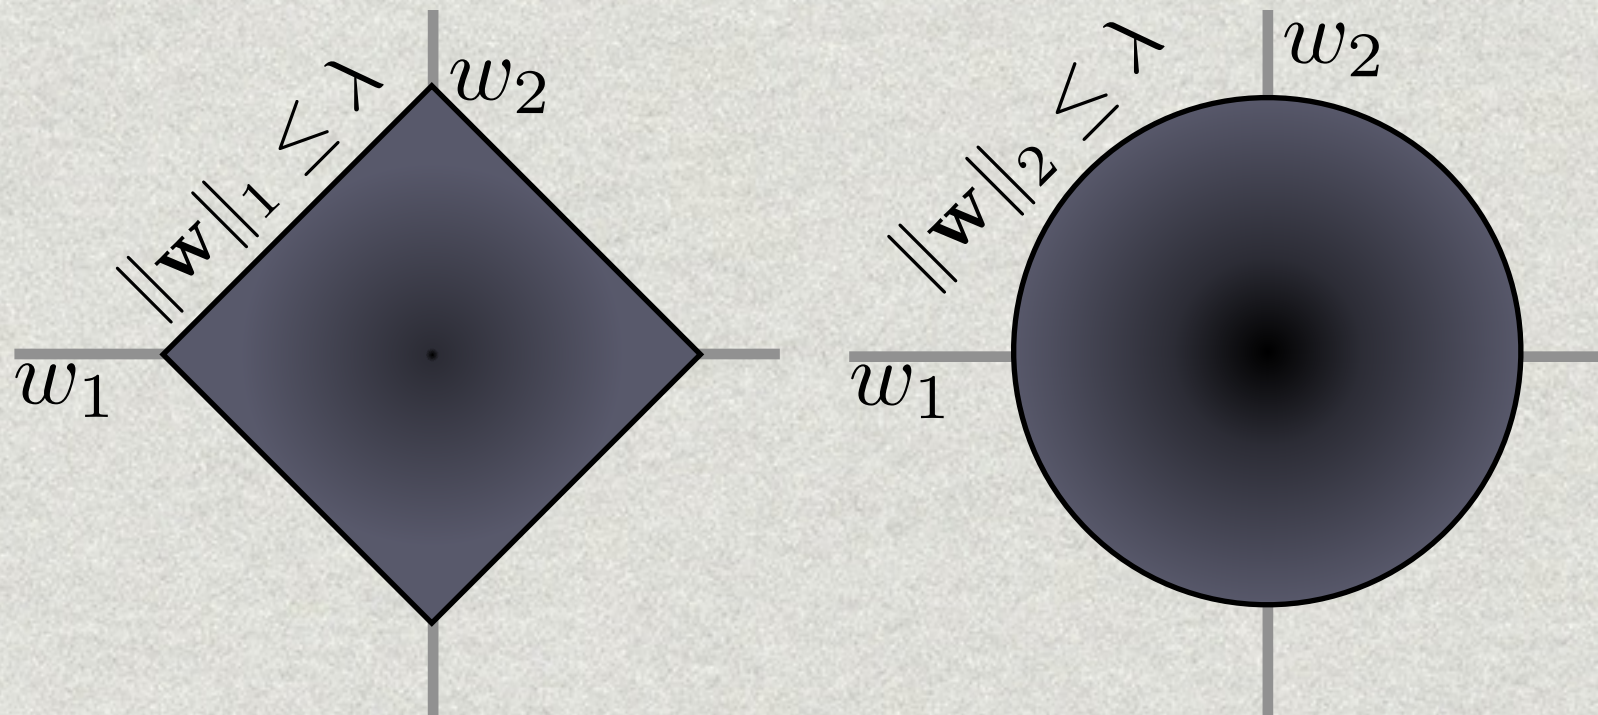
Overfitting

Yes, when number of features is very large & many are irrelevant

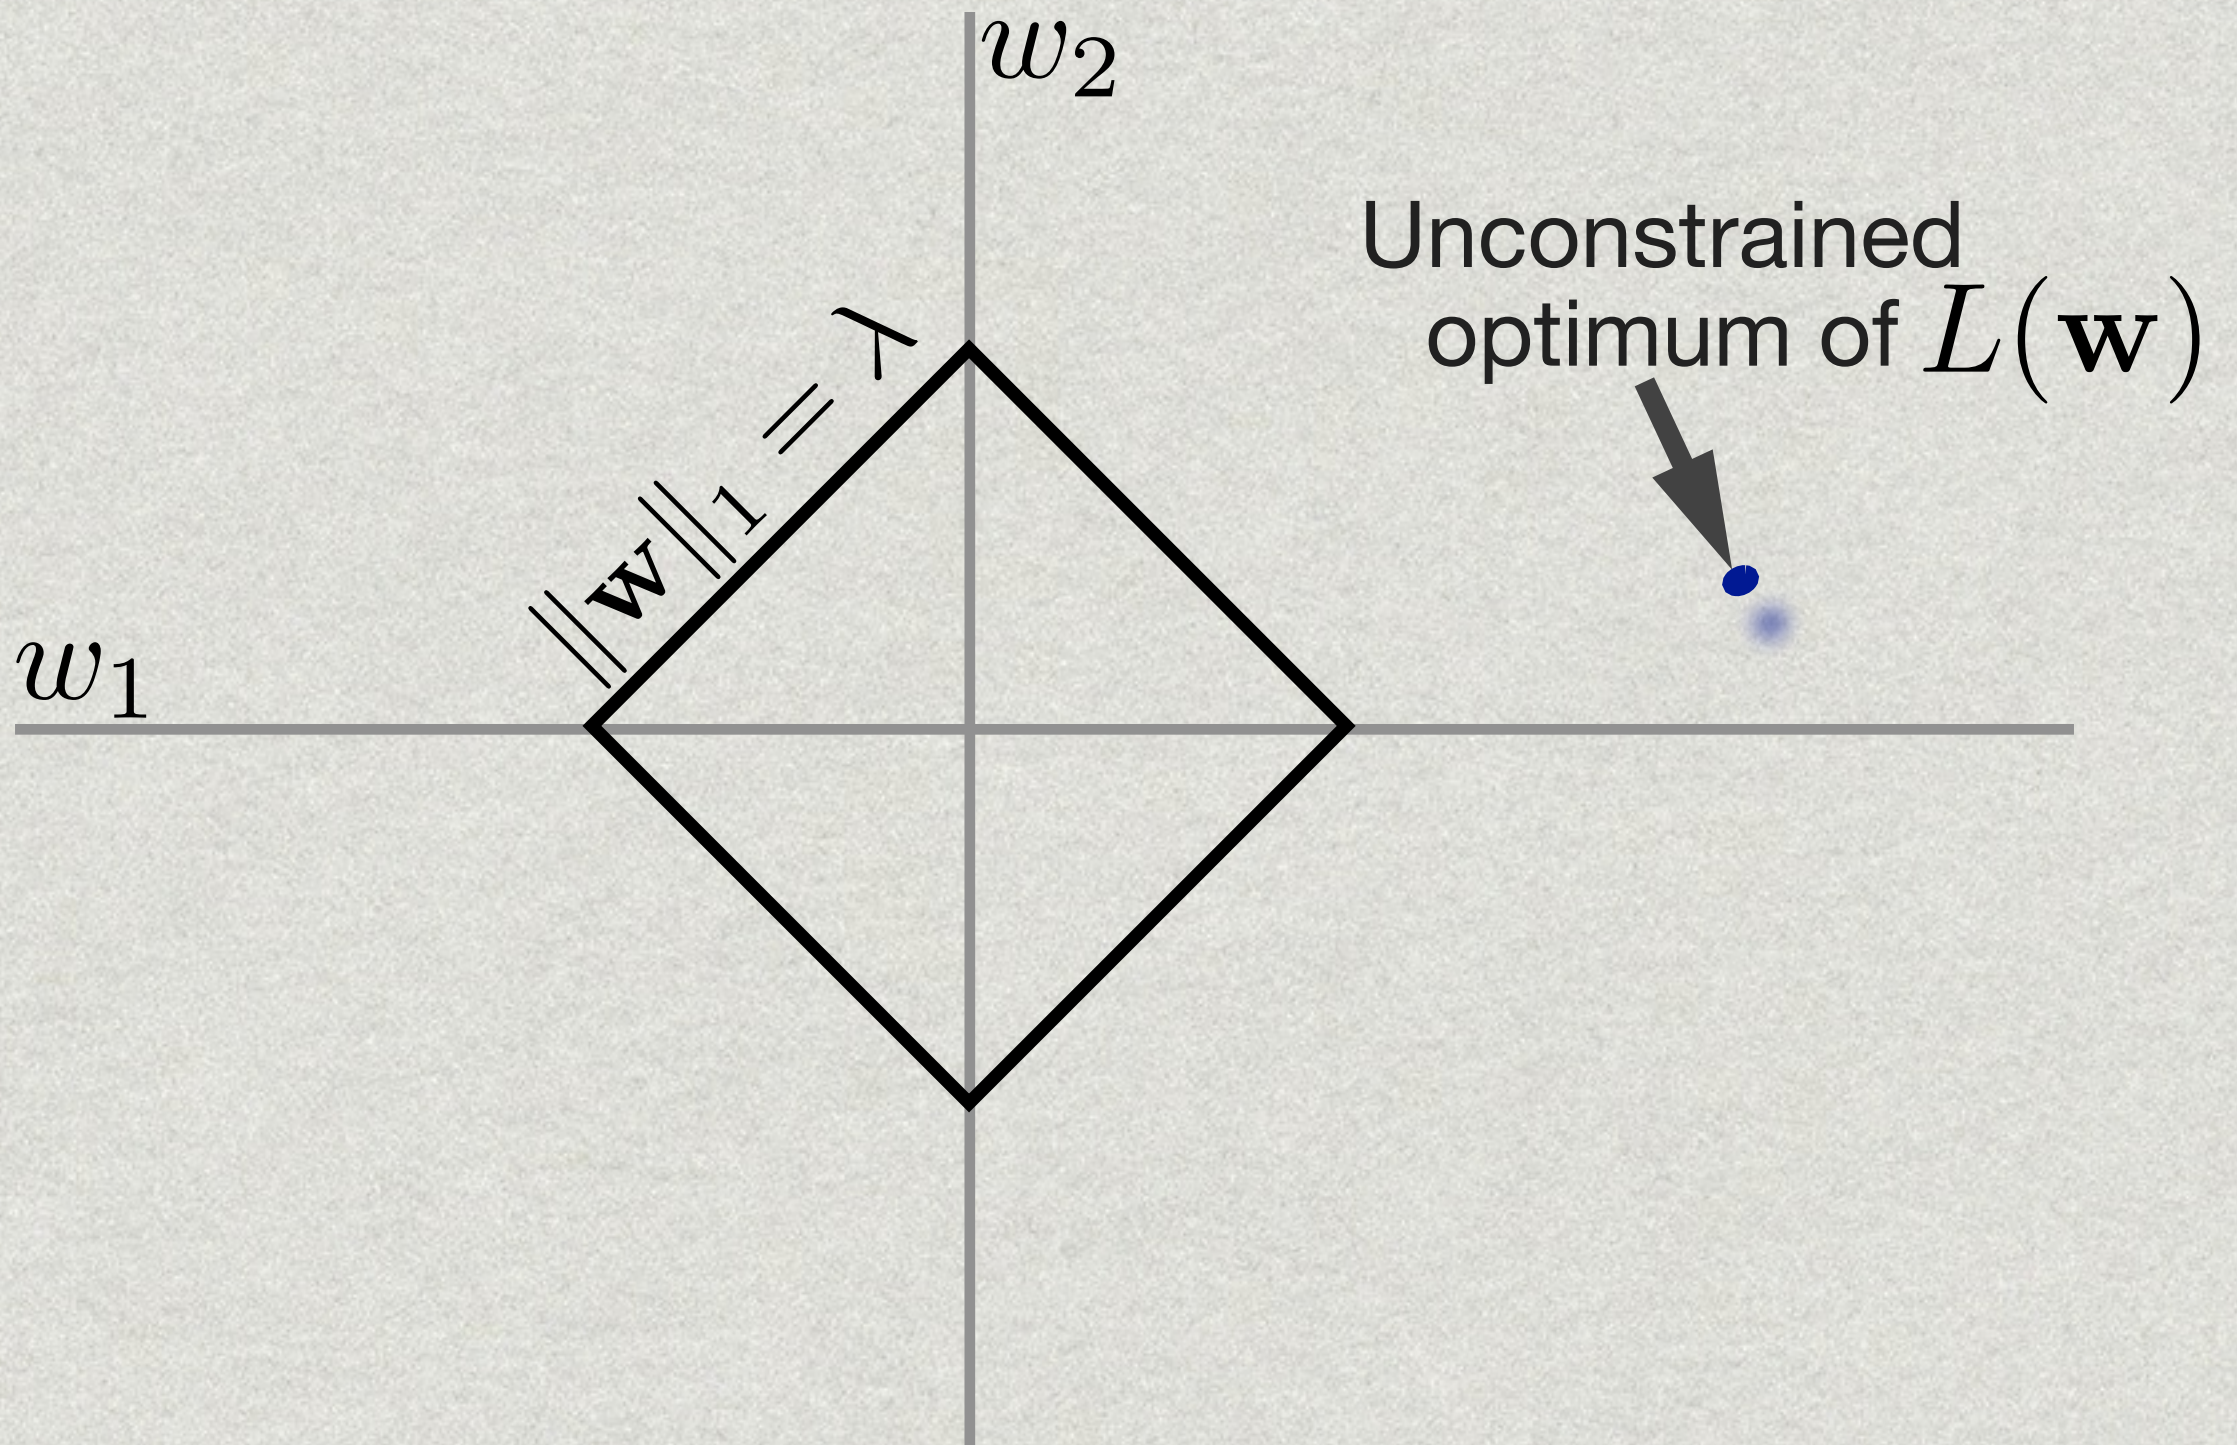


Preventing Overfitting

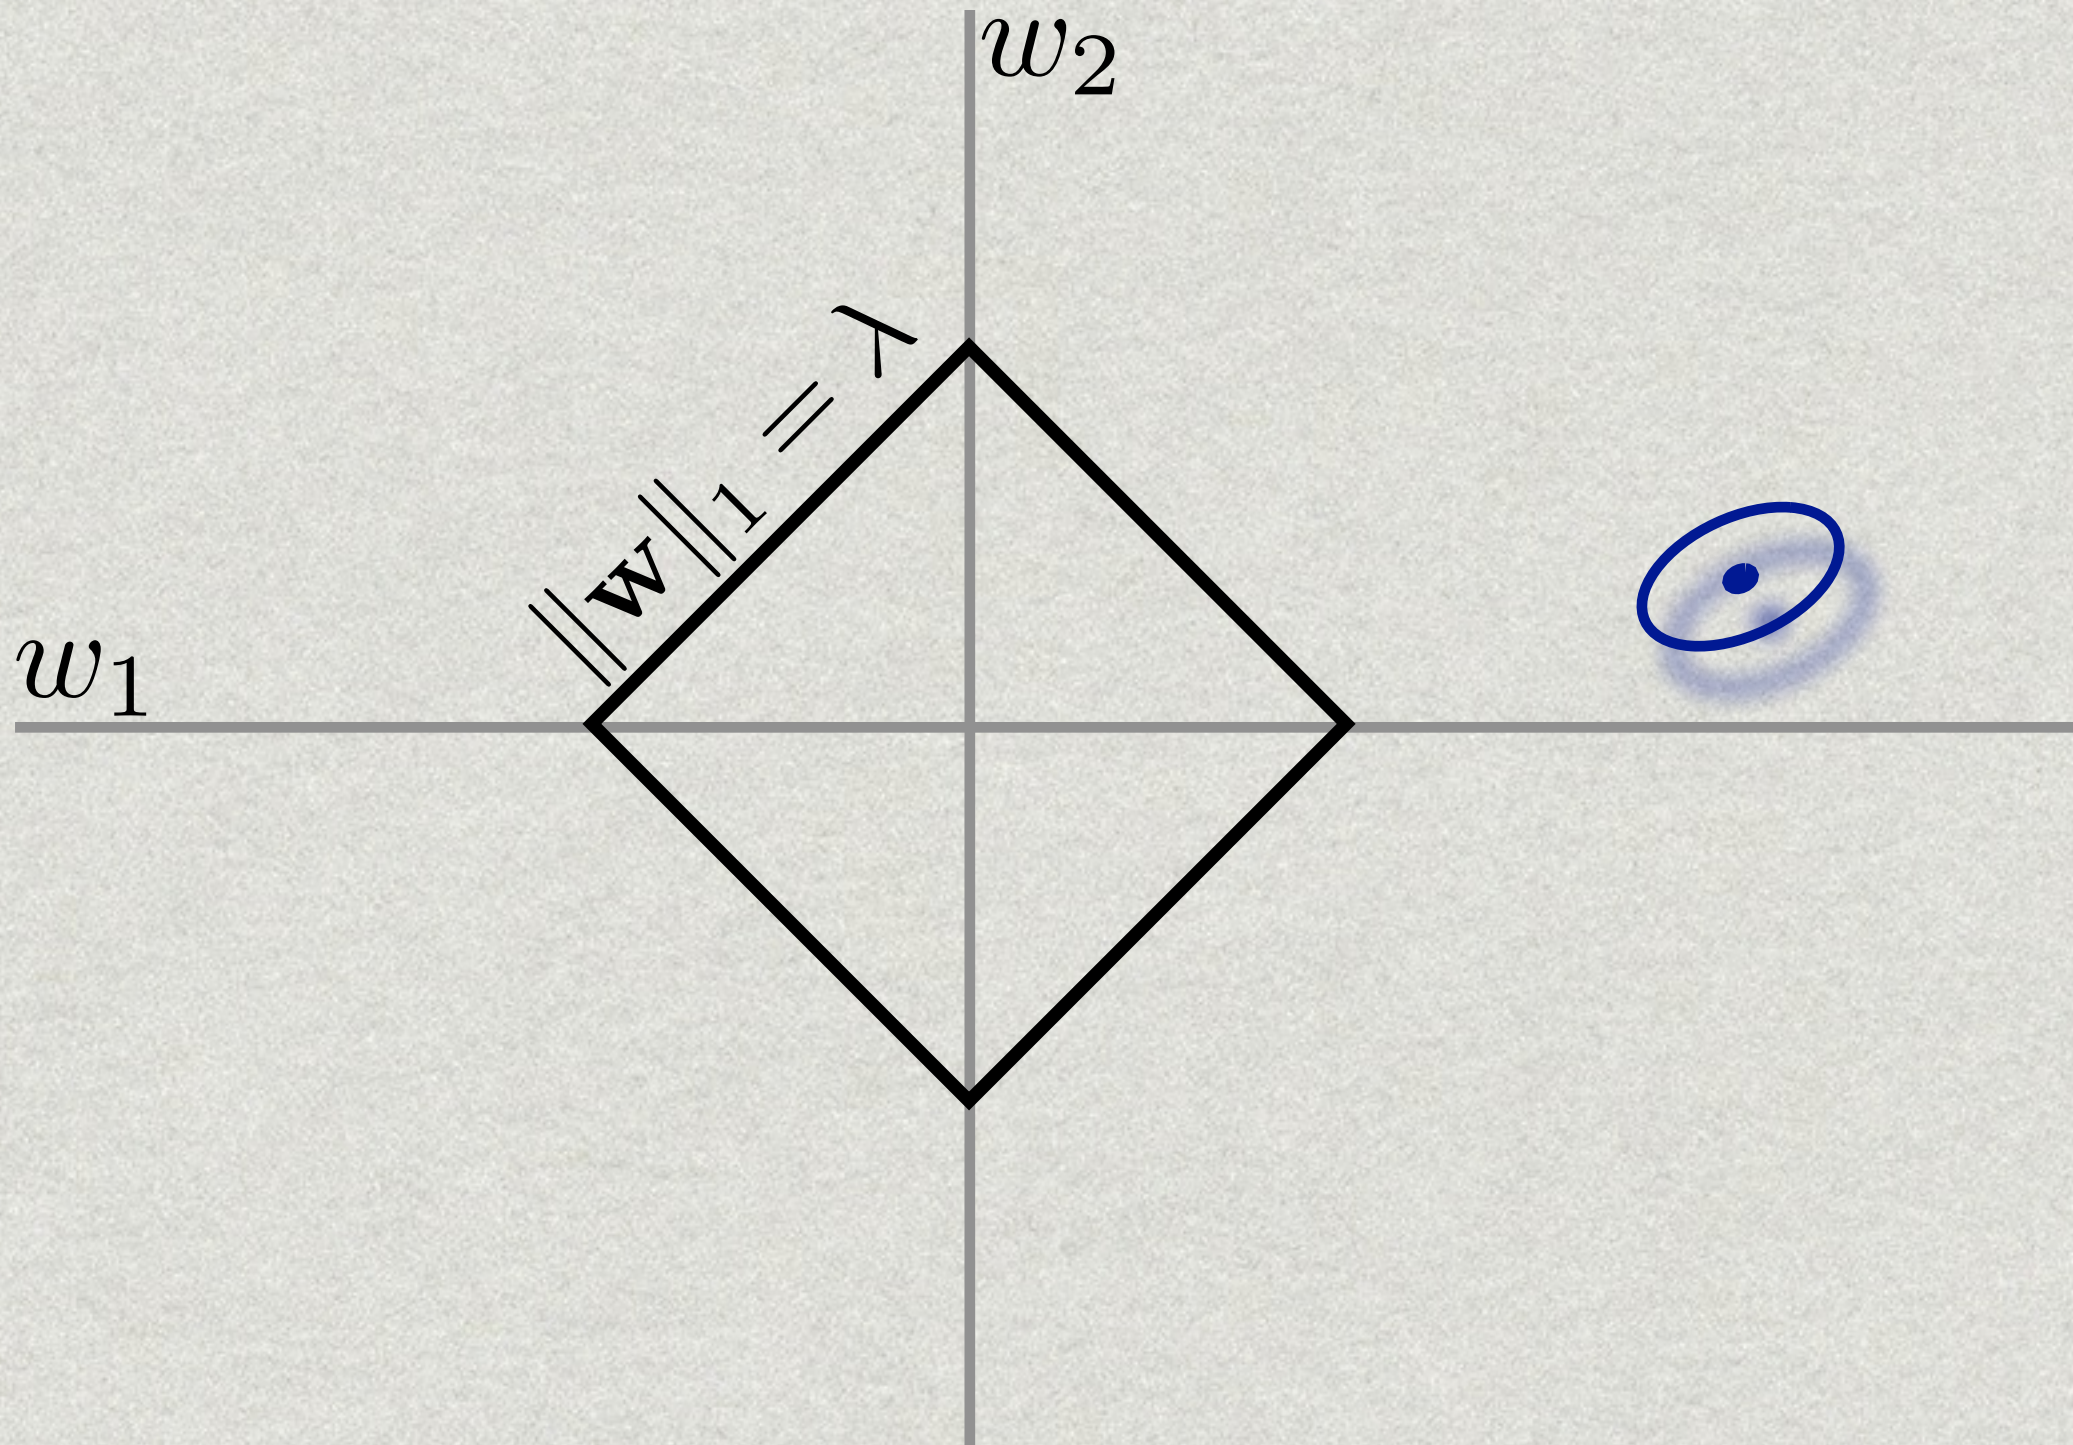
- To prevent overfitting we need to constrain the volume of the space of the possible linear predictors
- A common approach is to limit the p -norm of \mathbf{W}



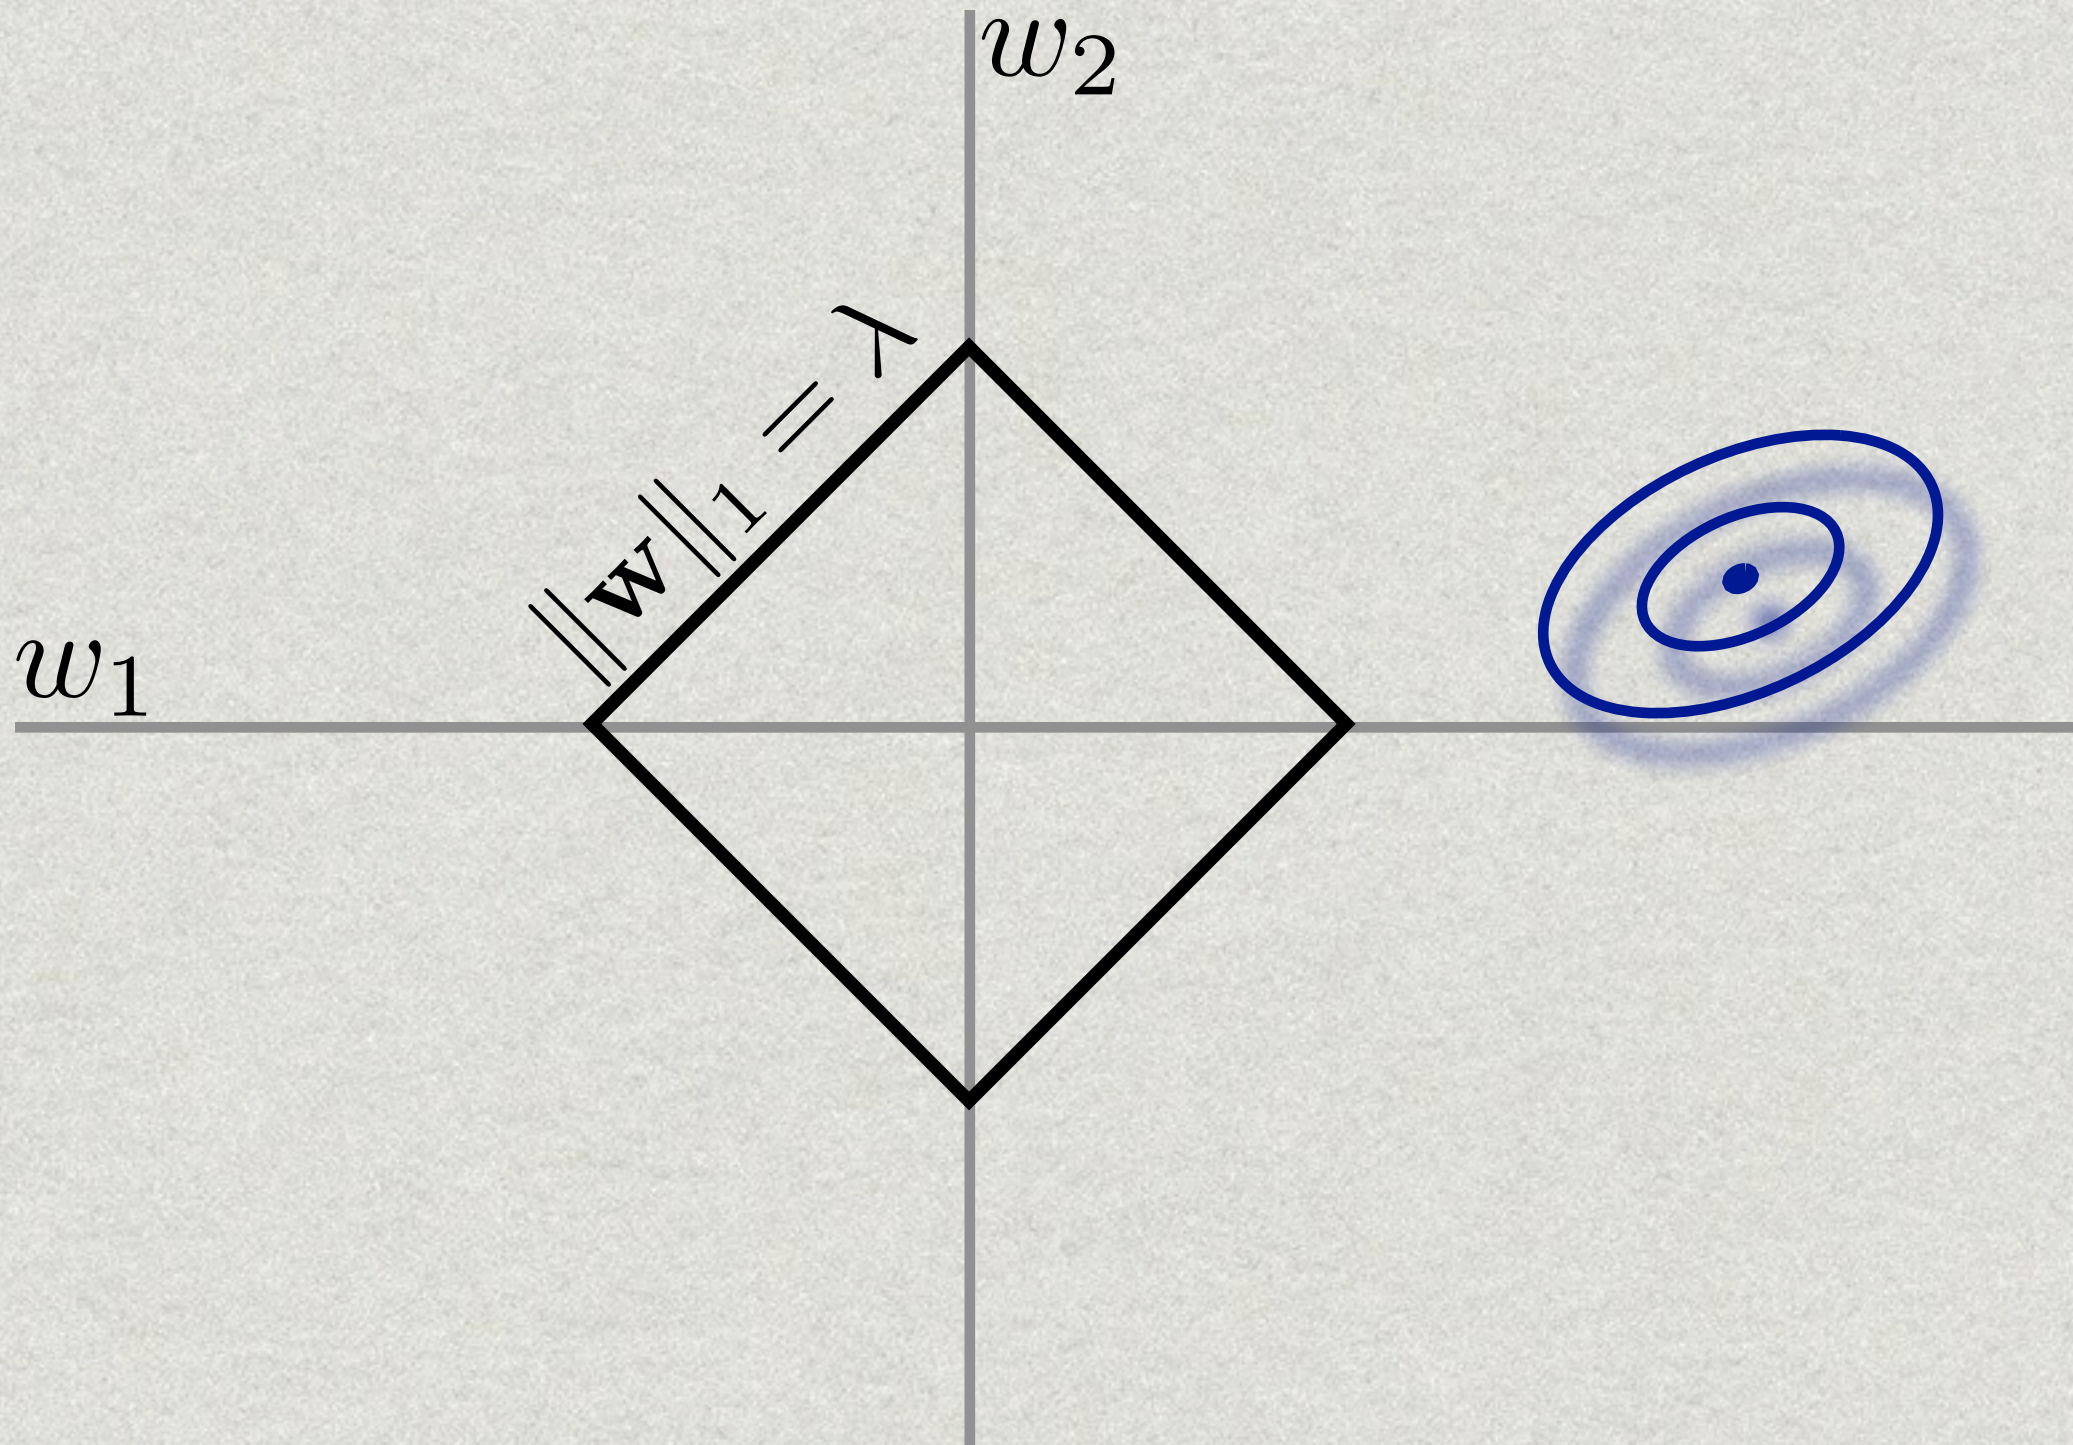
Achieving Sparsity using 1-norm




Achieving Sparsity using 1-norm



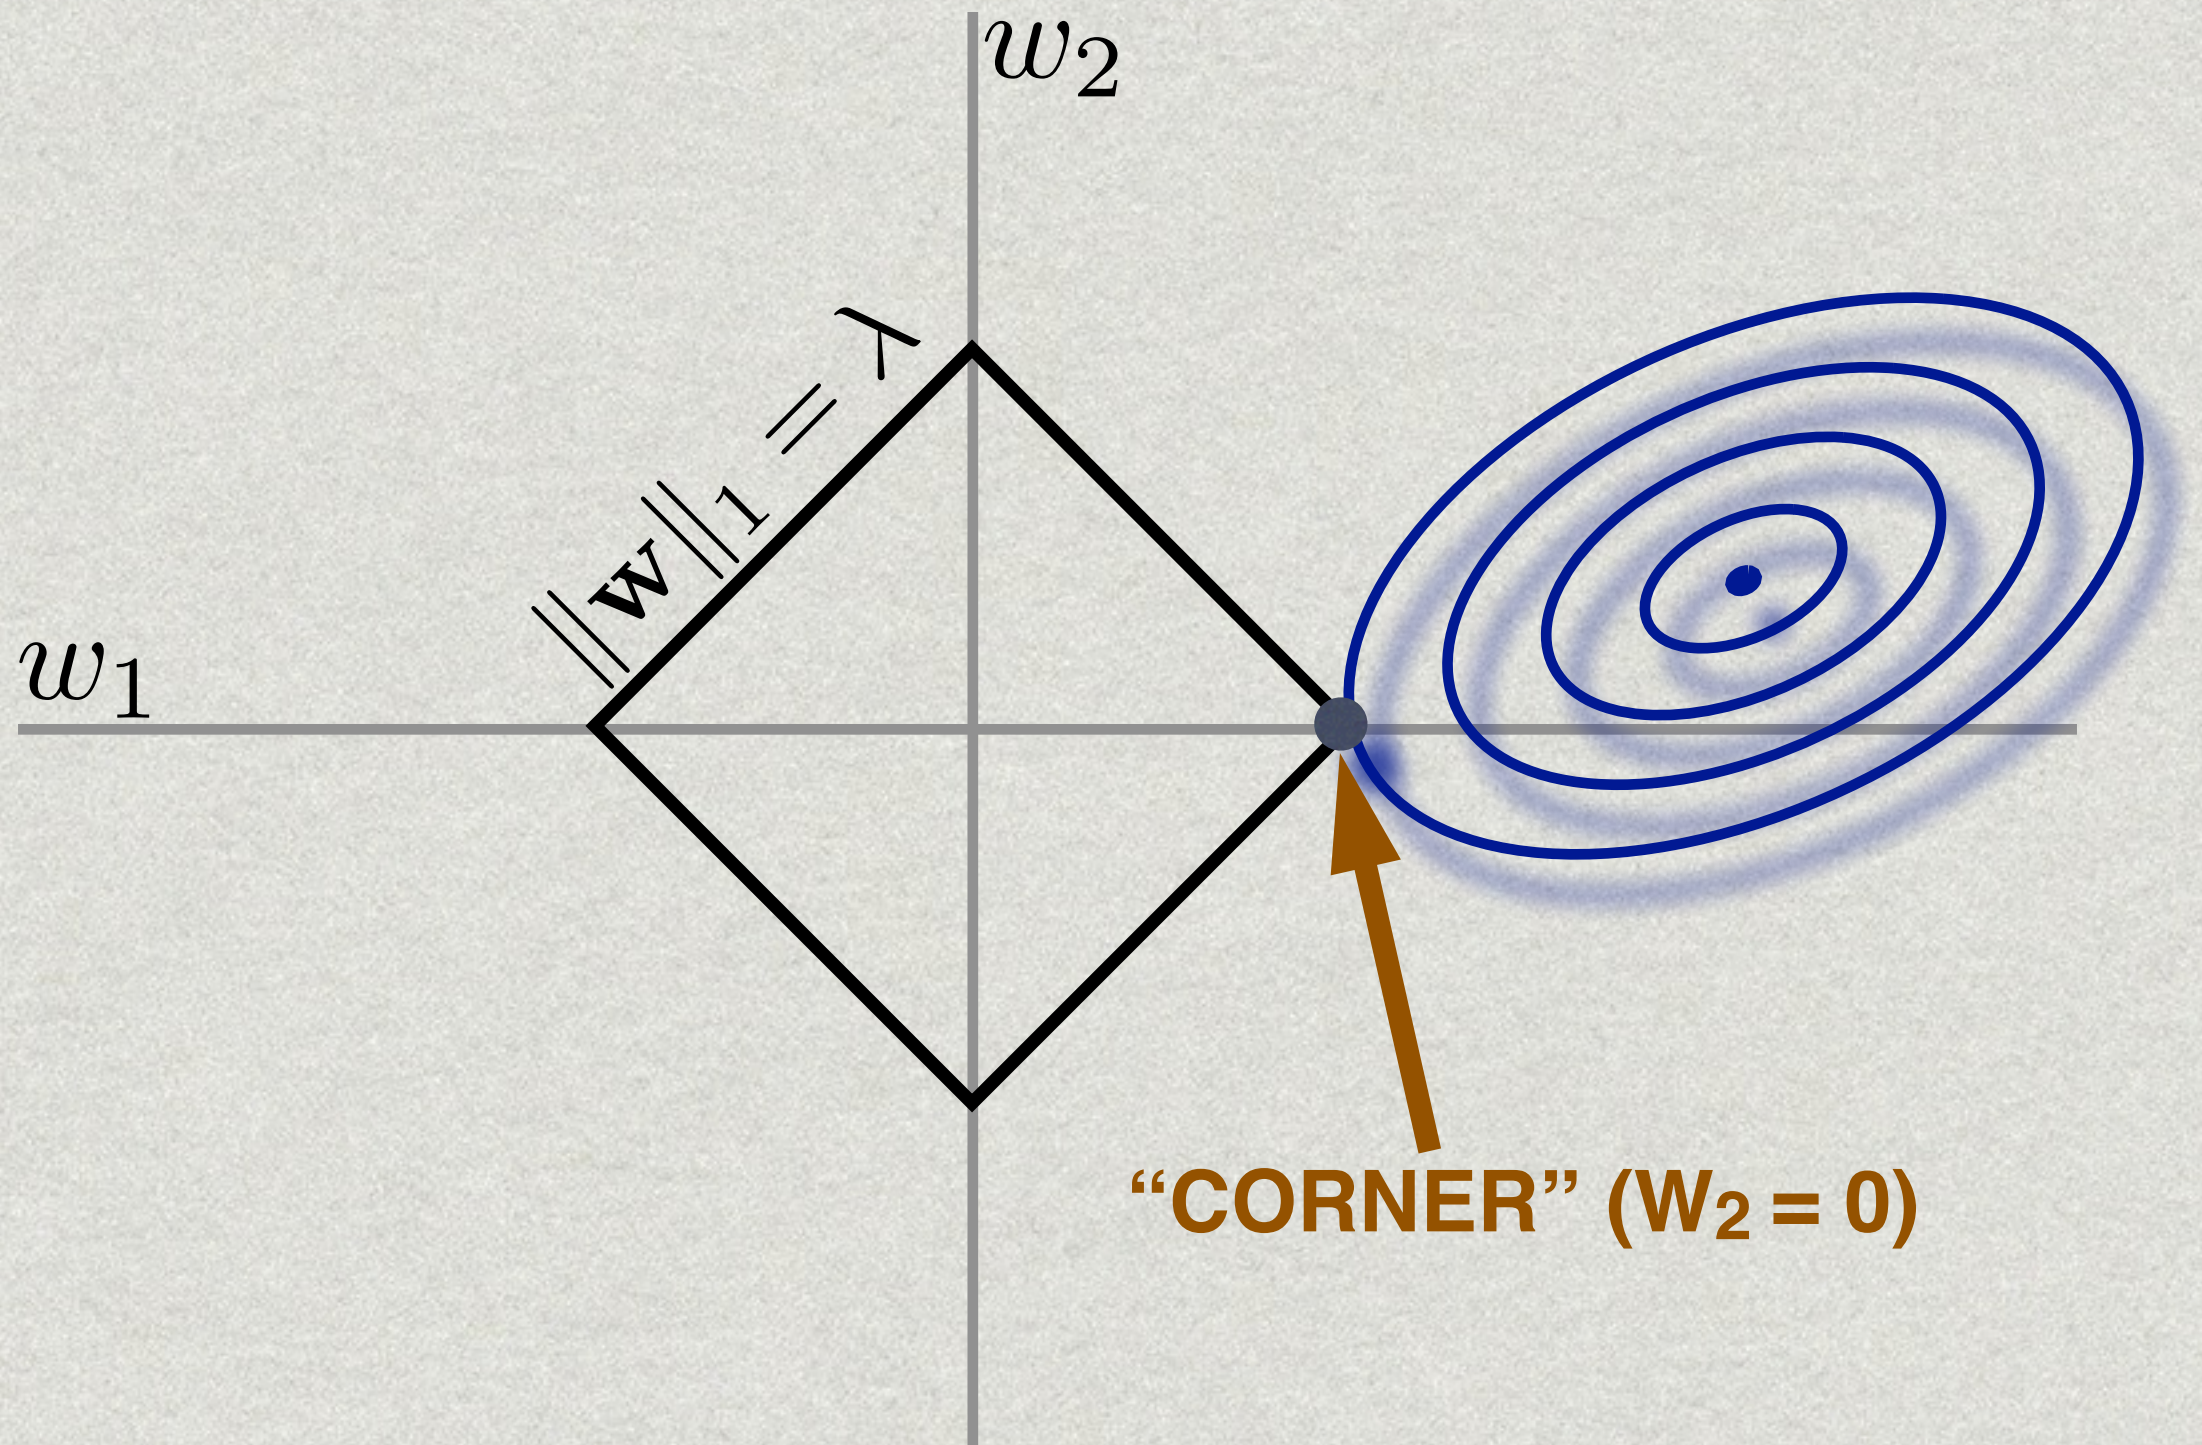
Achieving Sparsity using 1-norm




Achieving Sparsity using 1-norm



Achieving Sparsity using 1-norm



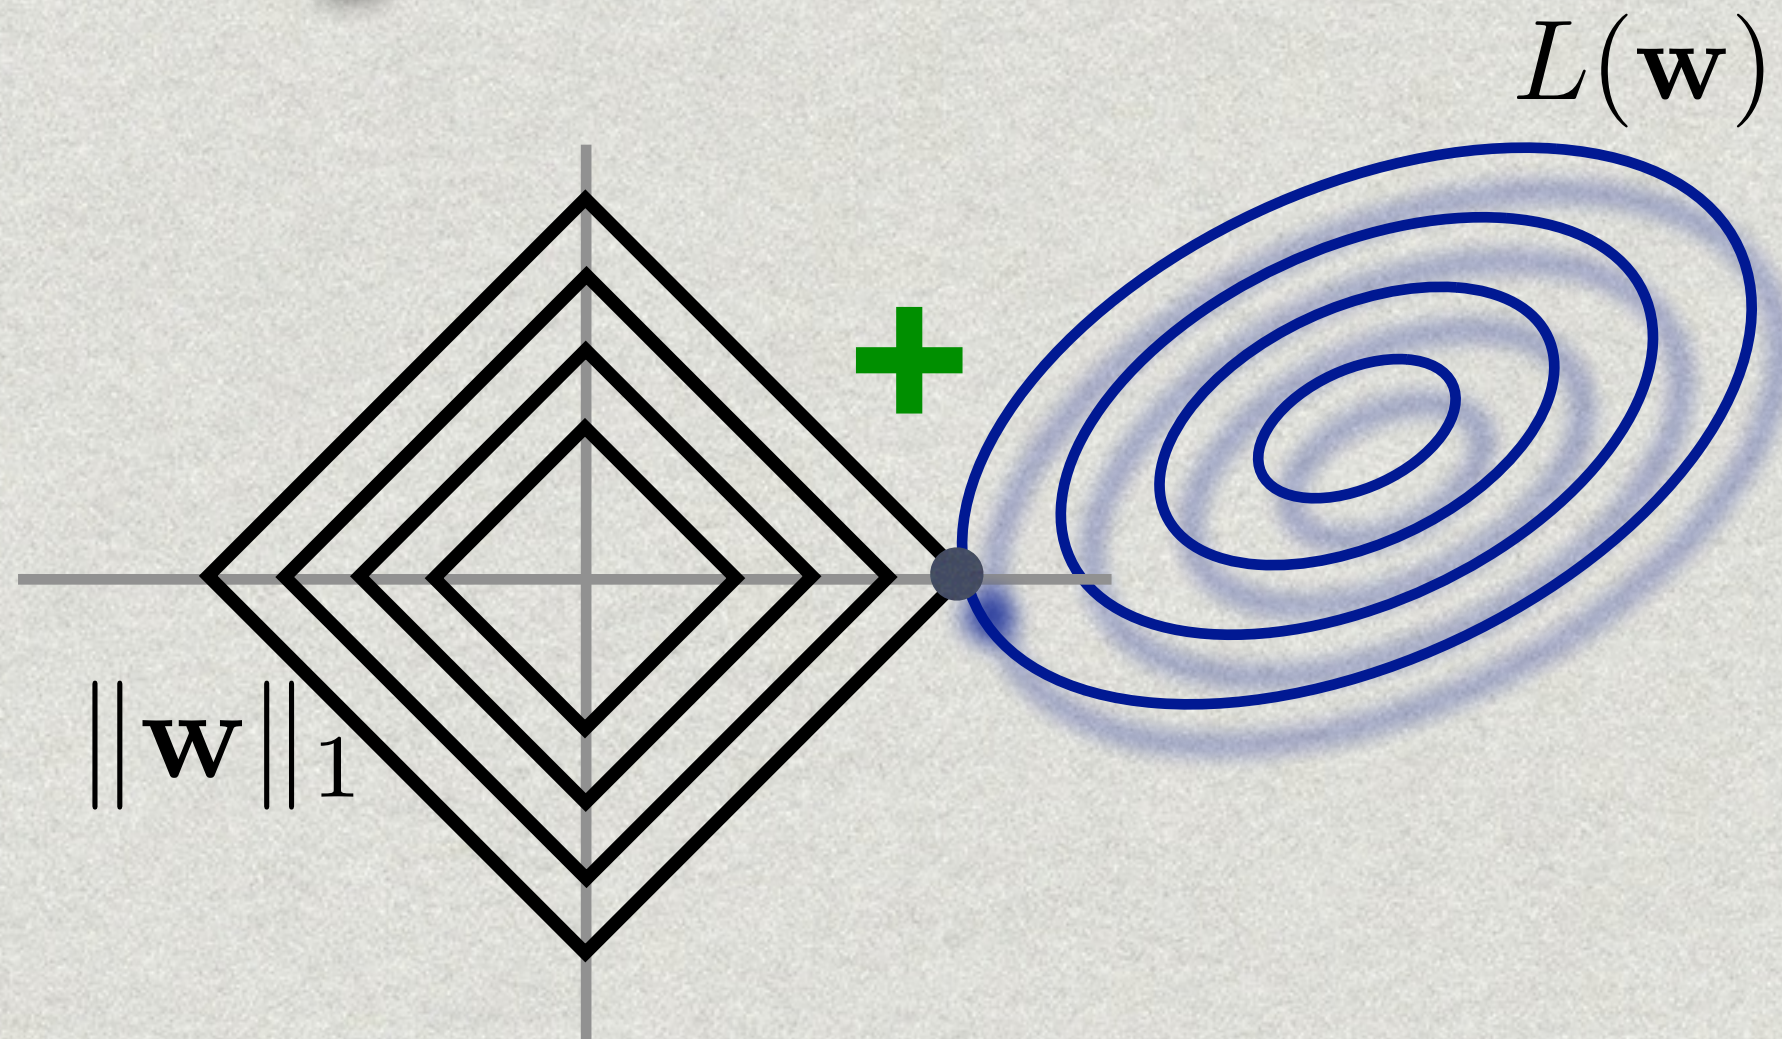
Achieving Sparsity using 1-norm



See e.g. Candes'06, Donoho'06


Penalized Form

$$\min_{\mathbf{w}} L(\mathbf{w}) + \lambda \|\mathbf{w}\|_1$$




SPARSITY PROPERTIES ARE ANALOGOUS TO CONSTRAINED FORM


Gradient Descent




Gradient Descent




Gradient Descent




Gradient Descent




Gradient Descent



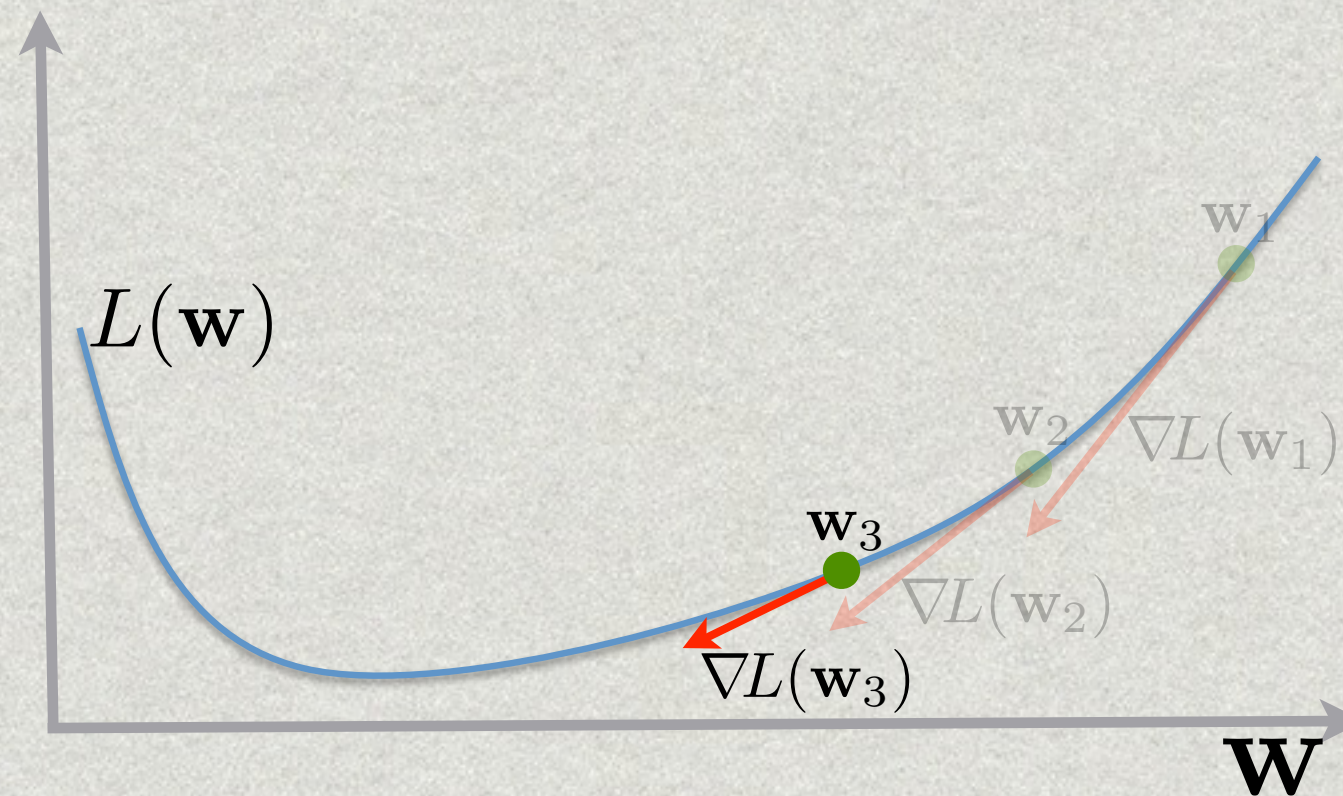
Gradient Descent



Gradient Descent



Gradient Descent




- Gradient descent main loop:

- Compute gradient $\nabla_t L = \frac{1}{|S|} \sum_{i \in S} \frac{\partial}{\partial \mathbf{w}} \ell(\mathbf{w}_t; (\mathbf{x}_i, y_i))$

- Update

$$\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t - \eta_t \nabla_t L$$

Gradient Descent



- Gradient descent main loop:

- Compute gradient $\nabla_t L = \frac{1}{|S|} \sum_{i \in S} \frac{\partial}{\partial \mathbf{w}} \ell(\mathbf{w}_t; (\mathbf{x}_i, y_i))$


- Update

$$\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t - \overset{\text{STEP SIZE}}{\eta_t} \nabla_t L$$


Lipschitz & Smooth Convex Losses



Lipschitz & Smooth Convex Losses



Lipschitz & Smooth Convex Losses



Lipschitz Losses

- Domain $\Omega \subset \mathbb{R}^d$ Loss function: $\mathcal{L} : \mathbb{R}^d \rightarrow \mathbb{R}_+$
- Lipschitz losses change sufficiently “slow”
$$\beta - \text{Lipschitz} \iff |\mathcal{L}(\mathbf{w}) - \mathcal{L}(\mathbf{v})| \leq \beta \|\mathbf{w} - \mathbf{v}\|$$
- $|x|$ is Lipschitz over the entire reals but x^2 is not!
- Homework Q.1: what is the Lipschitz constant for $\log(1+\exp(x))$ and what is the domain
- Homework Q.2: if L and Q are Lipschitz functions with constants β_1 & β_2 , what is the Lipschitz constant for $L(Q(\mathbf{w}))$ [Note that L is a scalar function while Q is a vector function]

Smooth Losses

- A loss is β -smooth if its gradient is β -Lipschitz
[Note that we extended Lipschitz to vector functions]

$$\|\nabla \mathcal{L}(\mathbf{w}) - \nabla \mathcal{L}(\mathbf{v})\| \leq \beta \|\mathbf{w} - \mathbf{v}\|$$

- Homework Q.3: show that if a loss is β -smooth then

$$\mathcal{L}(\mathbf{w}) \leq \mathcal{L}(\mathbf{v}) + \nabla \mathcal{L}(\mathbf{v}) \cdot (\mathbf{w} - \mathbf{v}) + \frac{\beta}{2} \|\mathbf{w} - \mathbf{v}\|^2$$

Gradient Descent for Lipschitz Losses

- Assume that loss function is β -Lipschitz
- Perform the following updates:

$$\mathbf{w}^{t+1} = \mathbf{w}^t - \eta_t \nabla \mathcal{L}(\mathbf{w}^t) \quad \text{where} \quad \eta_t = \tilde{O}\left(1/\sqrt{t}\right)$$

- Let w^* be the minimizer of the loss over the domain $\{w \text{ s.t. } \|w\| < r\}$
- Let u be the average of w^t from $t=1$ through T
- Then, the gap between u and w^* w.r.t loss is

$$\mathcal{L}(\mathbf{u}) - \mathcal{L}(\mathbf{w}^*) = \mathcal{L}\left(\frac{1}{T} \sum_{t=1}^T \mathbf{w}^t\right) - \mathcal{L}(\mathbf{w}^*) \leq \frac{r\beta}{\sqrt{T}}$$

Proof Outline

- Use convexity to upper bound the difference between the loss at u and the loss at w^*
- Use the distance between w^t and w^* as potential
- Find a learning rate that minimizes at each iteration a bound on the potential
- Important comments on smoothness and stochastic optimization to follow the proof
- See also Section 14.1 in:
Understanding Machine Learning: From Theory to Algorithms
by Shai Shalev-Shwartz & Shai Ben-David

Stochastic Optimization

Training set is large and the source is i.i.d then we can sub-sample S to obtain an estimate of the gradient

$$\hat{\nabla}_t L = \frac{1}{|S'|} \sum_{i \in S'} \frac{\partial}{\partial \mathbf{w}} \ell(\mathbf{w}_t; (\mathbf{x}_i, y_i))$$

Stochastic Optimization

Training set is large and the source is i.i.d then we can sub-sample S to obtain an estimate of the gradient

$$\hat{\nabla}_t L = \frac{1}{|S'|} \sum_{i \in S'} \frac{\partial}{\partial \mathbf{w}} \ell(\mathbf{w}_t; (\mathbf{x}_i, y_i)) \quad S' \subset S$$

Stochastic Optimization

Training set is large and the source is i.i.d then we can sub-sample S to obtain an estimate of the gradient


$$\hat{\nabla}_t L = \frac{1}{|S'|} \sum_{i \in S'} \frac{\partial}{\partial \mathbf{w}} \ell(\mathbf{w}_t; (\mathbf{x}_i, y_i)) \quad S' \subset S$$

Convergence Rate still holds in expectation {over S' } !

Subgradients

Subgradient set of a function f at x_0


$$\partial f(x_0) = \{g : f(x) \geq f(x_0) + g^\top(x - x_0)\}$$



Subgradients

Subgradient set of a function f at x_0

$$\partial f(x_0) = \{g : f(x) \geq f(x_0) + g^\top(x - x_0)\}$$



Minimization using Subgradients

Minimize

$$\min_{\mathbf{w}} L(\mathbf{w}) + \lambda \|\mathbf{w}\|_1$$

- Unconstrained stochastic **sub**gradient descent

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta_t \mathbf{g}_t \quad \mathbf{g}_t \in \hat{\nabla}_t L + \partial \|\mathbf{w}_t\|_1$$

Minimization using Subgradients

Minimize

$$\min_{\mathbf{w}} L(\mathbf{w}) + \lambda \|\mathbf{w}\|_1$$

- Unconstrained stochastic **sub**gradient descent

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta_t \mathbf{g}_t \quad \mathbf{g}_t \in \hat{\nabla}_t L + \partial \|\mathbf{w}_t\|_1$$



$$\partial |w_{t,j}| = \text{sign}(w_{t,j})$$

Subgradients: Caveat

Subgradients are “non-informative” at singularities

Subgradients: Caveat

Subgradients are “non-informative” at singularities




Subgradients: Caveat

Subgradients are “non-informative” at singularities




Subgradients: Caveat

Subgradients are “non-informative” at singularities




Subgradients: Caveat

Subgradients are “non-informative” at singularities



Subgradients: Caveat

Subgradients are “non-informative” at singularities



- **DENSE SOLUTION FOR W**

Subgradients: Caveat


Subgradients are “non-informative” at singularities




- **DENSE SOLUTION FOR W**
- **SLOW CONVERGENCE**

Fobos


Two Step Approach





Two Step Approach





Two Step Approach



Two Step Approach



Two Step Approach



Fobos: Two Step Approach

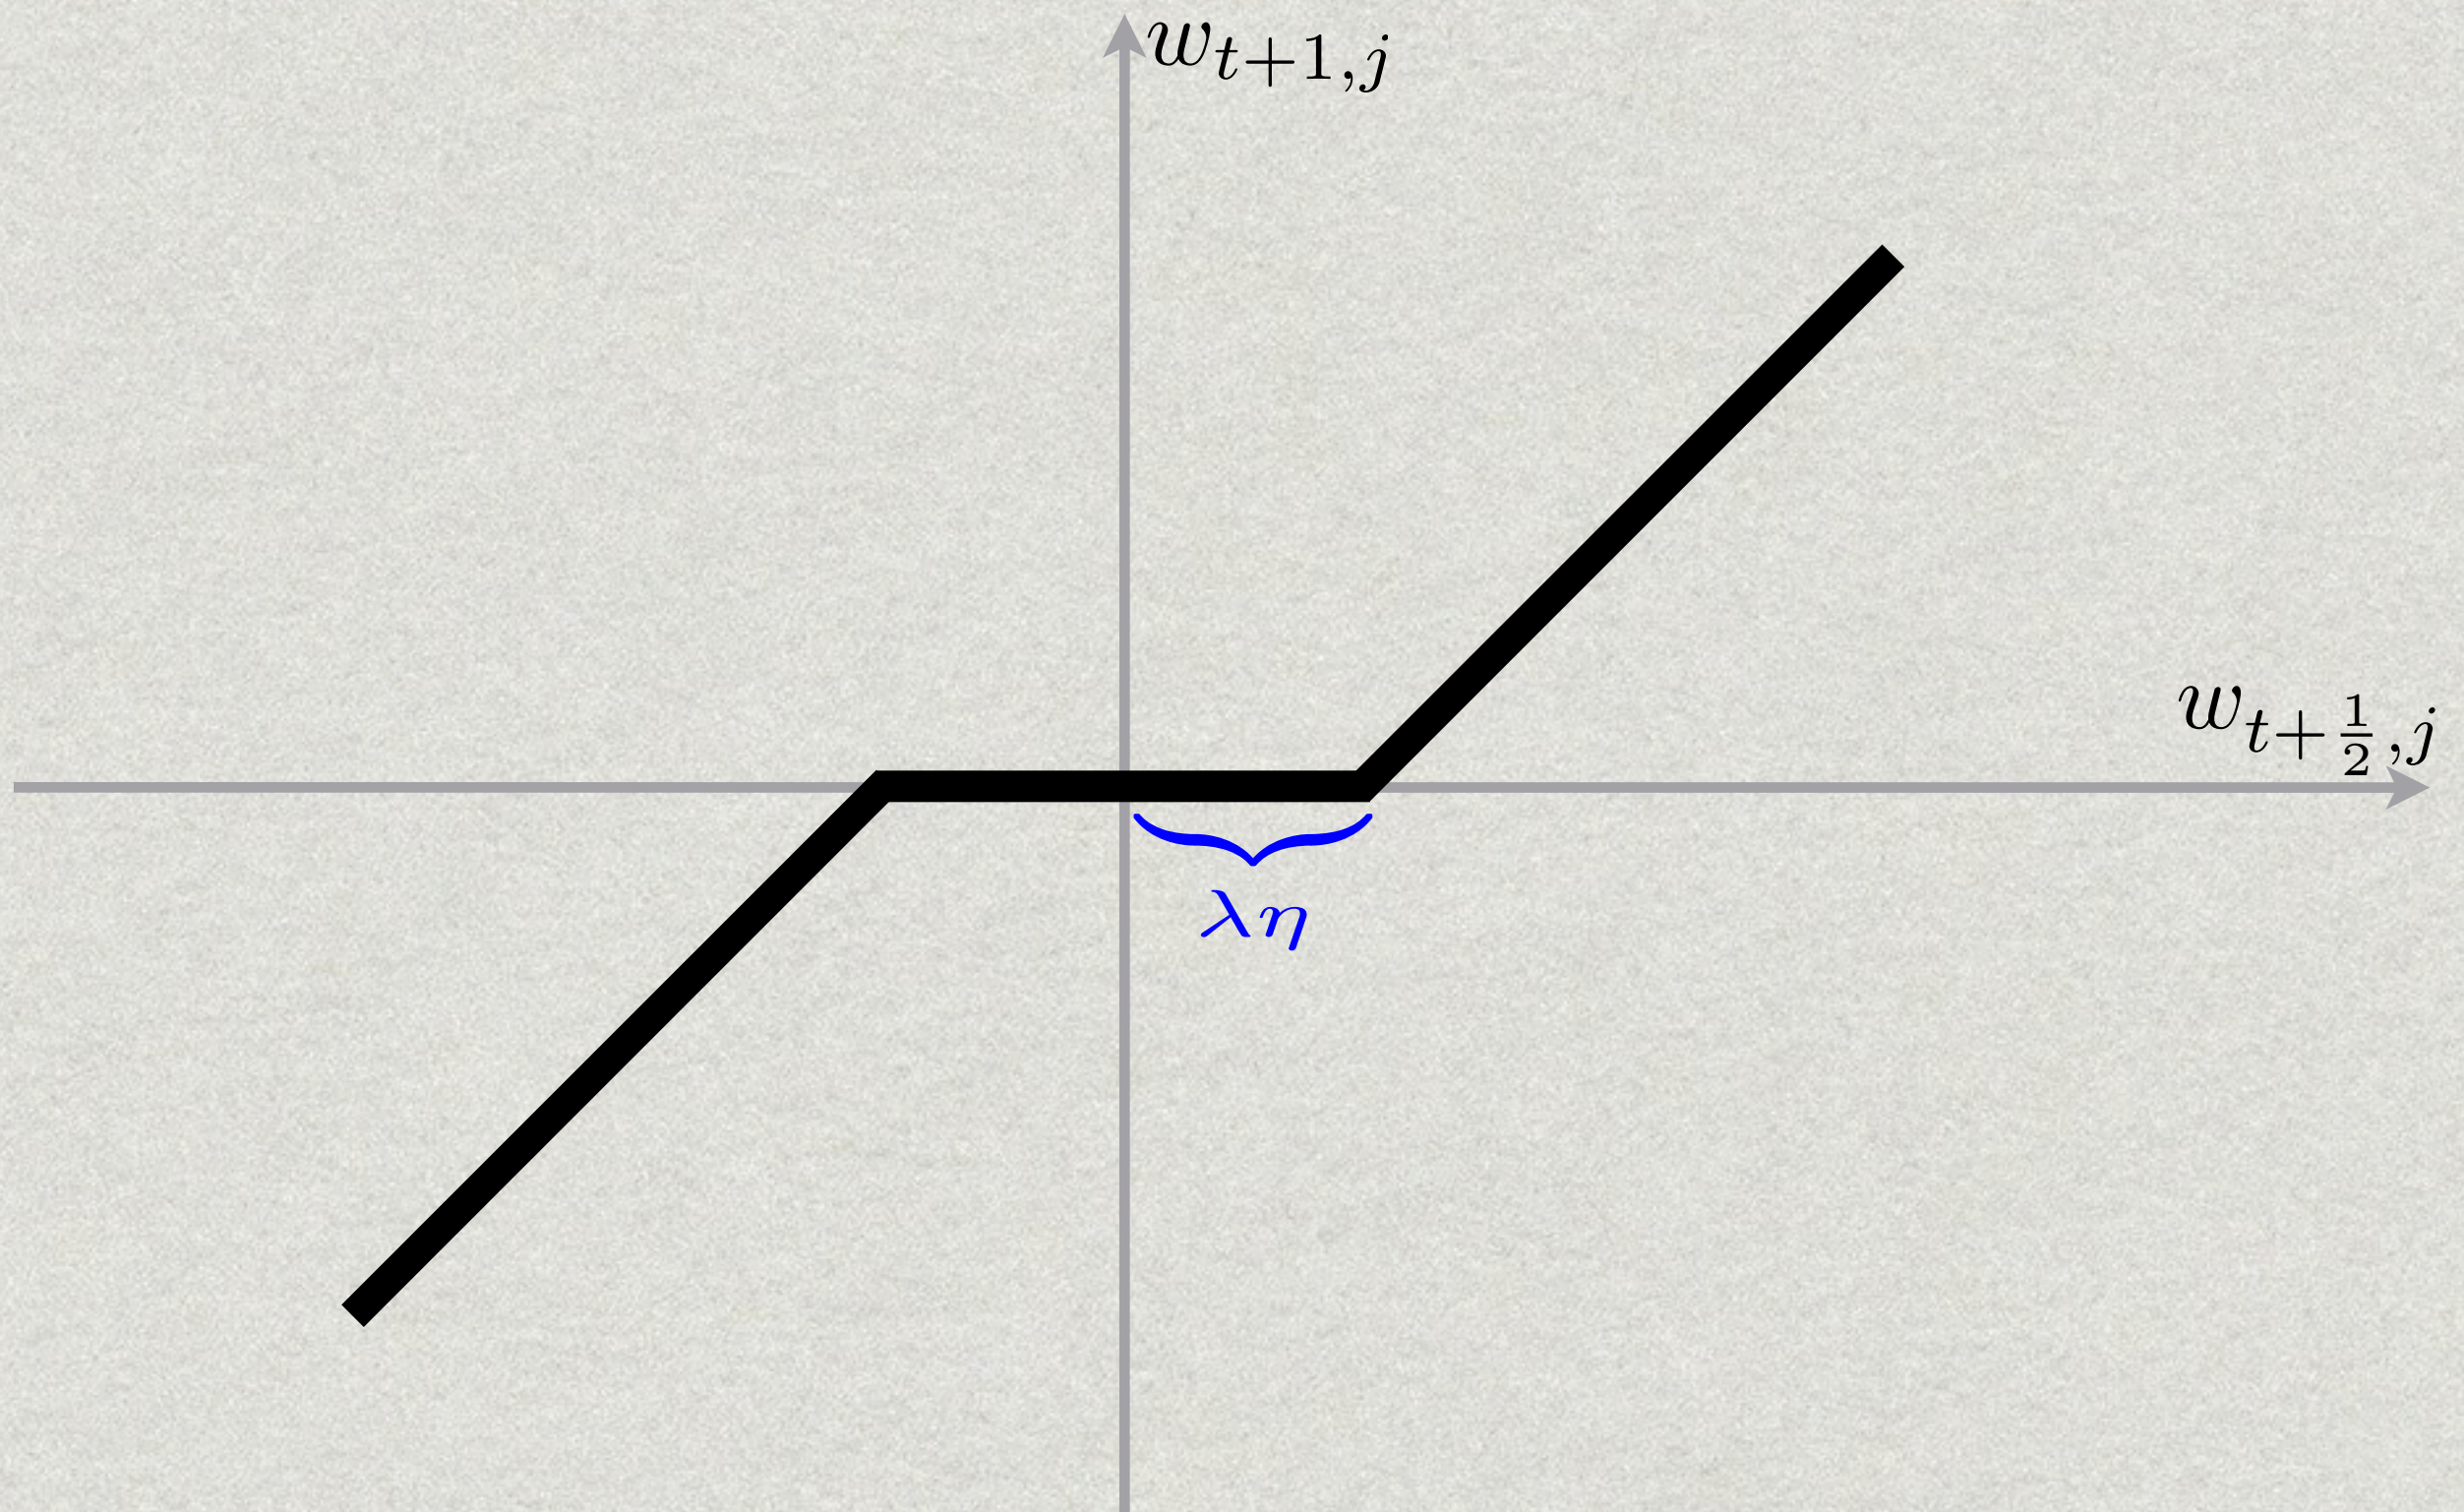
(1) Unconstrained stochastic gradient of loss

$$\mathbf{w}_{t+\frac{1}{2}} = \mathbf{w}_t - \eta \mathbf{g}_t$$


(2) Incorporate regularization and solve

$$\mathbf{w}_{t+1} = \operatorname{argmin}_{\mathbf{w}} \left\{ \frac{1}{2} \left\| \mathbf{w} - \mathbf{w}_{t+\frac{1}{2}} \right\|^2 + \eta \lambda R(\mathbf{w}) \right\}$$

Fobos for L_1 Regularization



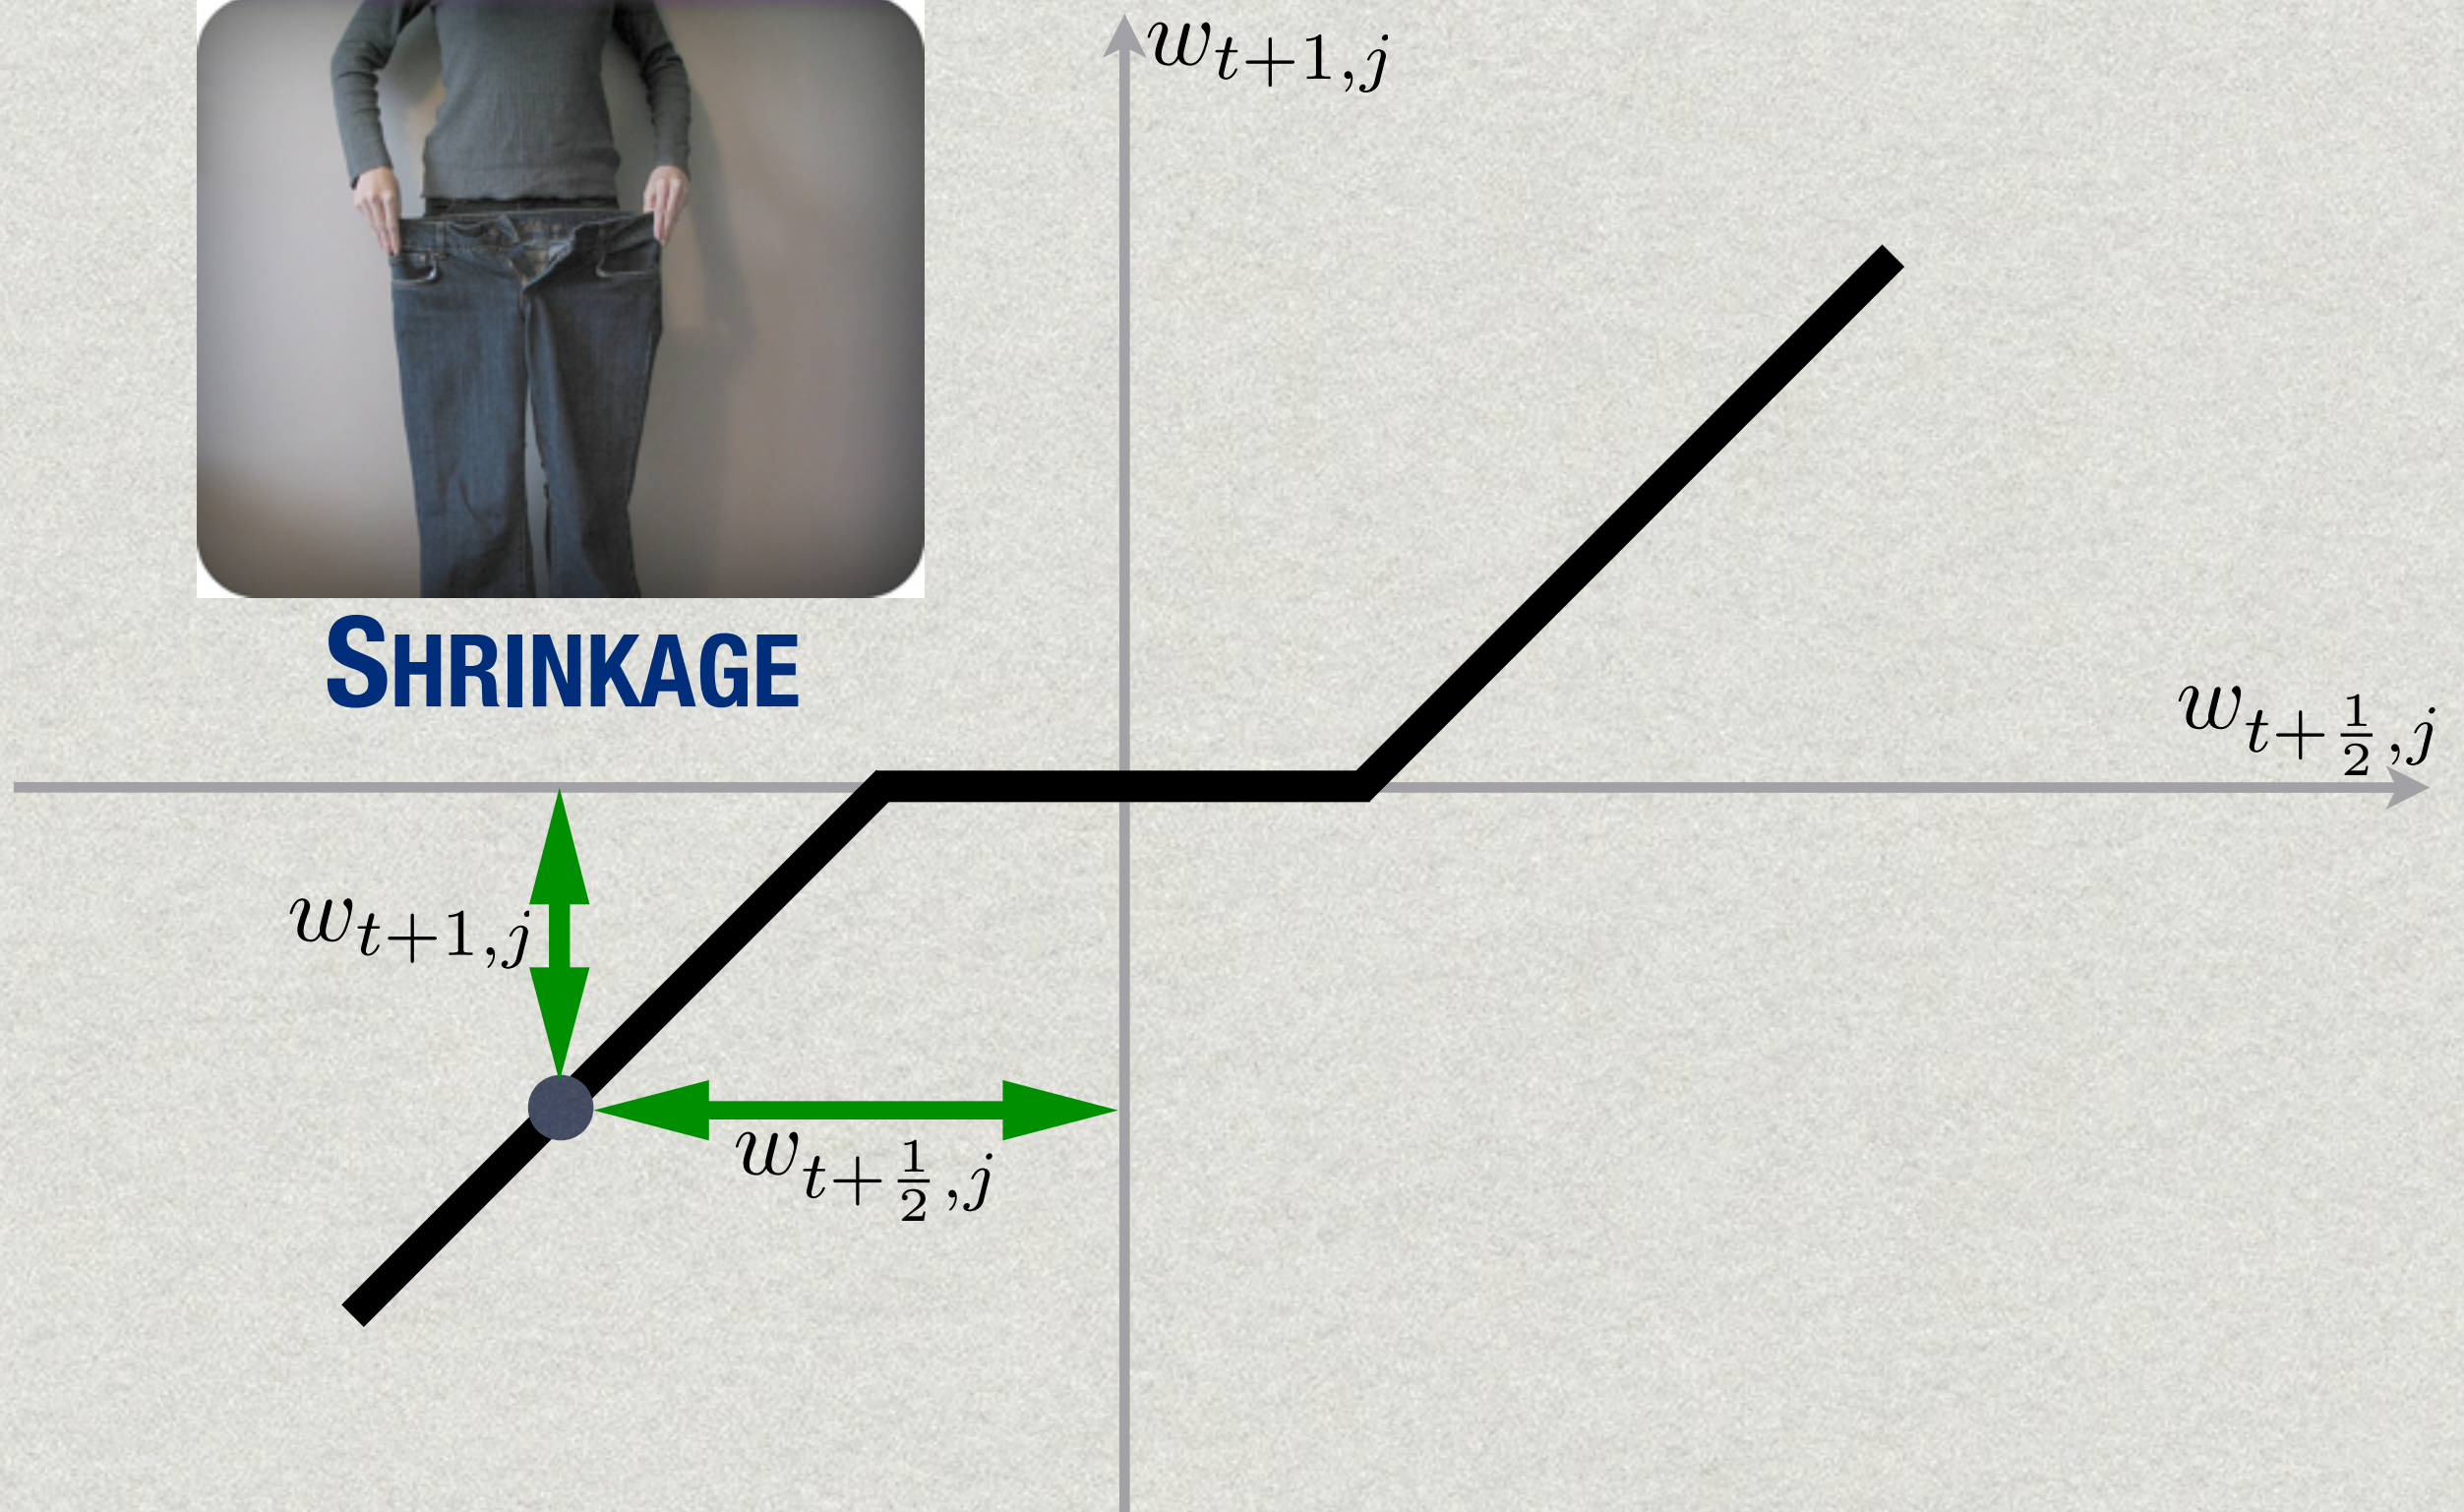
Fobos for L₁ Regularization




Fobos for L₁ Regularization



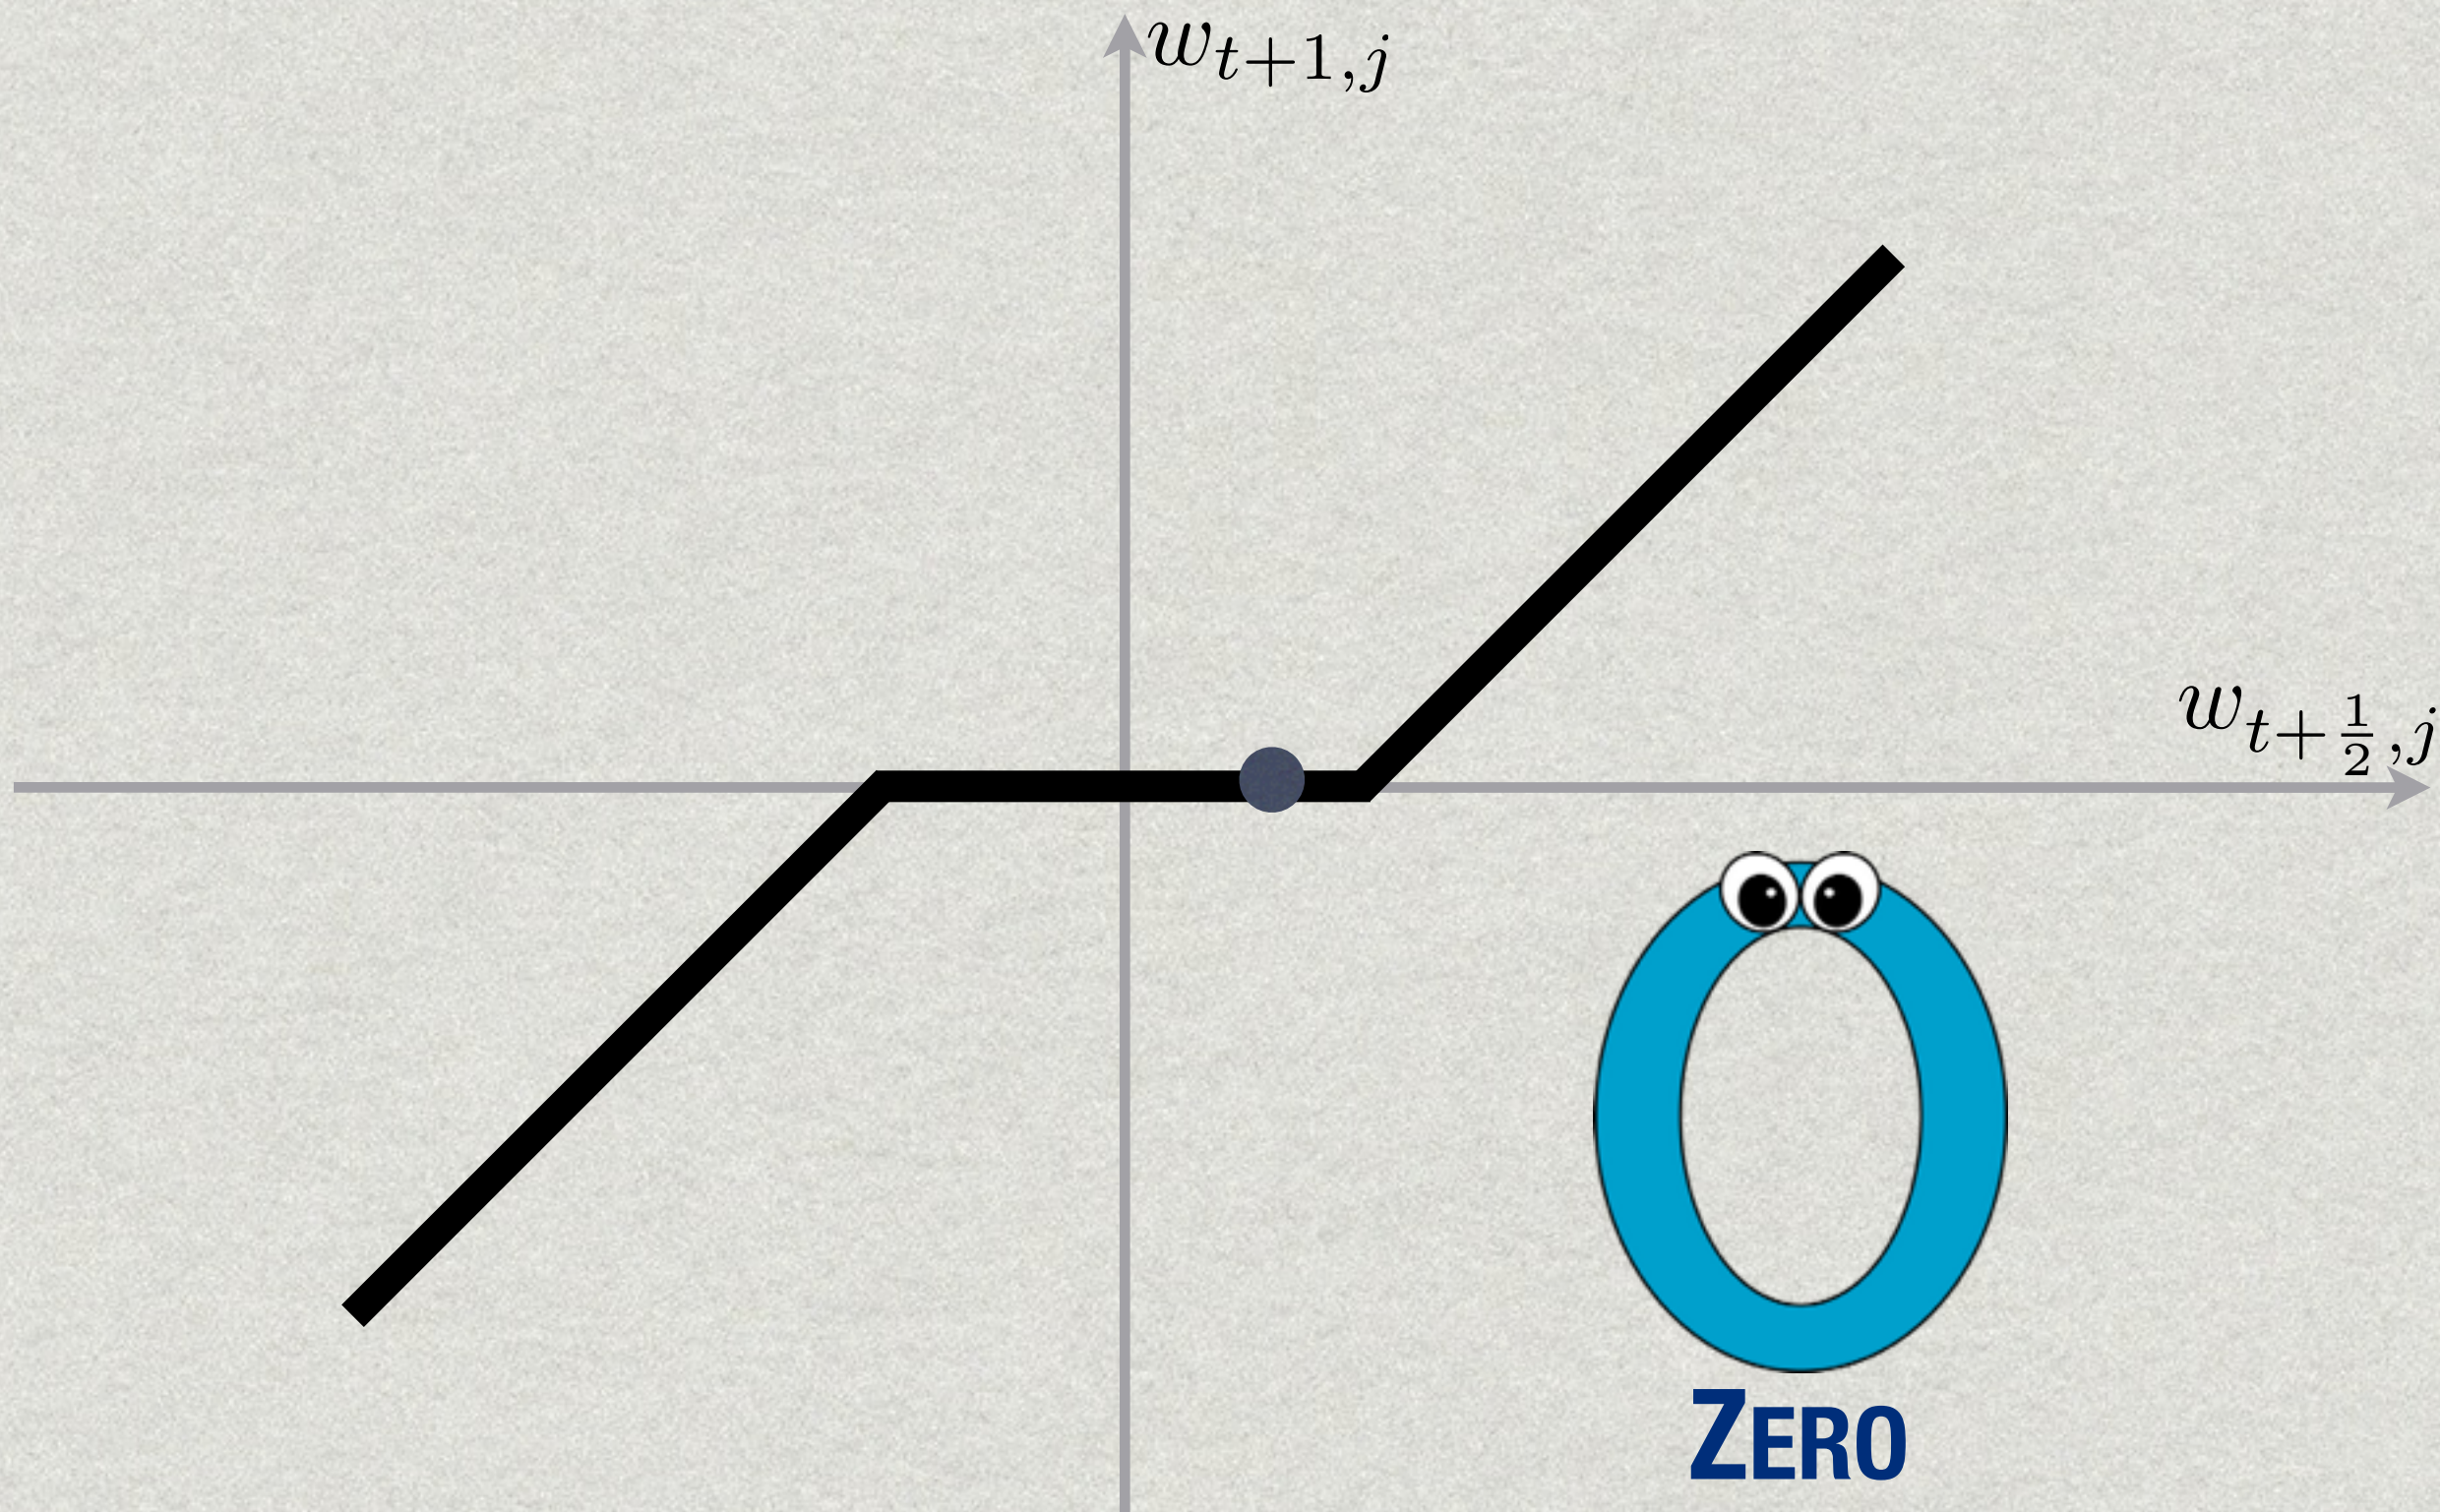
SHRINKAGE



Fobos for L₁ Regularization



Fobos for L_1 Regularization



Forward Looking Subgradient

- The optimum (\mathbf{w}_{t+1}) satisfies

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta \mathbf{g}_t^L - \eta \lambda \mathbf{g}_{t+1}^R$$

Forward Looking Subgradient

- The optimum (\mathbf{w}_{t+1}) satisfies

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta \mathbf{g}_t^L - \eta \lambda \mathbf{g}_{t+1}^R$$

**CURRENT GRADIENT
OF EMPIRICAL LOSS**



Forward Looking Subgradient

- The optimum (\mathbf{w}_{t+1}) satisfies

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta \mathbf{g}_t^L - \eta \lambda \mathbf{g}_{t+1}^R$$

**CURRENT GRADIENT
OF EMPIRICAL LOSS**



**FORWARD SUBGRADIENT
OF REGULARIZATION**



Forward Looking Subgradient

- The optimum (\mathbf{w}_{t+1}) satisfies

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta \mathbf{g}_t^L - \eta \lambda \mathbf{g}_{t+1}^R$$

**CURRENT GRADIENT
OF EMPIRICAL LOSS**



**FORWARD SUBGRADIENT
OF REGULARIZATION**




Yields very simple alternative analysis, in particular convergence to the optimum at a rate of

$$O\left(\frac{1}{\sqrt{T}}\right) \text{ or } O\left(\frac{\log(T)}{T}\right)$$

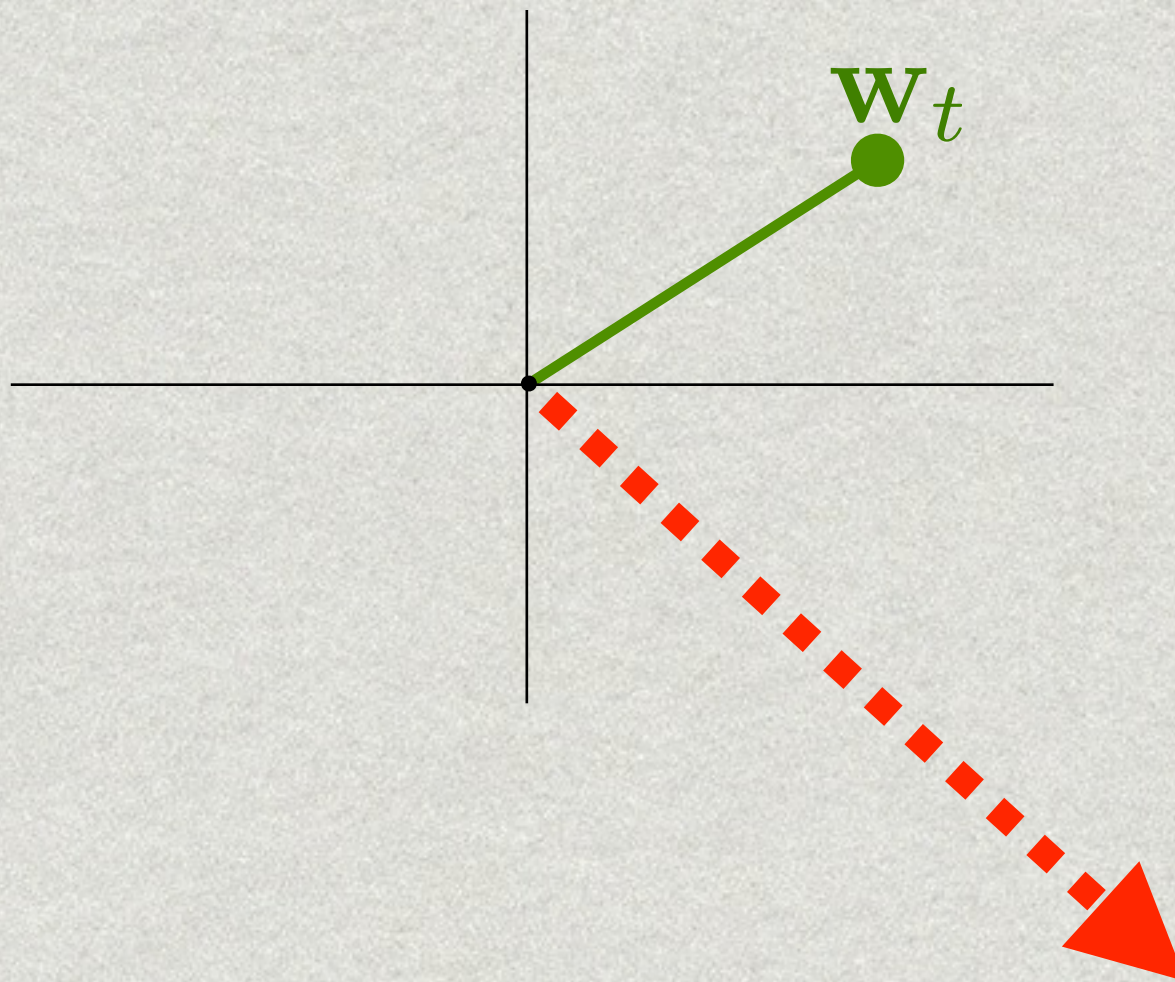
Proximal Operators




Proximal Operators




Proximal Operators



Proximal Operators




Proximal Operators



Cast a tradeoff:

- Maintaining proximity to weight vector
- Following the steepest descent direction

Proximal Operators



Cast a tradeoff:

- Maintaining proximity to weight vector
- Following the steepest descent direction


Proximal Operators



Cast a tradeoff:

- Maintaining proximity to weight vector
- Following the steepest descent direction

Proximal Operators



$$\mathbf{w}_{t+1} = \arg \min_{\mathbf{w} \in \Omega} \frac{1}{\eta} D(\mathbf{w} \parallel \mathbf{w}_{t+1}) + \langle \mathbf{w}, \mathbf{g}_t \rangle$$




proximity



gradient

Proximal Operators



$$\mathbf{w}_{t+1} = \arg \min_{\mathbf{w} \in \Omega} \frac{1}{\eta} D(\mathbf{w} \parallel \mathbf{w}_t \times) + \langle \mathbf{w}, \mathbf{g}_t \rangle$$

MY BAD!

proximity

gradient

Fobos & Proximal Operators

- Uses decomposition of the objective into an empirical risk minimization term and a regularization term
- Uses the squared Euclidian norm for proximity

$$\mathbf{w}_{t+1} = \arg \min_{\mathbf{w} \in \Omega} \frac{1}{2\eta} \|\mathbf{w} - \mathbf{w}_t\|^2 + \langle \mathbf{w}, \mathbf{g}_t \rangle + \lambda \|\mathbf{w}\|_1$$

EG & Proximal Operators

- If we constrain \mathbf{w} to the probability simplex, use relative entropy, we get Exponentiated Gradient (EG)

$$\mathbf{w}^{t+1} = \arg \min_{\mathbf{w} \in \Delta} \frac{1}{\eta} D_{\text{KL}}(\mathbf{w} \parallel \mathbf{w}^t) + \langle \mathbf{w}, \mathbf{g}^t \rangle$$

$$\mathbf{w}^{t+1} = \arg \min_{\mathbf{w} \in \Delta} \sum_{j=1}^d w_j \left(\log \left(\frac{w_j}{w_j^t} \right) + \eta g_j^t \right)$$

$$w_j^{t+1} = \frac{1}{Z} w_j^t \exp(-\eta g_j^t)$$

$$\text{where } Z = \sum_{l=1}^d w_l^t \exp(-\eta g_l^t)$$

High Dim Data \Leftrightarrow Sparse Gradients

	g	g	g	g
t=1	0	1.2	0	5.4
t=2	2	0	1.8	0
t=3	0	0	1.5	0
t=4	0	0	0	2
t=5	4.1	0	0	2
t=6	0	2.4	3.5	4
...				

For an efficient implementation computation should:

- Scale with the number of non-zeros
- Not with the full dimension

The following lemma to the rescue:

$$\mathcal{P}.1 : \quad \mathbf{w}_t = \arg \min_{\mathbf{w}} \frac{1}{2} \|\mathbf{w} - \mathbf{w}_{t-1}\|^2 + \lambda_t \|\mathbf{w}\|_q$$

$$\mathcal{P}.2 : \quad \mathbf{w}^* = \arg \min_{\mathbf{w}} \frac{1}{2} \|\mathbf{w} - \mathbf{w}_0\|^2 + \left(\sum_{t=1}^T \lambda_t \right) \|\mathbf{w}\|_q$$

$$T \times \mathcal{P}.1 \equiv \mathcal{P}.2 \quad q \in \{1, 2, \infty\}$$

$$\mathbf{w}_T (\mathcal{P}.1) = \mathbf{w}^* (\mathcal{P}.2)$$

Efficient High Dimensional Update

	g	g	g	g
t=1	0	1.2	0	5.4
t=2	2	0	1.8	0
t=3	0	0	1.5	0
t=4	0	0	0	2
t=5	4.1	0	0	2
t=6	0	2.4	3.5	4
...				

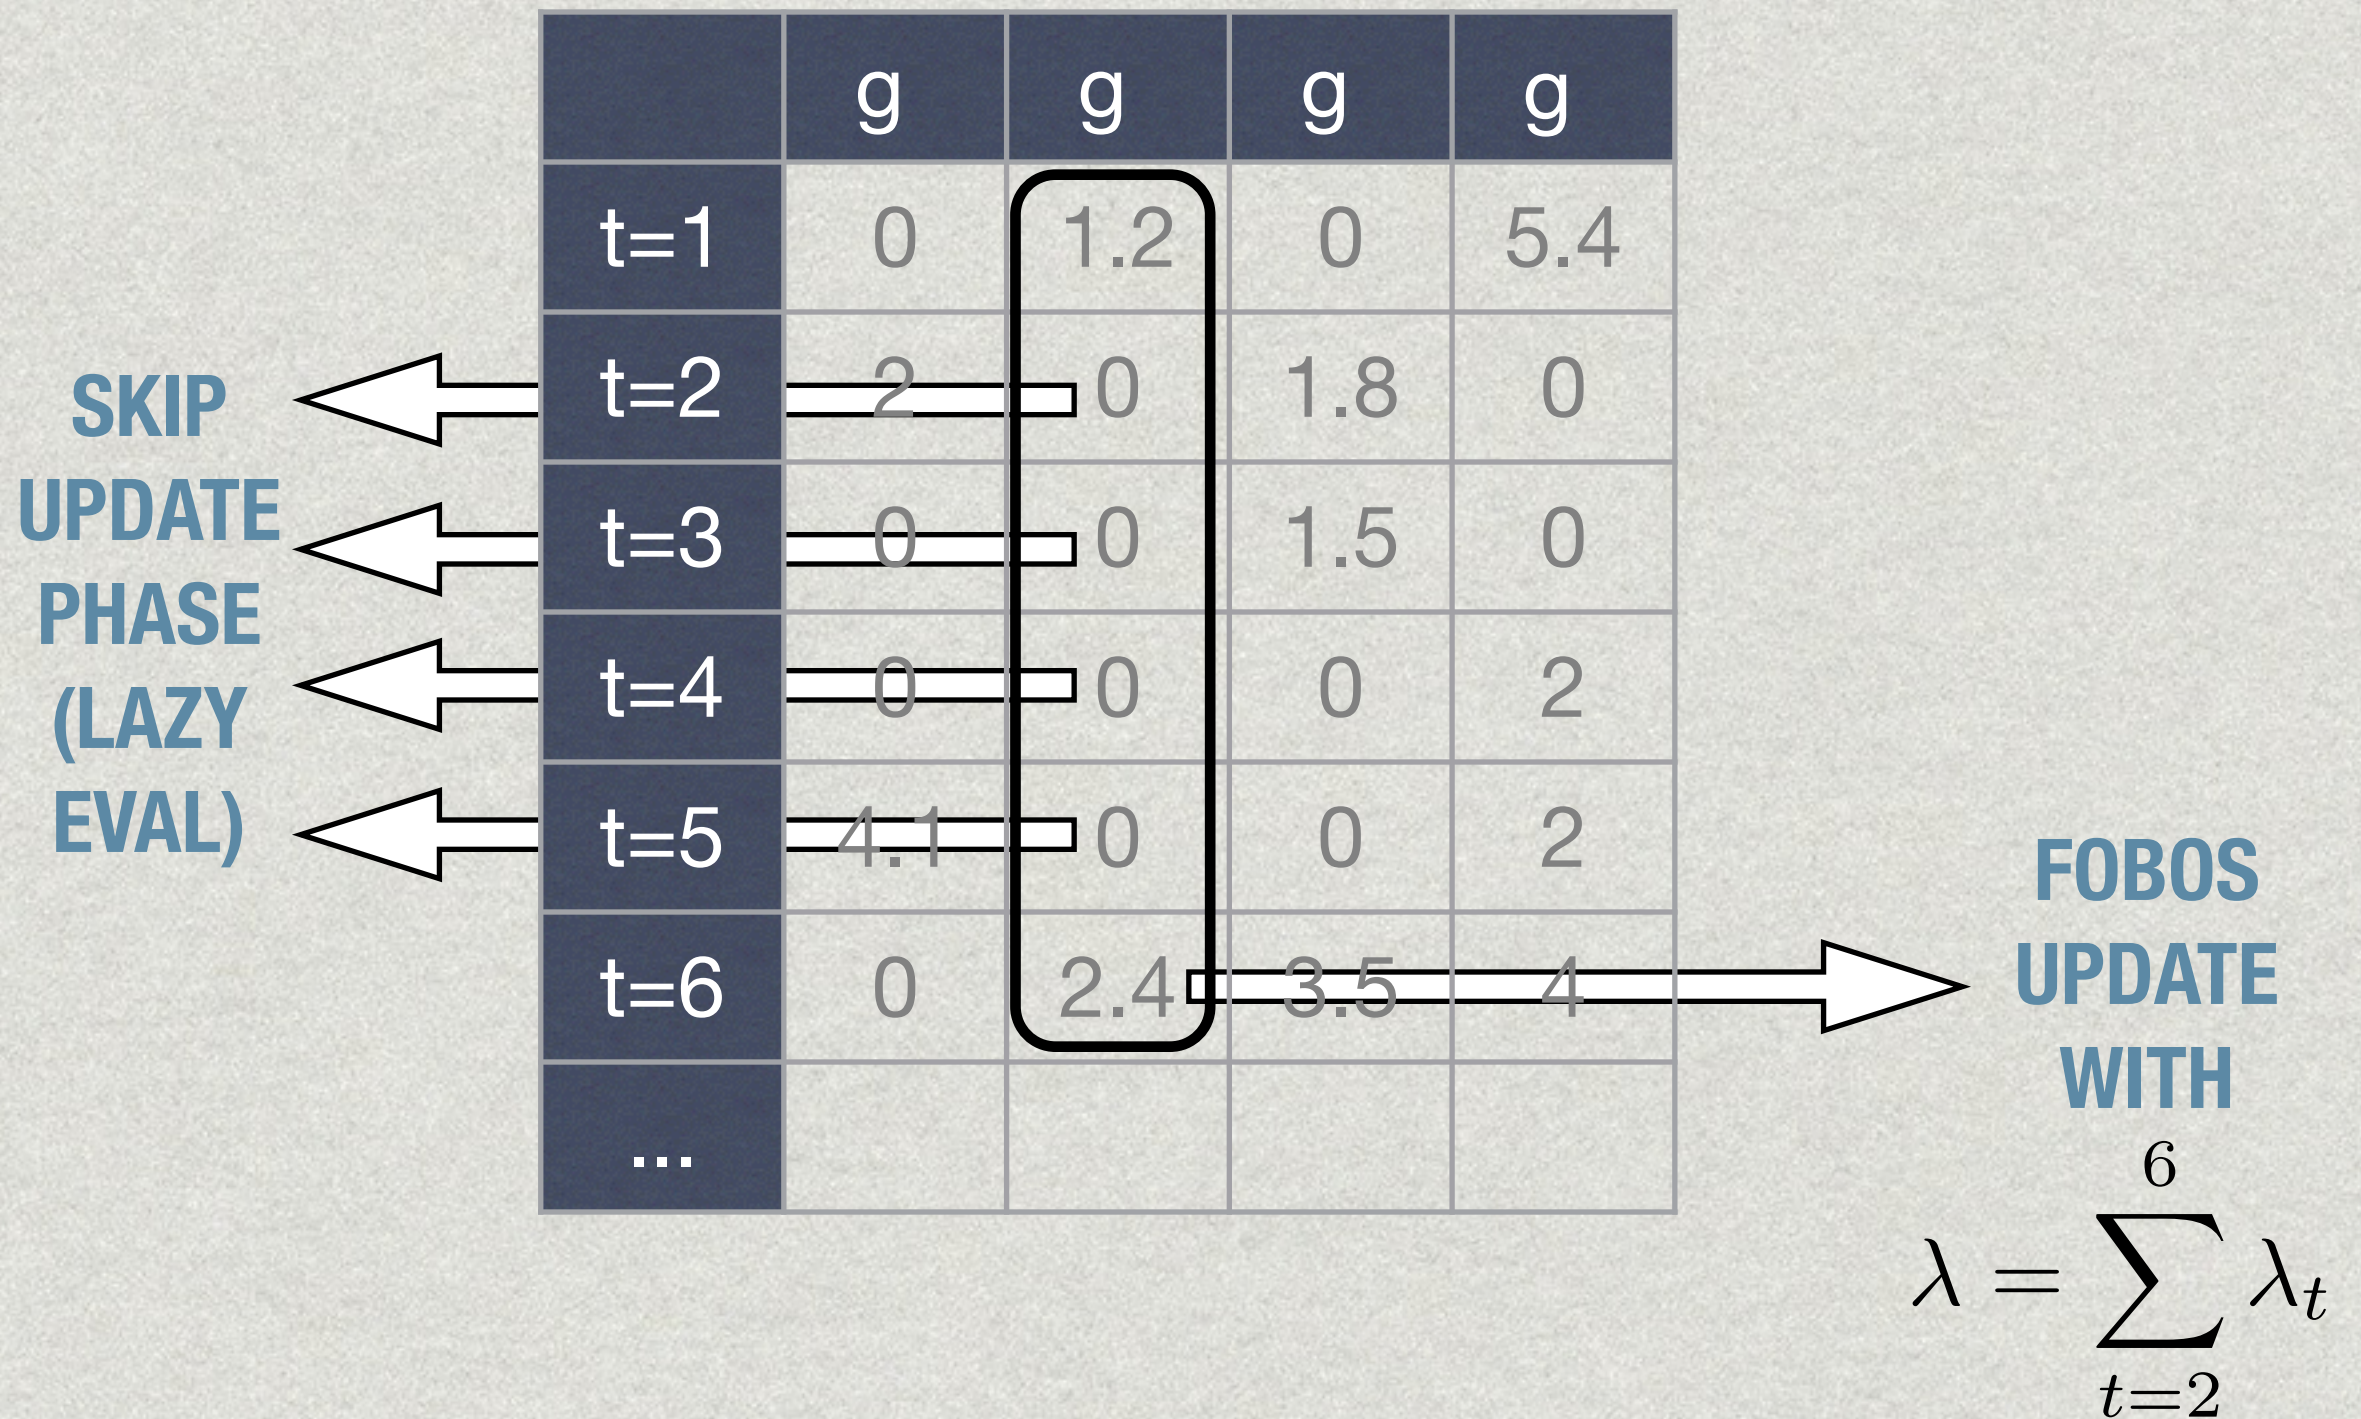
Efficient High Dimensional Update

	g	g	g	g
t=1	0	1.2	0	5.4
t=2	2	0	1.8	0
t=3	0	0	1.5	0
t=4	0	0	0	2
t=5	4.1	0	0	2
t=6	0	2.4	3.5	4
...				


Efficient High Dimensional Update

		g	g	g	g
	t=1	0	1.2	0	5.4
SKIP UPDATE PHASE (LAZY EVAL)	t=2	2	0	1.8	0
	t=3	0	0	1.5	0
	t=4	0	0	0	2
	t=5	4.1	0	0	2
	t=6	0	2.4	3.5	4
	...				

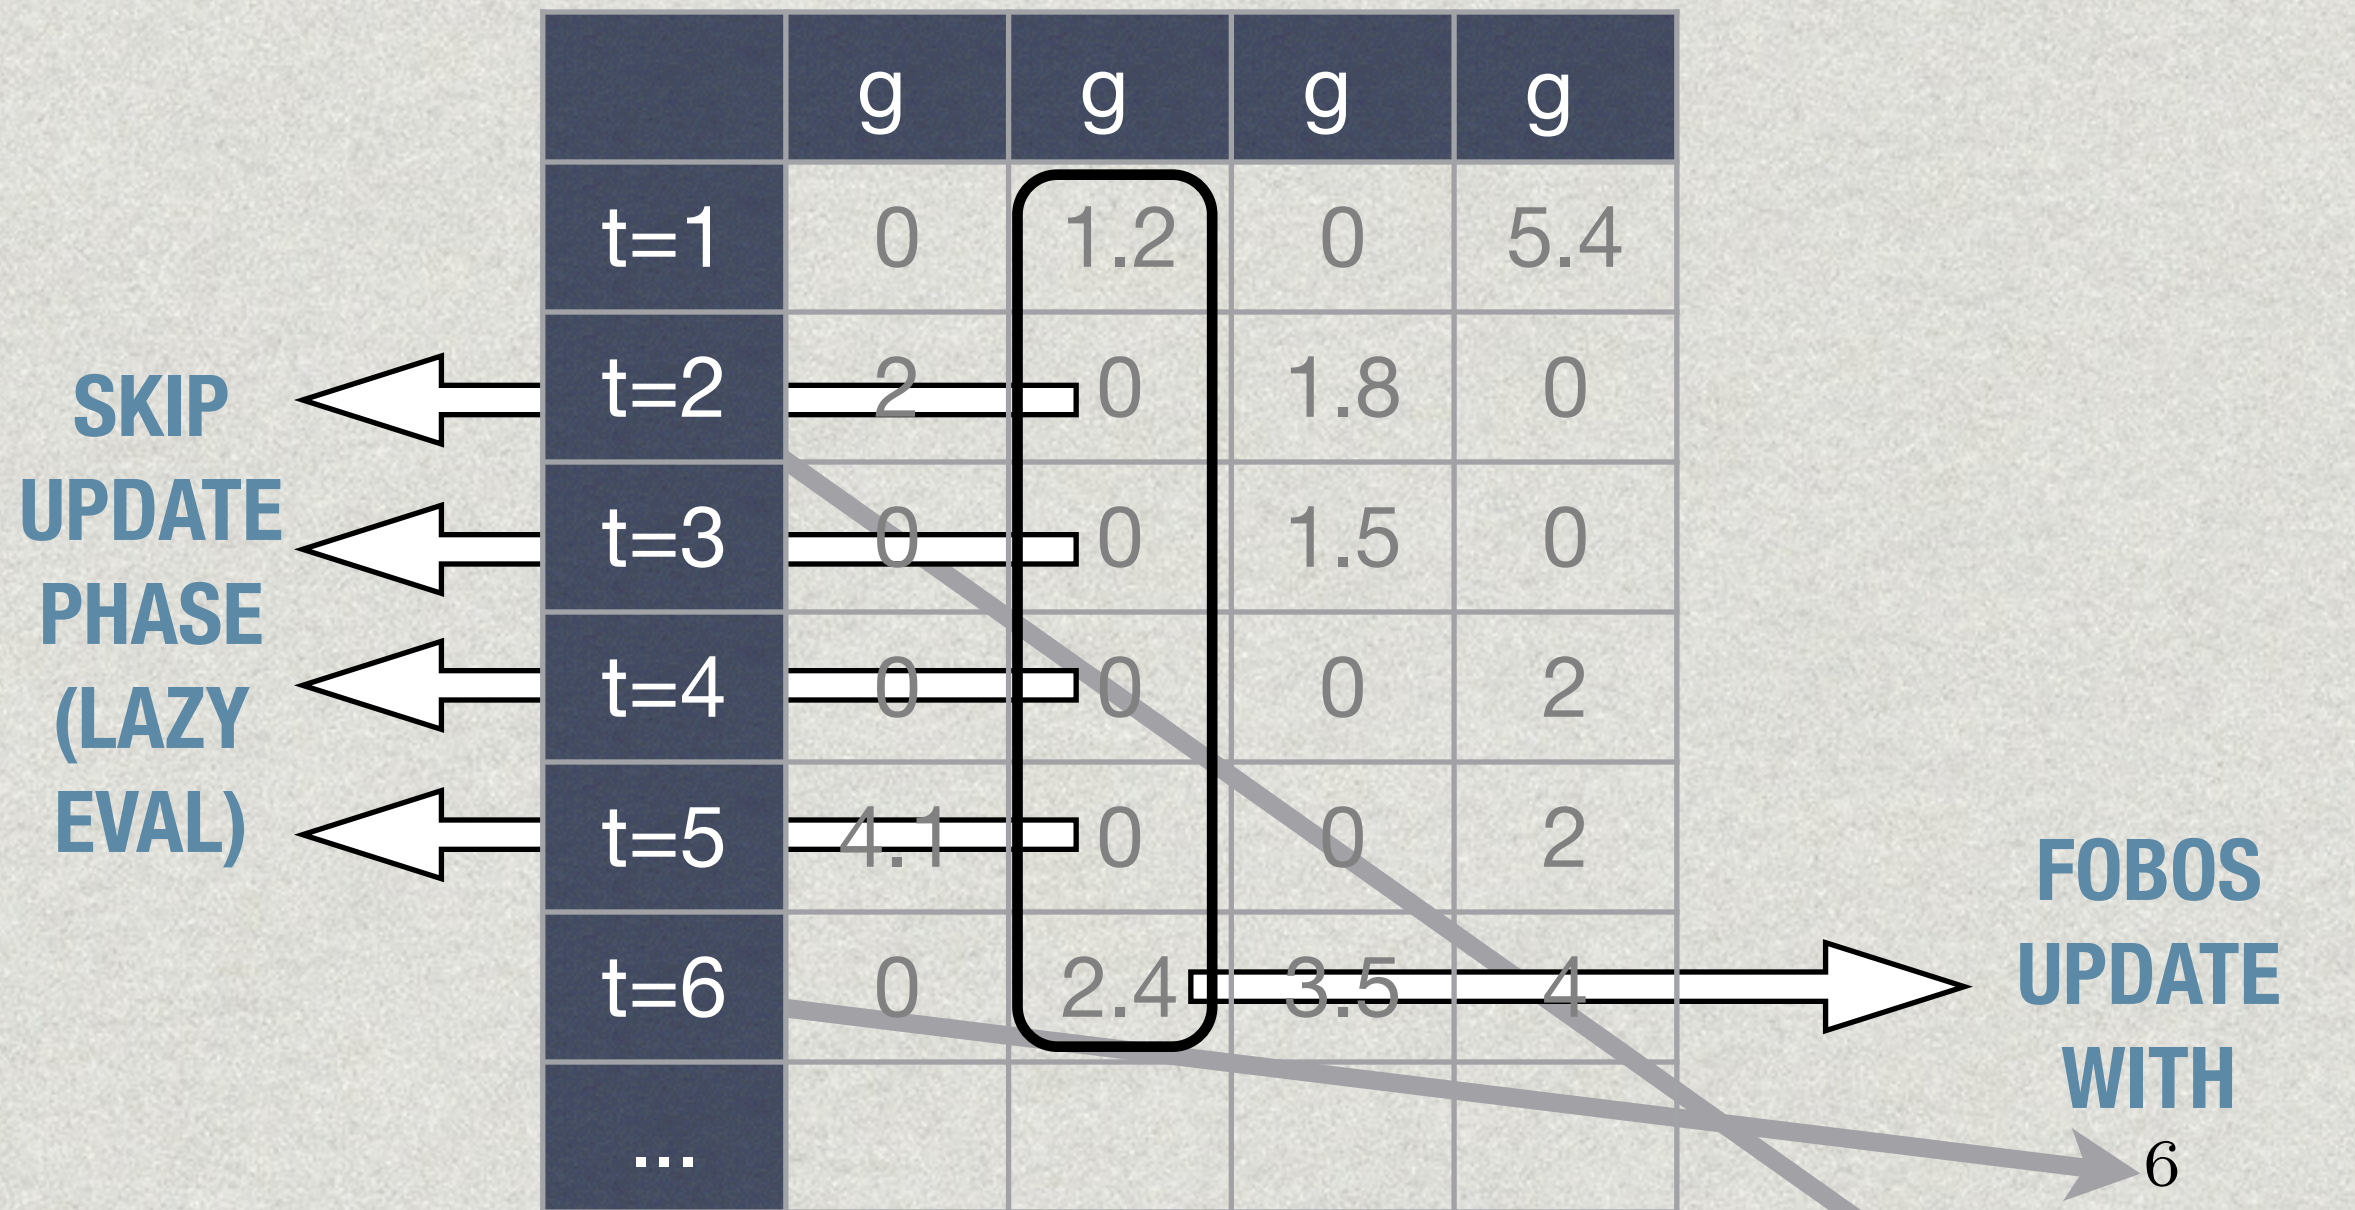
Efficient High Dimensional Update



Efficient High Dimensional Update



Efficient High Dimensional Update



Just-in-time update for each new sample:

- Accumulated proximal update
- Stochastic gradient step (w/o further ops)

$$\lambda = \sum_{t=2}^6 \lambda_t$$