## Announcements

If you have not already, please take this anonymous poll (also linked to on Slack). Thank you! https://tinyurl.com/ybhr5dfn

Start thinking about projects, dates are up

# Review: Cross-Validation 

Machine Learning - CSE546 Kevin Jamieson University of Washington October 12, 2016

## Use $k$-fold cross validation

- Randomly divide training data into $k$ equal parts
$D_{1}, \ldots, D_{k}$
- For each $i$
$\square$ Learn classifier $f_{D \mid D i}$ using data point not in $D_{i}$
$\square$ Estimate error of $f_{D I D i}$ on validation set $D_{i}$ :

$$
\operatorname{error}_{\mathcal{D}_{i}}=\frac{1}{\left|\mathcal{D}_{i}\right|} \sum_{\left(x_{j}, y_{j}\right) \in \mathcal{D}_{i}}\left(y_{j}-f_{\mathcal{D} \backslash \mathcal{D}_{i}}\left(x_{j}\right)\right)^{2}
$$

- $k$-fold cross validation error is average over data splits:

$$
\text { error }_{k-\text { fold }=\frac{1}{k} \sum_{i=1}^{k} \text { error }_{\mathcal{D}_{i}} \text { }}
$$

- $k$-fold cross validation properties:
$\square$ Much faster to compute than LOO
$\square$ More (pessimistically) biased - using much less data, only $n(k-1) / k$
- Usually, k=10


## Recap

- Given a dataset, begin by splitting into
- Model selection: Use k-fold cross-validation on TRAIN to train predictor and choose magic parameters such as $\lambda$


## VAL-3 TRAIN-3

- Model assessment: Use TEST to assess the accuracy of the model you output
- Never ever ever ever ever train or choose parameters based on the test data


## Bootstrap: basic idea

Given dataset drawn iid samples with CDF $F_{Z}$ :

$$
\mathcal{D}=\left\{z_{1}, \ldots, z_{n}\right\} \stackrel{i . i . d .}{\sim} F_{Z} \quad \widehat{\theta}=t(\mathcal{D})
$$

For $\mathrm{b}=1, \ldots, \mathrm{~B}$, samples sampled with replacement from $D$

$$
\mathcal{D}^{* b}=\left\{z_{1}^{* b}, \ldots, z_{n}^{* b}\right\} \stackrel{i . i . d .}{\sim} \widehat{F}_{Z, n} \quad \theta^{* b}=t\left(\mathcal{D}^{* b}\right)
$$



$$
\sup \left|\widehat{\boldsymbol{F}}_{n}(x)-\boldsymbol{H}(x)\right| \rightarrow 0 \quad a \mathrm{X}, \quad \rightarrow \quad \infty
$$

## Applications

Common applications of the bootstrap:

- Estimate parameters that escape simple analysis like the variance or median of an estimate
- Confidence intervals
- Estimates of error for a particular example:


Figures from Hastie et al

## Takeaways

Advantages:

- Bootstrap is very generally applicable. Build a confidence interval around anything
- Very simple to use
- Appears to give meaningful results even when the amount of data is very small
- Very strong asymptotic theory (as num. examples goes to infinity)

Disadvantages

- Very few meaningful finite-sample guarantees
- Potentially computationally intensive
- Reliability relies on test statistic and rate of convergence of empirical CDF to true CDF, which is unknown
- Poor performance on "extreme statistics" (e.g., the max)

Not perfect, but better than nothing.

## Recap

- Learning is...
$\square$ Collect some data
- E.g., housing info and sale price
$\square$ Randomly split dataset into TRAIN, VAL, and TEST
- E.g., 80\%, 10\%, and 10\%, respectively
$\square$ Choose a hypothesis class or model
- E.g., linear with non-linear transformations
$\square$ Choose a loss function
- E.g., least squares with ridge regression penalty on TRAIN
$\square$ Choose an optimization procedure
- E.g., set derivative to zero to obtain estimator, cross-validation on VAL to pick num. features and amount of regularization
$\square$ Justifying the accuracy of the estimate
- E.g., report TEST error with Bootstrap confidence interval


# Simple Variable Selection LASSO: Sparse Regression 

Machine Learning - CSE546 Kevin Jamieson University of Washington

October 11, 2016

$$
\widehat{w}_{L S}=\arg \min _{w} \sum_{i=1}^{n}\left(y_{i}-x_{i}^{T} w\right)^{2}
$$

- Vector w is sparse, if many entries are zero
- Very useful for many tasks, e.g.,

Efficiency: If size(w) = 100 Billion, each prediction is expensive:

- If part of an online system, too slow
- If w is sparse, prediction computation only depends on number of non-zeros


## Sparsity

$$
\widehat{w}_{L S}=\arg \min _{w} \sum_{i=1}^{n}\left(y_{i}-x_{i}^{T} w\right)^{2}
$$

- Vector w is sparse, if many entries are zero
- Very useful for many tasks, e.g.,
$\square$ Efficiency: If size(w) = 100 Billion, each prediction is expensive:
- If part of an online system, too slow
- If wis sparse, prediction computation only depends on number of non-zeros

Interpretability: What are the relevant dimension to make a prediction?

- E.g., what are the parts of the brain associated with particular words?


Superior temporal sulcus (posterior) ( $x-12 \mathrm{~mm}$ )

## Sparsity

$$
\widehat{w}_{L S}=\arg \min _{w} \sum_{i=1}^{n}\left(y_{i}-x_{i}^{T} w\right)^{2}
$$

- Vector w is sparse, if many entries are zero
- Very useful for many tasks, e.g.,
$\square$ Efficiency: If size(w) = 100 Billion, each prediction is expensive:
- If part of an online system, too slow
- If wis sparse, prediction computation only depends on number of non-zeros

Interpretability: What are the relevant dimension to make a prediction?

- E.g., what are the parts of the brain associated with particular words?
- How do we find "best" subset among all possible?



## Greedy model selection algorithm

- Pick a dictionary of features
$\square$ e.g., cosines of random inner products
Greedy heuristic:
$\square$ Start from empty (or simple) set of features $F_{0}=\varnothing$
$\square$ Run learning algorithm for current set of features $F_{t}$
- Obtain weights for these features
$\square$ Select next best feature $\mathbf{h}_{\mathbf{i}}(\mathbf{x})^{*}$
- e.g., $h_{j}(x)$ that results in lowest training error learner when using $F_{t}+\left\{h_{j}(x)^{*}\right\}$
$\square F_{t+1} \leftarrow F_{t}+\left\{\mathrm{h}_{\mathrm{i}}(\mathrm{x})^{*}\right\}$
$\square$ Recurse


## Greedy model selection

- Applicable in many other settings:
$\square$ Considered later in the course:
- Logistic regression: Selecting features (basis functions)
- Naïve Bayes: Selecting (independent) features $\mathrm{P}\left(\mathrm{X}_{\mathrm{i}} \mid \mathrm{Y}\right)$
- Decision trees: Selecting leaves to expand
- Only a heuristic!

Finding the best set of $k$ features is computationally intractable!
$\square$ Sometimes you can prove something strong about it...

## When do we stop???

Greedy heuristic:
$\square$ Select next best feature $\mathbf{X}_{i}^{*}$

- E.g. $\mathrm{h}_{\mathrm{j}}(\mathrm{x})$ that results in lowest training error learner when using $F_{t}+\left\{\mathrm{h}_{\mathrm{j}}(\mathrm{x})^{*}\right\}$
$\square$ Recurse
When do you stop???
- When training error is low enough?
- When test set error is low enough?
- Using cross validation?

Is there a more principled approach?

## Recall Ridge Regression

- Ridge Regression objective:

$$
\widehat{w}_{\text {ridge }}=\arg \min _{w} \sum_{i=1}^{n}\left(y_{i}-x_{i}^{T} w\right)^{2}+\lambda\|w\|_{2}^{2}
$$



## Ridge vs. Lasso Regression

- Ridge Regression objective:

$$
\widehat{w}_{\text {ridge }}=\arg \min _{w} \sum_{i=1}^{n}\left(y_{i}-x_{i}^{T} w\right)^{2}+\lambda\|w\|_{2}^{2}
$$


$+\lambda$

- Lasso Ridge objective:

$$
\begin{aligned}
& \widehat{w}_{\text {lasso }}=\arg \min _{w} \sum_{i=1}^{n}\left(y_{i}-x_{i}^{T} w\right)^{2}+\lambda\|w\|_{1} \\
& +\ldots+ \\
& +\lambda
\end{aligned}
$$



## Penalized Least Squares

$$
\begin{aligned}
& \text { Ridge }: r(w)=\|w\|_{2}^{2} \quad \text { Lasso }: r(w)=\|w\|_{1} \\
& \qquad \widehat{w}_{r}=\arg \min _{w} \sum_{i=1}^{n}\left(y_{i}-x_{i}^{T} w\right)^{2}+\lambda r(w)
\end{aligned}
$$

## Penalized Least Squares

$$
\begin{aligned}
& \text { Ridge : } r(w)=\|w\|_{2}^{2} \quad \text { Lasso : } r(w)=\|w\|_{1} \\
& \qquad \widehat{w}_{r}=\arg \min _{w} \sum_{i=1}^{n}\left(y_{i}-x_{i}^{T} w\right)^{2}+\lambda r(w)
\end{aligned}
$$

For any $\lambda \geq 0$ for which $\widehat{w}_{r}$ achieves the minimum, there exists a $\nu \geq 0$ such that

$$
\widehat{w}_{r}=\arg \min _{w} \sum_{i=1}^{n}\left(y_{i}-x_{i}^{T} w\right)^{2} \quad \text { subject to } r(\lambda) \leq \nu
$$

## Penalized Least Squares

$$
\begin{aligned}
& \text { Ridge }: r(w)=\|w\|_{2}^{2} \quad \text { Lasso }: r(w)=\|w\|_{1} \\
& \qquad \widehat{w}_{r}=\arg \min _{w} \sum_{i=1}^{n}\left(y_{i}-x_{i}^{T} w\right)^{2}+\lambda r(w)
\end{aligned}
$$

For any $\lambda \geq 0$ for which $\widehat{w}_{r}$ achieves the minimum, there exists a $\nu \geq 0$ such that

$$
\widehat{w}_{r}=\arg \min _{w} \sum_{i=1}^{n}\left(y_{i}-x_{i}^{T} w\right)^{2} \quad \text { subject to } r(\lambda) \leq \nu
$$



## Optimizing the LASSO Objective

- LASSO solution:

$$
\begin{aligned}
& \widehat{w}_{\text {lasso }}, \widehat{b}_{\text {lasso }}=\arg \min _{w, b} \sum_{i=1}^{n}\left(y_{i}-\left(x_{i}^{T} w+b\right)\right)^{2}+\lambda\|w\|_{1} \\
& \left.\widehat{b}_{\text {lasso }}=\arg \min _{w, b} \frac{1}{n} \sum_{i=1}^{n}\left(y_{i}-x_{i}^{T} \widehat{w}_{\text {lasso }}\right)\right)
\end{aligned}
$$

## Optimizing the LASSO Objective

- LASSO solution:

$$
\begin{aligned}
& \widehat{w}_{\text {lasso }}, \widehat{b}_{\text {lasso }}=\arg \min _{w, b} \sum_{i=1}^{n}\left(y_{i}-\left(x_{i}^{T} w+b\right)\right)^{2}+\lambda\|w\|_{1} \\
& \left.\widehat{b}_{\text {lasso }}=\arg \min _{w, b} \frac{1}{n} \sum_{i=1}^{n}\left(y_{i}-x_{i}^{T} \widehat{w}_{\text {lasso }}\right)\right)
\end{aligned}
$$

So as usual, preprocess to make sure that $\frac{1}{n} \sum_{i=1}^{n} y_{i}=0, \frac{1}{n} \sum_{i=1}^{n} x_{i}=\mathbf{0}$
so we don't have to worry about an offset.

## Optimizing the LASSO Objective

- LASSO solution:

$$
\begin{aligned}
& \widehat{w}_{\text {lasso }}, \widehat{b}_{\text {lasso }}=\arg \min _{w, b} \sum_{i=1}^{n}\left(y_{i}-\left(x_{i}^{T} w+b\right)\right)^{2}+\lambda\|w\|_{1} \\
& \left.\widehat{b}_{\text {lasso }}=\arg \min _{w, b} \frac{1}{n} \sum_{i=1}^{n}\left(y_{i}-x_{i}^{T} \widehat{w}_{\text {lasso }}\right)\right)
\end{aligned}
$$

So as usual, preprocess to make sure that $\frac{1}{n} \sum_{i=1}^{n} y_{i}=0, \frac{1}{n} \sum_{i=1}^{n} x_{i}=\mathbf{0}$
so we don't have to worry about an offset.

$$
\widehat{w}_{l a s s o}=\arg \min _{w} \sum_{i=1}^{n}\left(y_{i}-x_{i}^{T} w\right)^{2}+\lambda\|w\|_{1}
$$

How do we solve this?

## Coordinate Descent

- Given a function, we want to find minimum
- Often, it is easy to find minimum along a single coordinate:
- How do we pick next coordinate?
- Super useful approach for *many* problems
$\square$ Converges to optimum in some cases, such as LASSO


## Optimizing LASSO Objective One Coordinate at a Time

$$
\begin{array}{r}
\sum_{i=1}^{n}\left(y_{i}-x_{i}^{T} w\right)^{2}+\lambda\|w\|_{1}=\sum_{i=1}^{n}\left(y_{i}-\sum_{k=1}^{d} x_{i, k} w_{k}\right)^{2}+\lambda \sum_{k=1}^{d}\left|w_{k}\right| \\
=\sum_{i=1}^{n}\left(\left(y_{i}-\sum_{k \neq j} x_{i, k} w_{k}\right)-x_{i, j} w_{j}\right)^{2}+\lambda \sum_{k \neq j}\left|w_{k}\right|+\lambda\left|w_{j}\right|
\end{array}
$$

Equivalently:

$$
\widehat{w}_{j}=\arg \min _{w_{j}} \sum_{i=1}^{n}\left(r_{i}^{(j)}-x_{i, j} w_{j}\right)^{2}+\lambda\left|w_{j}\right|
$$

## Convex Functions

- Equivalent definitions of convexity:

$f$ convex:

$$
\begin{array}{lll}
f(\lambda x+(1-\lambda) y) \leq \lambda f(x)+(1-\lambda) f(y) & \forall x, y, \lambda \in[0,1] \\
f(y) \geq f(x)+\nabla f(x)^{T}(y-x) & \forall x, y &
\end{array}
$$

- Gradients lower bound convex functions and are unique at $\mathbf{x}$ iff function differentiable at $\mathbf{x}$
- Subgradients generalize gradients to non-differentiable points:
$\square$ Any supporting hyperplane at $\mathbf{x}$ that lower bounds entire function

$$
g \text { is a subgradient at } x \text { if } f(y) \geq f(x)+g^{T}(y-x)
$$

## Taking the Subgradient $\left.\hat{w}_{j}=\operatorname{argminin}_{\min _{i}}^{n} \sum_{i=1}^{n}\left(\sigma_{i}\right)-x_{i,} w_{j}\right)^{2}+\lambda w_{j}$

- Convex function is minimized at $w$ if 0 is a sub-gradient at $w$.

$$
\partial_{w_{j}}\left|w_{j}\right|=
$$

$$
\partial_{w_{j}} \sum_{i=1}^{n}\left(r_{i}^{(j)}-x_{i, j} w_{j}\right)^{2}=
$$

## Setting Subgradient to 0

$$
\begin{gathered}
\partial_{w_{j}}\left(\sum_{i=1}^{n}\left(r_{i}^{(j)}-x_{i, j} w_{j}\right)^{2}+\lambda\left|w_{j}\right|\right)= \begin{cases}a_{j} w_{j}-c_{j}-\lambda & \text { if } w_{j}<0 \\
{\left[-c_{j}-\lambda,-c_{j}+\lambda\right]} & \text { if } w_{j}=0 \\
a_{j} w_{j}-c_{j}+\lambda & \text { if } w_{j}>0\end{cases} \\
a_{j}=\left(\sum_{i=1}^{n} x_{i, j}^{2}\right) \quad c_{j}=2\left(\sum_{i=1}^{n} r_{i}^{(j)} x_{i, j}\right) \\
\widehat{w}_{j}=\arg \min _{w_{j}} \sum_{i=1}^{n}\left(r_{i}^{(j)}-x_{i, j} w_{j}\right)^{2}+\lambda\left|w_{j}\right| \\
\widehat{w}_{j}= \begin{cases}\left(c_{j}+\lambda\right) / a_{j} & \text { if } c_{j}<-\lambda \\
0 & \text { if }\left|c_{j}\right| \leq \lambda \\
\left(c_{j}-\lambda\right) / a_{j} & \text { if } c_{j}>\lambda\end{cases}
\end{gathered}
$$

## Soft Thresholding

$$
\begin{aligned}
\widehat{w}_{j}= \begin{cases}\left(c_{j}+\lambda\right) / a_{j} & \text { if } c_{j}<-\lambda \\
0 & \text { if }\left|c_{j}\right| \leq \lambda \\
\left(c_{j}-\lambda\right) / a_{j} & \text { if } c_{j}>\lambda\end{cases} \\
a_{j}=\sum_{i=1}^{n} x_{i, j}^{2}
\end{aligned}
$$



From
Kevin Murphy textbook

## Coordinate Descent for LASSO (aka Shooting Algorithm)

- Repeat until convergence
$\square$ Pick a coordinate / at (random or sequentially)
- Set:

$$
\widehat{w}_{j}= \begin{cases}\left(c_{j}+\lambda\right) / a_{j} & \text { if } c_{j}<-\lambda \\ 0 & \text { if }\left|c_{j}\right| \leq \lambda \\ \left(c_{j}-\lambda\right) / a_{j} & \text { if } c_{j}>\lambda\end{cases}
$$



$$
c_{j}=2 \sum_{i=1}^{n}\left(y_{i}-\sum_{k \neq j} x_{i, k} w_{k}\right) x_{i, j}
$$

$\square$ For convergence rates, see Shalev-Shwartz and Tewari 2009

- Other common technique = LARS
$\square$ Least angle regression and shrinkage, Efron et al. 2004


## Recall: Ridge Coefficient Path



From
Kevin Murphy textbook

- Typical approach: select $\lambda$ using cross validation


## Now: LASSO Coefficient Path



From
Kevin Murphy textbook

## What you need to know

- Variable Selection: find a sparse solution to learning problem
- $L_{1}$ regularization is one way to do variable selection
$\square$ Applies beyond regression
$\square$ Hundreds of other approaches out there
- LASSO objective non-differentiable, but convex $\rightarrow$ Use subgradient
- No closed-form solution for minimization $\rightarrow$ Use coordinate descent
- Shooting algorithm is simple approach for solving LASSO


# Classification Logistic Regression 

Machine Learning - CSE546 Kevin Jamieson University of Washington

October 12, 2016

## THUS FAR, REGRESSION: PREDICT A CONTINUOUS VALUE GIVEN SOME INPUTS

## Weather prediction revisted

## .



## Reading Your Brain, Simple Example

## Pairwise classification accuracy: 85\%

[Mitchell et al.]

Person


## Classification

- Learn: f:X $->Y$
$\square \mathbf{X}$ - features
$\square$ Y - target classes
- Conditional probability: $\mathrm{P}(\mathrm{Y} \mid \mathbf{X})$
- Suppose you know $\mathrm{P}(\mathrm{Y} \mid \mathrm{X})$ exactly, how should you classify?

Bayes optimal classifier:

- How do we estimate $\mathrm{P}(\mathrm{Y} \mid \mathrm{X})$ ?


## Link Functions

- Estimating $\mathrm{P}(\mathrm{Y} \mid \mathbf{X})$ : Why not use standard linear regression?
- Combining regression and probability?
$\square$ Need a mapping from real values to $[0,1]$
A link function!

Logistic Regression
Logistic

Learn $\mathrm{P}(\mathrm{Y} \mid \mathbf{X})$ directly

- Assume a particular functional form for link function
Sigmoid applied to a linear function of the input features:

$$
P(Y=0 \mid X, W)=\frac{1}{1+\exp \left(w_{0}+\sum_{i} w_{i} X_{i}\right)}
$$



## Understanding the sigmoid

$$
g\left(w_{0}+\sum_{i} w_{i} x_{i}\right)=\frac{1}{1+e^{w_{0}+\sum_{i} w_{i} x_{i}}}
$$

$w_{0}=-2, w_{1}=-1$
$w_{0}=0, w_{1}=-1$
$w_{0}=0, w_{1}=-0.5$



## Very convenient!

$$
P\left(Y=0\left\llcorner\mid X=<X_{1}, \ldots X_{n}>\right)=\frac{1}{1+\exp \left(w_{0}+\sum_{i} w_{i} X_{i}\right)}\right.
$$

implies

$$
\left.P(Y=1) \mid X=<X_{1}, \ldots X_{n}>\right)=\frac{\exp \left(w_{0}+\sum_{i} w_{i} X_{i}\right)}{1+\exp \left(w_{0}+\sum_{i} w_{i} X_{i}\right)}
$$

## Very convenient!

$$
P\left(Y=0 ᄂ \mid X=<X_{1}, \ldots X_{n}>\right)=\frac{1}{1+\exp \left(w_{0}+\sum_{i} w_{i} X_{i}\right)}
$$

implies

$$
\left.P(Y=1) \mid X=<X_{1}, \ldots X_{n}>\right)=\frac{\exp \left(w_{0}+\sum_{i} w_{i} X_{i}\right)}{1+\exp \left(w_{0}+\sum_{i} w_{i} X_{i}\right)}
$$

implies

$$
\frac{P(Y=1 \| X)}{P(Y=0 . \mid X)}=\exp \left(w_{0}+\sum_{i} w_{i} X_{i}\right)
$$

linear classification rule!

$$
\ln \frac{\left.P\left(Y={ }^{1}\right) \mid X\right)}{P(Y=0 \backslash \mid X)}=w_{0}+\sum_{i} w_{i} X_{i}
$$

## Logistic Regression a Linear classifier

$$
\frac{1}{1+\exp (-z)}
$$

$$
\begin{gathered}
g\left(w_{0}+\sum_{i} w_{i} x_{i}\right)=\frac{1}{1+e^{w_{0}+\sum_{i} w_{i} x_{i}}} \\
\ln \frac{P(Y=0 \mid X)}{P(Y=1 \mid X)}=w_{0}+\sum_{i} w_{i} X_{i}
\end{gathered}
$$

## Loss function: Conditional Likelihood

- Have a bunch of iid data of the form: $\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{n} \quad x_{i} \in \mathbb{R}^{d}, \quad y_{i} \in\{-1,1\}$

$$
\begin{aligned}
P(Y=-1 \mid x, w) & =\frac{1}{1+\exp \left(w^{T} x\right)} \\
P(Y=1 \mid x, w) & =\frac{\exp \left(w^{T} x\right)}{1+\exp \left(w^{T} x\right)}
\end{aligned}
$$

- This is equivalent to:

$$
P(Y=y \mid x, w)=\frac{1}{1+\exp \left(-y w^{T} x\right)}
$$

- So we can compute the maximum likelihood estimator:

$$
\widehat{w}_{M L E}=\arg \max _{w} \prod_{i=1}^{n} P\left(y_{i} \mid x_{i}, w\right)
$$

## Loss function: Conditional Likelihood

- Have a bunch of iid data of the form: $\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{n} \quad x_{i} \in \mathbb{R}^{d}, \quad y_{i} \in\{-1,1\}$

$$
\widehat{w}_{M L E}=\arg \max _{w} \prod_{i=1}^{n} P\left(y_{i} \mid x_{i}, w\right) \quad P(Y=y \mid x, w)=\frac{1}{1+\exp \left(-y w^{T} x\right)}
$$

## Loss function: Conditional Likelihood

- Have a bunch of iid data of the form: $\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{n} \quad x_{i} \in \mathbb{R}^{d}, \quad y_{i} \in\{-1,1\}$

$$
\begin{aligned}
\widehat{w}_{M L E} & =\arg \max _{w} \prod_{i=1}^{n} P\left(y_{i} \mid x_{i}, w\right) \quad P(Y=y \mid x, w)=\frac{1}{1+\exp \left(-y w^{T} x\right)} \\
& =\arg \min _{w} \sum_{i=1}^{n} \log \left(1+\exp \left(-y_{i} x_{i}^{T} w\right)\right)
\end{aligned}
$$

## Loss function: Conditional Likelihood

- Have a bunch of iid data of the form: $\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{n} \quad x_{i} \in \mathbb{R}^{d}, \quad y_{i} \in\{-1,1\}$

$$
\begin{aligned}
\widehat{w}_{M L E} & =\arg \max _{w} \prod_{i=1}^{n} P\left(y_{i} \mid x_{i}, w\right) \quad P(Y=y \mid x, w)=\frac{1}{1+\exp \left(-y w^{T} x\right)} \\
& =\arg \min _{w} \sum_{i=1}^{n} \log \left(1+\exp \left(-y_{i} x_{i}^{T} w\right)\right)
\end{aligned}
$$

Logistic Loss: $\ell_{i}(w)=\log \left(1+\exp \left(-y_{i} x_{i}^{T} w\right)\right)$
Squared error Loss: $\ell_{i}(w)=\left(y_{i}-x_{i}^{T} w\right)^{2} \quad$ (MLE for Gaussian noise)

## Loss function: Conditional Likelihood

- Have a bunch of iid data of the form: $\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{n} \quad x_{i} \in \mathbb{R}^{d}, \quad y_{i} \in\{-1,1\}$

$$
\begin{aligned}
\widehat{w}_{M L E} & =\arg \max _{w} \prod_{i=1}^{n} P\left(y_{i} \mid x_{i}, w\right) \quad P(Y=y \mid x, w)=\frac{1}{1+\exp \left(-y w^{T} x\right)} \\
& =\arg \min _{w} \sum_{i=1}^{n} \log \left(1+\exp \left(-y_{i} x_{i}^{T} w\right)\right)=J(w)
\end{aligned}
$$

What does $J(w)$ look like? Is it convex?

## Loss function: Conditional Likelihood

- Have a bunch of iid data of the form: $\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{n} \quad x_{i} \in \mathbb{R}^{d}, \quad y_{i} \in\{-1,1\}$

$$
\begin{aligned}
\widehat{w}_{M L E} & =\arg \max _{w} \prod_{i=1}^{n} P\left(y_{i} \mid x_{i}, w\right) \quad P(Y=y \mid x, w)=\frac{1}{1+\exp \left(-y w^{T} x\right)} \\
& =\arg \min _{w} \sum_{i=1}^{n} \log \left(1+\exp \left(-y_{i} x_{i}^{T} w\right)\right)=J(w)
\end{aligned}
$$

Good news: $J(\mathbf{w})$ is convex function of $\mathbf{w}$, no local optima problems Bad news: no closed-form solution to maximize $J(\mathbf{w})$

Good news: convex functions easy to optimize (next time)

## Linear Separability

$$
\arg \min _{w} \sum_{i=1}^{n} \log \left(1+\exp \left(-y_{i} x_{i}^{T} w\right)\right) \quad \text { When is this loss small? }
$$

# Large parameters $\rightarrow$ Overfitting 



- If data is linearly separable, weights go to infinity
$\square$ In general, leads to overfitting:
- Penalizing high weights can prevent overfitting...


## Regularized Conditional Log Likelihood

- Add regularization penalty, e.g., $\mathrm{L}_{2}$ :

$$
\arg \min _{w} \sum_{i=1}^{n} \log \left(1+\exp \left(-y_{i} x_{i}^{T} w\right)\right)+\lambda\|w\|_{2}^{2}
$$

- Practical note about $\mathrm{w}_{0}$ :

