Statistical Learning Theory: Expert Learning

CS4780/5780 – Machine Learning
Fall 2014

Thorsten Joachims
Cornell University

Reading: Mitchell Chapter 7.5

Generalization Error Bound: Infinite H, Non-Zero Error

- Setting
  - Sample of n labeled instances *S*
  - Learning Algorithm *L* using a hypothesis space *H* with *VCDim(H)=d*
  - *L* returns hypothesis *ĥ=L(S)* with lowest training error
- Definition: The VC-Dimension of H is equal to the maximum number d of examples that can be split into two sets in all 2^d ways using functions from H (shattering).
- Given hypothesis space *H* with *VCDim(H)* equal to *d* and an i.i.d. sample *S* of size *n*, with probability (1-δ) it holds that

$$Err_P(h_{\mathcal{L}(S)}) \leq Err_S(h_{\mathcal{L}(S)}) + \sqrt{\frac{d\left(\ln\left(\frac{2n}{d}\right)+1\right) - \ln\left(\frac{\delta}{4}\right)}{n}}$$

Outline

- Online learning
- Review of perceptron and mistake bound
- Expert model
  - Halving Algorithm
  - Weighted Majority Algorithm
  - Exponentiated Gradient Algorithm
- Bandit model
  - EXP3 Algorithm

Online Classification Model

- Setting
  - Classification
  - Hypothesis space H with h: X→Y
  - Measure misclassifications (i.e. zero/one loss)
- Interaction Model
  - Initialize hypothesis $h \in H$
  - FOR t from 1 to T
    - Receive x_t
    - Make prediction $\hat{y}_t = h(x_t)$
    - Receive true label y_t
    - Record if prediction was correct (e.g., $\hat{y}_t = y_t$)
    - Update h

(Online) Perceptron Algorithm

- Input: $S = ((\vec{x}_1, y_1), ..., (\vec{x}_n, y_n))$, $\vec{x}_i \in \Re^N$, $y_i \in \{-1, 1\}$
- Algorithm:
  - $\vec{w}_0 = \vec{0}$, $k = 0$
  - FOR i=1 TO n
    - IF $y_i(\vec{w}_k \cdot \vec{x}_i) \leq 0$ ### makes mistake
      - $\vec{w}_{k+1} = \vec{w}_k + y_i \vec{x}_i$
      - $k = k + 1$
    - ENDIF
  - ENDFOR
- Output: $\vec{w}_k$

Perceptron Mistake Bound

Theorem: For any sequence of training examples $S = ((\vec{x}_1, y_1), \dots, (\vec{x}_n, y_n))$ with

$$R = \max \|\vec{x}_i\|,$$

if there exists a weight vector $\vec{w}_{opt}$ with $\|\vec{w}_{opt}\| = 1$ and

$$y_i \left(\vec{w}_{opt} \cdot \vec{x}_i\right) \geq \delta$$

for all $1 \leq i \leq n$, then the Perceptron makes at most

$$\frac{R^2}{\delta^2}$$

errors.

Expert Learning Model

- Setting
  - N experts named $\mathrm{H} = \{h_1, \dots, h_N\}$
  - Each expert h_i takes an action $y = h_i(x_t)$ in each round t and incurs loss $\Delta_{t,i}$
  - Algorithm can select which expert's action to follow in each round
- Interaction Model
  - FOR t from 1 to T
    - Algorithm selects expert h_{i_t} according to strategy A_{w_t} and follows its action y
    - Experts incur losses $\Delta_{t,1} \dots \Delta_{t,N}$
    - Algorithm incurs loss Δ_{t,i_t}
    - Algorithm updates w_t to w_{t+1} based on $\Delta_{t,1} \dots \Delta_{t,N}$

Halving Algorithm

- Setting
  - N experts named $\mathrm{H} = \{h_1, \dots, h_N\}$
  - Binary actions $y = \{+1, -1\}$ given input x, zero/one loss
  - Perfect expert exists in H
- Algorithm
  - $VS_1 = H$
  - FOR t = 1 TO T
    - Predict the same y as majority of $h_i \in VS_t$
    - $VS_{t+1} = VS_t$ minus those $h_i \in VS_t$ that were wrong
- Mistake Bound
  - How many mistakes can the Halving algorithm make before predicting perfectly?