# Machine Teaching for Bayesian Learners in the Exponential Family 

Xiaojin Zhu

Department of Computer Sciences, University of Wisconsin-Madison (jerryzhu@cs.wisc.edu)

Machine Teaching
Machine teaching: finding the best training set.


- World: test items $x \stackrel{i i d}{\sim} p\left(x \mid \theta^{*}\right)$.
- Learner: hypothesis space $\Theta$
- Teacher: knows $\theta^{*}, \Theta$, learning algorithm, teaches by creating a training set $\mathcal{D}$.


## Optimal Teaching Key Idea

$\min _{\mathcal{D}} \operatorname{loss}\left(\widehat{f_{\mathcal{D}}}, \theta^{*}\right)+\operatorname{effort}(\mathcal{D})$

- effort() of the teacher/learner to work with $\mathcal{D}$. - Not regularized estimation: $\theta^{*}$ given.
- Hard combinatorial optimization
- Objective called Teaching Impedance $T I(\mathcal{D})$

Teaching Bayesian Learners

- Teacher knows learner prior $p_{0}(\theta)$ and likelihood $p(\mathcal{D} \mid \theta)$, can design non-iid $\mathcal{D}$
${ }^{*} \operatorname{loss}\left(\widehat{\mathcal{D}_{\mathcal{D}}}, \theta^{*}\right)=K L\left(\delta_{\theta^{*}} \| p(\theta \mid \mathcal{D})\right)$
- Teaching is to

$$
\min _{\mathcal{D}}-\log p\left(\theta^{*} \mid \mathcal{D}\right)+\operatorname{effort}(\mathcal{D}) .
$$

- Not MAP estimate! Still hard.

Teaching Bayesian Learners in the Exponential Family

## - Exponential family

$p(x \mid \theta)=h(x) \exp \left(\theta^{\top} T(x)-A(\theta)\right)$

- For $\mathcal{D}=\left\{x_{1}, \ldots, x_{n}\right\}$ the likelihood is

$$
p(\mathcal{D} \mid \theta)=\prod_{i=1}^{n} h\left(x_{i}\right) \exp \left(\theta^{\top} \mathbf{s}-A(\theta)\right)
$$

with aggregate sufficient statistics

$$
\mathbf{s} \equiv \sum_{i=1}^{n} T\left(x_{i}\right)
$$

Two-step algorithm:
© ifinding aggregate sufficient statistics
© unpacking

## Step 1: Sufficient Statistics

- Conjugate prior $p\left(\theta \mid \lambda_{1}, \lambda_{2}\right)=$

$$
h_{0}(\theta) \exp \left(\lambda_{1}^{\top} \theta-\lambda_{2} A(\theta)-A_{0}\left(\lambda_{1}, \lambda_{2}\right)\right)
$$

- $\mathcal{D}$ enters the posterior only via s and $n$ : $\exp \left(\left(\lambda_{1}+\mathbf{s}\right)^{\top} \theta-\left(\lambda_{2}+n\right) A(\theta)-A_{0}\left(\lambda_{1}+\mathbf{s}, \lambda_{2}+n\right)\right)$ - Optimal teaching problem

$$
\begin{aligned}
\min _{n, \mathbf{s}} & -\theta^{* \top}\left(\lambda_{1}+\mathbf{s}\right)+A\left(\theta^{*}\right)\left(\lambda_{2}+n\right) \\
& +A_{0}\left(\lambda_{1}+\mathbf{s}, \lambda_{2}+n\right)+\operatorname{effort}(n, \mathbf{s})
\end{aligned}
$$

- Convex relaxation: $n \in \mathbb{R}$ and $\mathbf{s} \in \mathbb{R}^{D}$
Step 2: Unpacking
(1) Round $n \leftarrow \max (0,[n])$
© Find $n$ teaching examples whose aggregate
sufficient statistics is approximately $\mathbf{s}$
- initialize $x_{i} \stackrel{i i d}{\sim} p\left(x \mid \theta^{*}\right), i=1 \ldots n$
- solve $\min _{x_{1}, \ldots, x_{n}}\left\|\mathbf{s}-\sum_{i=1}^{n} T\left(x_{i}\right)\right\|^{2}$ (nonconvex)

Some unpacking examples:

- Exponential dist $T(x)=x: x_{i}=\frac{\mathrm{s}}{n}$
- Poisson dist $T(x)=x$ (integers): rounding
- Gaussian dist $T(x)=\left(x, x^{2}\right), n=3, \mathbf{s}=(3,5)$ :
$\left\{x_{1}=0, x_{2}=1, x_{3}=2\right\}$ or
$\left\{x_{1}=\frac{1}{2}, x_{2}=\frac{5+\sqrt{13}}{4}, x_{3}=\frac{5-\sqrt{13}}{4}\right\}$


## Example 1

Teaching a 1D threshold classifier.

- Learner $p_{0}(\theta)=1, p(y=1 \mid x, \theta)=1_{x \geq \theta}$
- $p(\theta \mid \mathcal{D})$ uniform in $\left[\max _{i: y_{i}=-1}\left(x_{i}\right), \min _{i: y_{i}=1}\left(x_{i}\right)\right]$
- $\operatorname{effort}(\mathcal{D})=c|\mathcal{D}|$
- The optimal teaching problem becomes
$\min _{n,\left(x_{i}, y_{i}\right) \mathrm{In}}-\log \left(\frac{1}{\min _{i y_{i}=1}\left(x_{i}\right)-\max _{i y_{i}=-1}\left(x_{i}\right)}\right)+c n$.
- One solution: $\mathcal{D}=\left\{\left(\theta^{*}-\epsilon / 2,-1\right),\left(\theta^{*}+\epsilon / 2,1\right)\right\}$
as $\epsilon \rightarrow 0$ with $T I=\log (\epsilon)+2 c \rightarrow-\infty$

Example 2
Learner can't tell similar items

$$
\operatorname{effort}(\mathcal{D})=\frac{c}{\min _{x_{i}, x_{j} \in \mathcal{D}}\left|x_{i}-x_{j}\right|}
$$

- With $\mathcal{D}=\left\{\left(\theta^{*}-\epsilon / 2,-1\right),\left(\theta^{*}+\epsilon / 2,1\right)\right\}$, $T I=\log (\epsilon)+c / \epsilon$ with minimum at $\epsilon=c$. - $\mathcal{D}=\left\{\left(\theta^{*}-c / 2,-1\right),\left(\theta^{*}+c / 2,1\right)\right\}$

Example 3
Teaching to pick a Gaussian out of two
$-\Theta=\left\{\theta_{A}=N\left(-\frac{1}{4}, \frac{1}{2}\right), \theta_{B}=N\left(\frac{1}{4}, \frac{1}{2}\right)\right\}, \theta^{*}=\theta_{A}$, $p_{0}\left(\theta_{A}\right)=p_{0}\left(\theta_{B}\right)=\frac{1}{2}$

- $\operatorname{loss}(\mathcal{D})=\log \left(1+\prod_{i=1}^{n} \exp \left(x_{i}\right)\right)$ minimized by $x_{i} \rightarrow-\infty$, weird items.
- Box constraints $x_{i} \in[-d, d]$ :

$$
\min _{n, x_{1 n}} \log \left(1+\prod_{i=1}^{n} \exp \left(x_{i}\right)\right)+c n+\sum_{i=1}^{n} \mathbb{I}\left(\left|x_{i}\right| \leq d\right)
$$

- Solution: $n=\max \left(0,\left[\frac{1}{d} \log \left(\frac{d}{c}-1\right)\right]\right), x_{1: n}=-d$
- Note $n=0$ when $c \geq \frac{d}{2}$ : the effort of teaching outweighs the benefit. The teacher will choose not to teach, leaving learner with its prior $p_{0}$ !


## Example 4

Teaching the mean of a univariate Gaussian

- The world is $N\left(x ; \mu^{*}, \sigma^{2}\right)$
- Learner's prior $p_{0}(\mu)=N\left(\mu \mid \mu_{0}, \sigma_{0}^{2}\right)$, knows $\sigma^{2}$ - $T(x)=x$
- Aggregate sufficient statistics solution

$$
s=\frac{\sigma^{2}}{\sigma_{0}^{2}}\left(\mu^{*}-\mu_{0}\right)+\mu^{*} n
$$

Note $\frac{s}{n} \neq \mu^{*}$ : compensating for the learner's (wrong) prior belief $\mu_{0}$.

- $n$ is the solution to

$$
n-\frac{1}{2 \operatorname{effort}^{\prime}(n)}+\frac{\sigma^{2}}{\sigma_{0}^{2}}=0
$$

When effort $(n)=c n, n=\frac{1}{2 c}-\frac{\sigma^{2}}{\sigma_{0}^{2}}$

- Unpacking $s$ is trivial, e.g. $x_{1}=\ldots=x_{n}=s / n$
- Teacher will choose not to teach if the learner
initially had a "narrow mind": $\sigma_{0}^{2}<2 c \sigma^{2}$.

Example 5
Teaching a multinomial distribution.

$$
\begin{aligned}
\min _{\mathbf{s}} & -\log \Gamma\left(\sum_{k=1}^{K}\left(\beta_{k}+s_{k}\right)\right)+\sum_{k=1}^{K} \log \Gamma\left(\beta_{k}+s_{k}\right) \\
& -\sum_{k=1}^{K}\left(\beta_{k}+s_{k}-1\right) \log \pi_{k}^{*}+\operatorname{effort}(\mathbf{s})
\end{aligned}
$$

- Example: world $\pi^{*}=\left(\frac{1}{10}, \frac{3}{10}, \frac{6}{10}\right)$
- Learner "wrong" Dirichlet prior $\beta=(6,3,1)$

If effort(s) $=0$, "brute-force teaching"
$\mathbf{s}=(317,965,1933)$
If effort(s) $=0.3 \sum_{k=1}^{K} s_{k}$,

- $\mathbf{s}=(0,2,8), T I=2.65$.
- Not $\mathbf{s}=(1,3,6), T I=4.51$. doesn't correct prior
$\cdot \operatorname{Not} \mathbf{s}=(317,965,1933), T I=956.25$


## Example 6

Teaching a multivariate Gaussian.

* World $N\left(\mu^{*}=(\mathbf{0}, \mathbf{0}, \mathbf{0}), \Sigma^{*}=I\right)$
- Learner Normal-Inverse-Wishart prior $\mu_{0}=(1,1,1), \kappa_{0}=1, \nu_{0}=2+10^{-5}, \Lambda_{0}=10^{-5} I$.
"Expensive" effort $(\mathcal{D})=n$
Optimal $\mathcal{D}$ with $n=4$, unpacked into a tetrahedron


Teaching Dimension is a Special Case

- Given concept class $C=\{c\}$, define
$P\left(y=1 \mid x, \theta_{c}\right)=[c(x)=+]$ and $P(x)$ uniform. - The world has $\theta^{*}=\theta_{c^{*}}$
- The learner has $\Theta=\left\{\theta_{c} \mid c \in C\right\}, p_{0}(\theta)=\frac{1}{|C|}$
- $P\left(\theta_{c} \mid \mathcal{D}\right)=\frac{1}{\mid c \in C \text { consistent with } \mathcal{D} \mid}$ or 0 .

Teaching dimension [Goldman \& Kearns'95] $T D\left(c^{*}\right)$ is the minimum cardinality of $\mathcal{D}$ that uniquely identifies the target concept:

$$
\min _{\mathcal{D}}-\log P\left(\theta_{c^{\bullet}} \mid \mathcal{D}\right)+\gamma|\mathcal{D}|
$$

where $\gamma \leq \frac{1}{|C|}$
The solution $\mathcal{D}$ is a minimum teaching set for $c^{*}$, and $|\mathcal{D}|=T D\left(c^{*}\right)$

