# Machine Learning I 

B. Andres, J. Irmai, J. Presberger, D. Stein, S. Zhao

Machine Learning for Computer Vision
TU Dresden


Winter Term 2023/2024

## Deciding with Linear Functions

Contents. This part of the course is about a special case of supervised learning: the supervised learning of linear functions by logistic regression.

- We state the problem by defining labeled data, the family of functions and a probability distribution whose maximization motivates a regularizer and a loss function
- We show: This supervised learning problem is convex and can thus be solved by means of the steepest descent algorithm.

Deciding with Linear Functions


We consider real attributes. More specifically, we consider some finite set $V \neq \emptyset$ and labeled data $T=(S, X, x, y)$ with $X=\mathbb{R}^{V}$.

Deciding with Linear Functions


We consider real attributes. More specifically, we consider some finite set $V \neq \emptyset$ and labeled data $T=(S, X, x, y)$ with $X=\mathbb{R}^{V}$.
Hence, $x: S \rightarrow \mathbb{R}^{V}$ and $y: S \rightarrow\{0,1\}$.

Deciding with Linear Functions


We consider linear functions. More specifically, we consider $\Theta=\mathbb{R}^{V}$ and $f: \Theta \rightarrow \mathbb{R}^{X}$ such that

$$
\begin{equation*}
\forall \theta \in \Theta \forall \hat{x} \in X: \quad f_{\theta}(\hat{x})=\langle\theta, \hat{x}\rangle=\sum_{v \in V} \theta_{v} \hat{x}_{v} \tag{1}
\end{equation*}
$$

## Deciding with Linear Functions



## Random Variables

- For any sample $s \in S$, let $X_{s}$ be a random variable whose value is a vector $x_{s} \in \mathbb{R}^{V}$, the attribute vector of $s$


## Deciding with Linear Functions



## Random Variables

- For any sample $s \in S$, let $X_{s}$ be a random variable whose value is a vector $x_{s} \in \mathbb{R}^{V}$, the attribute vector of $s$
- For any sample $s \in S$, let $Y_{s}$ be a random variable whose value is a binary number $y_{s} \in\{0,1\}$, the label of $s$


## Deciding with Linear Functions



## Random Variables

- For any sample $s \in S$, let $X_{s}$ be a random variable whose value is a vector $x_{s} \in \mathbb{R}^{V}$, the attribute vector of $s$
- For any sample $s \in S$, let $Y_{s}$ be a random variable whose value is a binary number $y_{s} \in\{0,1\}$, the label of $s$
- For any $v \in V$, let $\Theta_{v}$ be a random variable whose value is a real number $\theta_{v} \in \mathbb{R}$, a parameter of the linear function we seek to learn


## Deciding with Linear Functions



Factorization

$$
\begin{equation*}
P(X, Y, \Theta)=\prod_{s \in S}\left(P\left(Y_{s} \mid X_{s}, \Theta\right) P\left(X_{s}\right)\right) \prod_{v \in V} P\left(\Theta_{v}\right) \tag{2}
\end{equation*}
$$

## Deciding with Linear Functions



Factorization

## Deciding with Linear Functions



Factorization
$P(\Theta \mid X, Y)$

## Deciding with Linear Functions



Factorization

$$
P(\Theta \mid X, Y)=\frac{P(X, Y, \Theta)}{P(X, Y)}
$$

## Deciding with Linear Functions



Factorization

$$
\begin{aligned}
P(\Theta \mid X, Y) & =\frac{P(X, Y, \Theta)}{P(X, Y)} \\
& =\frac{P(Y \mid X, \Theta) P(X) P(\Theta)}{P(X, Y)}
\end{aligned}
$$

## Deciding with Linear Functions



Factorization

$$
\begin{aligned}
P(\Theta \mid X, Y) & =\frac{P(X, Y, \Theta)}{P(X, Y)} \\
& =\frac{P(Y \mid X, \Theta) P(X) P(\Theta)}{P(X, Y)} \\
& \propto P(Y \mid X, \Theta) P(\Theta)
\end{aligned}
$$

## Deciding with Linear Functions



Factorization

$$
\begin{aligned}
P(\Theta \mid X, Y) & =\frac{P(X, Y, \Theta)}{P(X, Y)} \\
& =\frac{P(Y \mid X, \Theta) P(X) P(\Theta)}{P(X, Y)} \\
& \propto P(Y \mid X, \Theta) P(\Theta) \\
& =\prod_{s \in S} P\left(Y_{s} \mid X_{s}, \Theta\right) \prod_{v \in V} P\left(\Theta_{v}\right)
\end{aligned}
$$

## Deciding with Linear Functions

## Distributions

- Logistic distribution

$$
\begin{equation*}
\forall s \in S: \quad p_{Y_{s} \mid X_{s}, \Theta}(1)=\frac{1}{1+2^{-f_{\theta}\left(x_{s}\right)}} \tag{3}
\end{equation*}
$$



## Deciding with Linear Functions

## Distributions

- Normal distribution with $\sigma \in \mathbb{R}^{+}$:

$$
\begin{equation*}
\forall v \in V: \quad p_{\Theta v}\left(\theta_{v}\right)=\frac{1}{\sigma \sqrt{2 \pi}} e^{-\theta_{v}^{2} / 2 \sigma^{2}} \tag{3}
\end{equation*}
$$



## Deciding with Linear Functions

Lemma. Estimating maximally probable parameters $\theta$, given attributes $x$ and labels $y$, i.e.,

$$
\underset{\theta \in \mathbb{R}^{m}}{\operatorname{argmax}} \quad p_{\Theta \mid X, Y}(\theta, x, y)
$$

is equivalent ot the supervised learning problem

$$
\begin{equation*}
\min _{\theta \in \Theta} \quad \lambda R(\theta)+\frac{1}{|S|} \sum_{s \in S} L\left(f_{\theta}\left(x_{s}\right), y_{s}\right) \tag{4}
\end{equation*}
$$

with $L, R$ and $\lambda$ such that

$$
\begin{align*}
\forall r \in \mathbb{R} \forall \hat{y} \in\{0,1\}: \quad L(r, \hat{y}) & =-\hat{y} r+\log \left(1+2^{r}\right)  \tag{5}\\
\forall \theta \in \Theta: \quad R(\theta) & =\|\theta\|_{2}^{2}  \tag{6}\\
\lambda & =\frac{\log e}{2 \sigma^{2}} . \tag{7}
\end{align*}
$$

## Deciding with Linear Functions

Lemma. Estimating maximally probable parameters $\theta$, given attributes $x$ and labels $y$, i.e.,

$$
\underset{\theta \in \mathbb{R}^{m}}{\operatorname{argmax}} \quad p_{\Theta \mid X, Y}(\theta, x, y)
$$

is equivalent ot the supervised learning problem

$$
\begin{equation*}
\min _{\theta \in \Theta} \quad \lambda R(\theta)+\frac{1}{|S|} \sum_{s \in S} L\left(f_{\theta}\left(x_{s}\right), y_{s}\right) \tag{4}
\end{equation*}
$$

with $L, R$ and $\lambda$ such that

$$
\begin{array}{rlrl}
\forall r \in \mathbb{R} \forall \hat{y} \in\{0,1\}: \quad L(r, \hat{y}) & =-\hat{y} r+\log \left(1+2^{r}\right) \\
\forall \theta \in \Theta: & R(\theta) & =\|\theta\|_{2}^{2} \\
\lambda & =\frac{\log e}{2 \sigma^{2}} . \tag{7}
\end{array}
$$

It is called the $l_{2}$-regularized logistic regression problem with respect to $x, y$ and $\sigma$.

## Deciding with Linear Functions

Proof. Firstly,

$$
\begin{align*}
\underset{\theta \in \mathbb{R}^{m}}{\operatorname{argmax}} & p_{\Theta \mid X, Y}(\theta, x, y) \\
=\underset{\theta \in \mathbb{R}^{m}}{\operatorname{argmax}} & \prod_{s \in S} p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}, \theta\right) \prod_{v \in V} p_{\Theta_{v}}\left(\theta_{v}\right) \\
=\underset{\theta \in \mathbb{R}^{m}}{\operatorname{argmax}} & \sum_{s \in S} \log p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}, \theta\right)+\sum_{v \in V} \log p_{\Theta_{v}}\left(\theta_{v}\right) \tag{8}
\end{align*}
$$

## Deciding with Linear Functions

Proof. Firstly,

$$
\begin{align*}
\underset{\theta \in \mathbb{R}^{m}}{\operatorname{argmax}} & p_{\Theta \mid X, Y}(\theta, x, y) \\
=\underset{\theta \in \mathbb{R}^{m}}{\operatorname{argmax}} & \prod_{s \in S} p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}, \theta\right) \prod_{v \in V} p_{\Theta_{v}}\left(\theta_{v}\right) \\
=\underset{\theta \in \mathbb{R}^{m}}{\operatorname{argmax}} & \sum_{s \in S} \log p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}, \theta\right)+\sum_{v \in V} \log p_{\Theta_{v}}\left(\theta_{v}\right) \tag{8}
\end{align*}
$$

Secondly,

$$
\begin{align*}
& \log p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}, \theta\right) \\
= & y_{s} \log p_{Y_{s} \mid X_{s}, \Theta}\left(1, x_{s}, \theta\right)+\left(1-y_{s}\right) \log p_{Y_{s} \mid X_{s}, \Theta}\left(0, x_{s}, \theta\right) \\
= & y_{s} \log \frac{p_{Y_{s} \mid X_{s}, \Theta}\left(1, x_{s}, \theta\right)}{p_{Y_{s} \mid X_{s}, \Theta}\left(0, x_{s}, \theta\right)}+\log p_{Y_{s} \mid X_{s}, \Theta}\left(0, x_{s}, \theta\right) \tag{9}
\end{align*}
$$

## Deciding with Linear Functions

Proof. Firstly,

$$
\begin{align*}
\underset{\theta \in \mathbb{R}^{m}}{\operatorname{argmax}} & p_{\Theta \mid X, Y}(\theta, x, y) \\
=\underset{\theta \in \mathbb{R}^{m}}{\operatorname{argmax}} & \prod_{s \in S} p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}, \theta\right) \prod_{v \in V} p_{\Theta_{v}}\left(\theta_{v}\right) \\
=\underset{\theta \in \mathbb{R}^{m}}{\operatorname{argmax}} & \sum_{s \in S} \log p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}, \theta\right)+\sum_{v \in V} \log p_{\Theta_{v}}\left(\theta_{v}\right) \tag{8}
\end{align*}
$$

Secondly,

$$
\begin{align*}
& \log p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}, \theta\right) \\
= & y_{s} \log p_{Y_{s} \mid X_{s}, \Theta}\left(1, x_{s}, \theta\right)+\left(1-y_{s}\right) \log p_{Y_{s} \mid X_{s}, \Theta}\left(0, x_{s}, \theta\right) \\
= & y_{s} \log \frac{p_{Y_{s} \mid X_{s}, \Theta}\left(1, x_{s}, \theta\right)}{p_{Y_{s} \mid X_{s}, \Theta}\left(0, x_{s}, \theta\right)}+\log p_{Y_{s} \mid X_{s}, \Theta}\left(0, x_{s}, \theta\right) \tag{9}
\end{align*}
$$

Thus, with (3) and (4):

$$
\begin{equation*}
\underset{\theta \in \mathbb{R}^{m}}{\operatorname{argmin}} \sum_{s \in S}\left(-y_{s}\left\langle\theta, x_{s}\right\rangle+\log \left(1+2^{\left\langle\theta, x_{s}\right\rangle}\right)\right)+\frac{\log e}{2 \sigma^{2}}\|\theta\|_{2}^{2} \tag{10}
\end{equation*}
$$

## Deciding with Linear Functions

Lemma. The objective function

$$
\begin{equation*}
\varphi(\theta)=\lambda R(\theta)+\frac{1}{|S|} \sum_{s \in S} L\left(f_{\theta}\left(x_{s}\right), y_{s}\right) \tag{11}
\end{equation*}
$$

of the $l_{2}$-regularized logistic regression problem is convex.

## Deciding with Linear Functions

Lemma. The objective function

$$
\begin{equation*}
\varphi(\theta)=\lambda R(\theta)+\frac{1}{|S|} \sum_{s \in S} L\left(f_{\theta}\left(x_{s}\right), y_{s}\right) \tag{11}
\end{equation*}
$$

of the $l_{2}$-regularized logistic regression problem is convex.

Proof. Exercise!

## Deciding with Linear Functions

Lemma. The objective function

$$
\begin{equation*}
\varphi(\theta)=\lambda R(\theta)+\frac{1}{|S|} \sum_{s \in S} L\left(f_{\theta}\left(x_{s}\right), y_{s}\right) \tag{11}
\end{equation*}
$$

of the $l_{2}$-regularized logistic regression problem is convex.

## Proof. Exercise!

The problem can be solved by the steepest descent algorithm with a tolerance parameter $\epsilon \in \mathbb{R}_{0}^{+}$:

$$
\begin{aligned}
& \theta:=0 \\
& \text { repeat } \\
& \quad d:=\nabla \varphi(\theta) \\
& \eta:=\operatorname{argmin} \\
& \quad \theta:=\theta-\eta d \\
& \quad \text { if }\|d\|<\epsilon \\
& \quad \quad \text { return } \theta
\end{aligned}
$$

## Deciding with Linear Functions

Lemma: Estimating maximally probable labels $y$, given attributes $x^{\prime}$ and parameters $\theta$, i.e.,

$$
\begin{equation*}
\underset{y \in\{0,1\}^{S}}{\operatorname{argmax}} \quad p_{Y \mid X, \Theta}\left(y, x^{\prime}, \theta\right) \tag{12}
\end{equation*}
$$

is equivalent to the inference problem

$$
\begin{equation*}
\min _{y^{\prime} \in\{0,1\}^{S}} \sum_{s \in S} L\left(f_{\theta}\left(x_{s}\right), y_{s}^{\prime}\right) \tag{13}
\end{equation*}
$$

It has the solution

$$
\forall s \in S^{\prime}: \quad y_{s}= \begin{cases}1 & \text { if } f_{\theta}\left(x_{s}^{\prime}\right)>0  \tag{14}\\ 0 & \text { otherwise }\end{cases}
$$

## Deciding with Linear Functions

Proof. Firstly,

$$
\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} \quad p_{Y \mid X, \Theta}\left(y, x^{\prime}, \theta\right)
$$

## Deciding with Linear Functions

Proof. Firstly,

$$
\begin{aligned}
\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} & p_{Y \mid X, \Theta}\left(y, x^{\prime}, \theta\right) \\
=\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} & \prod_{s \in S^{\prime}} p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}^{\prime}, \theta\right)
\end{aligned}
$$

## Deciding with Linear Functions

Proof. Firstly,

$$
\begin{aligned}
\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} & p_{Y \mid X, \Theta}\left(y, x^{\prime}, \theta\right) \\
=\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} & \prod_{s \in S^{\prime}} p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}^{\prime}, \theta\right) \\
=\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} & \sum_{s \in S^{\prime}} \log p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}^{\prime}, \theta\right)
\end{aligned}
$$

## Deciding with Linear Functions

Proof. Firstly,

$$
\begin{aligned}
& \underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} \\
&=\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} p_{Y \mid X, \Theta}\left(y, x^{\prime}, \theta\right) \\
&=\underset{y \in S^{\prime}}{\operatorname{argmax}} p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}^{\prime}, \theta\right) \\
&=\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} \sum_{s \in S^{\prime}} \log p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}^{\prime}, \theta\right) \\
& y \in\{0,1\}^{S^{\prime}} \sum_{s \in S^{\prime}}\left(y_{s} \log \frac{p_{Y_{s} \mid X_{s}, \Theta}\left(1, x_{s}^{\prime}, \theta\right)}{p_{Y_{s} \mid X_{s}, \Theta}\left(0, x_{s}^{\prime}, \theta\right)}+\log p_{Y_{s} \mid X_{s}, \Theta}\left(0, x_{s}^{\prime}, \theta\right)\right)
\end{aligned}
$$

## Deciding with Linear Functions

## Proof. Firstly,

$$
\begin{aligned}
& \underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} \\
&=\underset{y \mid X, \Theta}{ }\left(y, x^{\prime}, \theta\right) \\
&=\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} \prod_{s \in S^{\prime}} p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}^{\prime}, \theta\right) \\
&=\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} \sum_{s \in S^{\prime}} \log p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}^{\prime}, \theta\right) \\
&=\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} \sum_{s \in S^{\prime}}\left(y_{s} \log \frac{p_{Y_{s} \mid X_{s}, \Theta}\left(1, x_{s}^{\prime}, \theta\right)}{p_{Y_{s} \mid X_{s}, \Theta}\left(0, x_{s}^{\prime}, \theta\right)}+\log p_{Y_{s} \mid X_{s}, \Theta}\left(0, x_{s}^{\prime}, \theta\right)\right) \\
&= \underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmin}} \\
& \sum_{s \in S^{\prime}}\left(-y_{s} f_{\theta}\left(x_{s}^{\prime}\right)+\log \left(1+2^{f_{\theta}\left(x_{s}^{\prime}\right)}\right)\right)
\end{aligned}
$$

## Deciding with Linear Functions

Proof. Firstly,

$$
\begin{aligned}
& \operatorname{argmax} \quad p_{Y \mid X, \Theta}\left(y, x^{\prime}, \theta\right) \\
& y \in\{0,1\}^{S^{\prime}} \\
& =\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} \prod_{s \in S^{\prime}} p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}^{\prime}, \theta\right) \\
& =\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} \sum_{s \in S^{\prime}} \log p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}^{\prime}, \theta\right) \\
& =\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} \sum_{s \in S^{\prime}}\left(y_{s} \log \frac{p_{Y_{s} \mid X_{s}, \Theta}\left(1, x_{s}^{\prime}, \theta\right)}{p_{Y_{s} \mid X_{s}, \Theta}\left(0, x_{s}^{\prime}, \theta\right)}+\log p_{Y_{s} \mid X_{s}, \Theta}\left(0, x_{s}^{\prime}, \theta\right)\right) \\
& =\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmin}} \sum_{s \in S^{\prime}}\left(-y_{s} f_{\theta}\left(x_{s}^{\prime}\right)+\log \left(1+2^{f_{\theta}\left(x_{s}^{\prime}\right)}\right)\right) \\
& =\underset{y \in\{0,1\} S^{\prime}}{\operatorname{argmin}} \sum_{s \in S^{\prime}} L\left(f_{\theta}\left(x_{s}^{\prime}\right), y_{s}\right) .
\end{aligned}
$$

## Deciding with Linear Functions

Proof. Firstly,

$$
\begin{aligned}
& \underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}}
\end{aligned} p_{Y \mid X, \Theta}\left(y, x^{\prime}, \theta\right), \begin{aligned}
\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} & \prod_{s \in S^{\prime}} p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}^{\prime}, \theta\right) \\
=\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} & \sum_{s \in S^{\prime}} \log p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}^{\prime}, \theta\right) \\
=\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} & \sum_{s \in S^{\prime}}\left(y_{s} \log \frac{p_{Y_{s} \mid X_{s}, \Theta}\left(1, x_{s}^{\prime}, \theta\right)}{p_{Y_{s} \mid X_{s}, \Theta}\left(0, x_{s}^{\prime}, \theta\right)}+\log p_{Y_{s} \mid X_{s}, \Theta}\left(0, x_{s}^{\prime}, \theta\right)\right) \\
=\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmin}} & \sum_{s \in S^{\prime}}\left(-y_{s} f_{\theta}\left(x_{s}^{\prime}\right)+\log \left(1+2^{f_{\theta}\left(x_{s}^{\prime}\right)}\right)\right) \\
=\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmin}} & \sum_{s \in S^{\prime}} L\left(f_{\theta}\left(x_{s}^{\prime}\right), y_{s}\right) .
\end{aligned}
$$

Secondly,

$$
\min _{y \in\{0,1\}^{S^{\prime}}} \sum_{s \in S^{\prime}}\left(-y_{s} f_{\theta}\left(x_{s}^{\prime}\right)+\log \left(1+2^{f_{\theta}\left(x_{s}^{\prime}\right)}\right)\right)
$$

## Deciding with Linear Functions

Proof. Firstly,

$$
\begin{aligned}
& \operatorname{argmax} \quad p_{Y \mid X, \Theta}\left(y, x^{\prime}, \theta\right) \\
& y \in\{0,1\}^{S^{\prime}} \\
& =\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} \prod_{s \in S^{\prime}} p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}^{\prime}, \theta\right) \\
& =\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} \sum_{s \in S^{\prime}} \log p_{Y_{s} \mid X_{s}, \Theta}\left(y_{s}, x_{s}^{\prime}, \theta\right) \\
& =\underset{y \in\{0,1\}^{S^{\prime}}}{\operatorname{argmax}} \sum_{s \in S^{\prime}}\left(y_{s} \log \frac{p_{Y_{s} \mid X_{s}, \Theta}\left(1, x_{s}^{\prime}, \theta\right)}{p_{Y_{s} \mid X_{s}, \Theta}\left(0, x_{s}^{\prime}, \theta\right)}+\log p_{Y_{s} \mid X_{s}, \Theta}\left(0, x_{s}^{\prime}, \theta\right)\right) \\
& =\underset{y \in\{0,1\} S^{\prime}}{\operatorname{argmin}} \sum_{s \in S^{\prime}}\left(-y_{s} f_{\theta}\left(x_{s}^{\prime}\right)+\log \left(1+2^{f_{\theta}\left(x_{s}^{\prime}\right)}\right)\right) \\
& =\underset{y \in\{0,1\} S^{\prime}}{\operatorname{argmin}} \sum_{s \in S^{\prime}} L\left(f_{\theta}\left(x_{s}^{\prime}\right), y_{s}\right) .
\end{aligned}
$$

Secondly,

$$
\min _{y \in\{0,1\}} \sum_{s \in S^{\prime}}\left(-y_{s} f_{\theta}\left(x_{s}^{\prime}\right)+\log \left(1+2^{f_{\theta}\left(x_{s}^{\prime}\right)}\right)\right)=\sum_{s \in S^{\prime}} \max _{y_{s} \in\{0,1\}} y_{s} f_{\theta}\left(x_{s}^{\prime}\right)
$$

## Deciding with Linear Functions

## Summary.

- The $l_{2}$-regularized logistic regression problem is a supervised learning problem w.r.t. the family of linear functions.
- It is motivated by a Bayesian statistical model with the logistic distribution as the likelihood as the normal distribution as the prior.
- It is a convex optimization problem that can be solved, e.g., by the steepest descent algorithm.

