8. Supervised learning, MaxMin estimate

Given training data $\mathcal{I}_\ell = \{(x^j, s^j) \mid j = 1, .., \ell\}$ solve

$$\vec{u}_* = \operatorname*{argmax}_{\vec{u}} \min_j \log p_{\vec{u}}(x^j, s^j)$$

We write the task as a convex optimisation task

$$\log Z(\vec{u}) - c \to \min_{\vec{u}, c}$$

$$\text{s.t.} \quad c - \langle \vec{\varphi}^j, \vec{u} \rangle \leq 0 \quad \forall j = 1, .., \ell \tag{1}$$

where $\vec{\varphi}^j = \vec{\varphi}(x^j, s^j)$

Construct the Lagrange function and the dual task

$$L(c, \vec{u}, \vec{\lambda}) = \log Z(\vec{u}) - c + \sum_{j=1}^{\ell} \lambda_j \left[c - \langle \vec{\varphi}^j, \vec{u} \rangle \right]$$

$$\max_{\vec{\lambda} \geq 0} \min_{\vec{u}, c} L(c, \vec{u}, \vec{\lambda}) =$$

$$\max_{\vec{\lambda} \in \Delta^{\ell-1}} \underbrace{\min_{\vec{u}} \left[\log Z(\vec{u}) - \sum_j \lambda_j \langle \vec{\varphi}^j, \vec{u} \rangle \right]}_{\text{MLE}}$$

where $\Delta^{\ell-1}$ denotes the simplex $\Delta^{\ell-1} = \{\vec{\lambda} \in \mathbb{R}^\ell_+ \mid \sum_j \lambda_j = 1\}$

Algorithm Choose arbitrary $\vec{\lambda}^{(0)} \in \Delta^{n-1}$ e.g. $\lambda_j^{(0)} = \frac{1}{\ell}$

Iterate:

1) Solve the MLE task (see sec. 7) for $\vec{\lambda}^{(t)} \to \vec{u}^{(t)}$

2) find $i \in \operatorname*{argmin}_j p_{\vec{u}^{(t)}}(x^j, s^j) = \operatorname*{argmin}_j \langle \vec{\varphi}(x^j, s^j), \vec{u}^{(t)} \rangle$

set $\vec{\lambda}^{(t+1)} \sim (t+1)\vec{\lambda}^{(t)} + \vec{e}_i$

until

$$\sum_{j=1}^{\ell} \lambda_j \langle \vec{\varphi}^j, \vec{u} \rangle - \min_j \langle \vec{\varphi}^j, \vec{u} \rangle \leq \varepsilon$$

Theorem 1 (w/o proof)

The algorithm stops after a finite number of iterations

The obtained model $\vec{u}_*$ is ε-optimal, i.e.

$$\min_j \log p_{\vec{u}_0}(x^j, s^j) - \min_j \log p_{\vec{u}_*}(x^j, s^j) \leq \varepsilon,$$

where $\vec{u}_0$ is the optimal model.

9. Empirical risk minimisation

Given: • training data $\mathcal{T}_\ell = \{(x^j, s^j) \mid j = 1, .., \ell\}$

• loss function $C(s', s) = \mathbb{1}\{s' \neq s\}$, i.e.

Bayes optimal decision $q_{\vec{u}}(x) \in \operatorname{argmax}_{s \in K^n} p_{\vec{u}}(x, s)$

The task of empirical risk minimisation reads

$$\frac{1}{\ell} \sum_{j=1}^{\ell} \mathbb{1}\{s^j \neq q_{\vec{u}}(x^j)\} \to \min_{\vec{u}}$$

It is not tractable in general; the objective function is neither convex nor differentiable.

It is simple to solve if $\exists\, \vec{u}_*$ s.t.

$$s^j = \operatorname{argmax}_{s \in K^n} p_{\vec{u}_*}(x^j, s) \quad \forall\, j = 1, .., \ell$$

i.e. we must find a $\vec{u}$ s.t. inequalities

$$\langle \vec{\varphi}(x^j, s^j), \vec{u} \rangle > \langle \vec{\varphi}(x^j, s), \vec{u} \rangle \quad \forall\, s \neq s^j$$

hold for all $j = 1, .., \ell$

This task can be solved by the perceptron algorithm:

Start from arbitrary $\vec{u}$ and iterate:

- solve

$$\tilde{s}^j = \operatorname*{argmax}_{s \in K^n} \langle \vec{\varphi}(x^j, s), \vec{u} \rangle \qquad \text{(see sec. 4)}$$

$\forall\, j = 1, \dots, \ell$

- if for some j $\tilde{s}^j \neq s^j$, update $\vec{u}$ by

$$\vec{u} \to \vec{u} + \vec{\varphi}(x^j, s^j) - \vec{\varphi}(x^j, \tilde{s}^j)$$

Let us reconsider the general task

$$q_{\vec{u}}(x) \in \operatorname*{argmax}_{s \in K^n} \langle \vec{\varphi}(x, s), \vec{u} \rangle$$

$$\frac{1}{\ell} \sum_{j=1}^{\ell} \mathbb{1}\{ s^j \neq q_{\vec{u}}(x^j) \} \to \min_{\vec{u}}$$

Approximate the loss (as a function of $\vec{u}$) by a convex upper bound, e.g., as follows

$$\mathbb{1}\{ s' \neq q_{\vec{u}}(x) \} \leq \max_{s} \left\{ \mathbb{1}\{ s' \neq s \} - \langle \vec{\varphi}(x, s'), \vec{u} \rangle + \langle \vec{\varphi}(x, s), \vec{u} \rangle \right\}$$

The approx. task reads

$$\frac{1}{\ell} \sum_{j=1}^{\ell} \max_{s \in K^n} \left[\mathbb{1}\{ s' \neq s \} + \langle \vec{\varphi}(x^j, s) - \vec{\varphi}(x^j, s^j), \vec{u} \rangle \right] \to \min_{\vec{u}}$$

Solve it

- by subgradient descent or
- by cutting plane algorithm or
- ...