# 10-701 <br> Machine Learning 

Regression

## Where we are



Today

## Choosing a restaurant

- In everyday life we need to make decisions by taking into account lots of factors
- The question is what weight we put on each of these factors (how important are they with respect to the others).
- Assume we would like to build a recommender system for ranking potential restaurants based on an individuals' preferences

| Reviews <br> (out of 5 <br> stars) | $\$$ | Distance | Cuisine <br> (out of 10) |
| :--- | :--- | :--- | :--- |
| 4 | 30 | 21 | 7 |
| 2 | 15 | 12 | 8 |
| 5 | 27 | 53 | 9 |
| 3 | 20 | 5 | 6 |

- If we have many observations we may be able to recover the weights



## Linear regression

- Given an input $x$ we would like to compute an output y
- For example:
- Predict height from age
- Predict Google's price from

Yahoo's price

- Predict distance from wall using sensor readings


Note that now $Y$ can be continuous

## Linear regression

- Given an input $x$ we would like to compute an output y
- In linear regression we assume that $y$ and $x$ are related with the following equation:

What we are
Observed values


where w is a parameter and $\varepsilon$ represents measurement or other noise

## Linear regression

- Our goal is to estimate $w$ from a training data of $\left\langle x_{i}, y_{i}\right\rangle$ pairs
- One way to find such relationship is to minimize the a least squares error:

$$
\arg \min _{w} \sum_{i}\left(y_{i}-w x_{i}\right)^{2}
$$



- Several other approaches can be used as well
- So why least squares?
- minimizes squared distance between measurements and predicted line
- has a nice probabilistic interpretation
- easy to compute

If the noise is Gaussian with mean 0 then least squares is also the maximum likelihood estimate of w

## Solving linear regression using least squares minimization

- You should be familiar with this by now ...
- We just take the derivative w.r.t. to w and set to 0 :

$$
\begin{gathered}
\frac{\partial}{\partial w} \sum_{i}\left(y_{i}-w x_{i}\right)^{2}=2 \sum_{i}-x_{i}\left(y_{i}-w x_{i}\right) \Rightarrow \\
2 \sum_{i} x_{i}\left(y_{i}-w x_{i}\right)=0 \Rightarrow \\
\sum_{i} x_{i} y_{i}=\sum_{i} w x_{i}^{2} \Rightarrow \\
\sum_{i} x_{i} y_{i} \\
\sum_{i} x_{i}^{2}
\end{gathered}
$$

## Regression example

- Generated: w=2
- Recovered: w=2.03
- Noise: std=1



## Regression example

- Generated: w=2
- Recovered: w=2.05
- Noise: std=2



## Regression example

- Generated: w=2
- Recovered: w=2.08
- Noise: std=4



## Bias term

- So far we assumed that the line passes through the origin
- What if the line does not?
- No problem, simply change the model to

$$
y=w_{0}+w_{1} x+\varepsilon
$$

- Can use least squares to
 determine $\mathrm{w}_{0}, \mathrm{w}_{1}$

$$
w_{0}=\frac{\sum_{i} y_{i}-w_{1} x_{i}}{n}
$$

$$
w_{1}=\frac{\sum_{i} x_{i}\left(y_{i}-w_{0}\right)}{\sum_{i} x_{i}^{2}}
$$

## Bias term

- So far we assumed that the line passes through the origin
- What if the line does not?
- No problem, simply change the model to

$$
\begin{aligned}
& y=w \text { Just a second, we will soon } \\
& \text { give a simpler solution }
\end{aligned}
$$

- Can use least squares to determine $\mathrm{w}_{0}, \mathrm{w}_{1}$

$$
w_{0}=\frac{\sum_{i} y_{i}-w_{1} x_{i}}{n}
$$

$$
w_{1}=\frac{\sum_{i} x_{i}\left(y_{i}-w_{0}\right)}{\sum_{i} x_{i}^{2}}
$$

## Multivariate regression

- What if we have several inputs?
- Stock prices for Yahoo, Microsoft and Ebay for the Google prediction task
- This becomes a multivariate linear regression problem
- Again, its easy to model:

Google's stock price
Microsoft's stock price
Yahoo's stock price

## Multivariate regression

- What if we have several inputs?
- Stock prices for Yahoo. Microsoft and Ebay for
the Goc Not all functions can be
- This be approximated using the input
values directly
n problem
- Again, its easy to model:

$$
y=w_{0}+w_{1} x_{1}+\ldots+w_{k} x_{k}+\varepsilon
$$

$$
y=10+3 x_{1}^{2}-2 x_{2}^{2}+\varepsilon
$$

> In some cases we would like to use polynomial or other terms based on the input data, are these still linear regression problems?

Yes. As long as the coefficients are linear the equation is still a linear regression problem!

## Non-Linear basis function

- So far we only used the observed values
- However, linear regression can be applied in the same way to functions of these values
- As long as these functions can be directly computed from the observed values the parameters are still linear in the data and the problem remains a linear regression problem

$$
y=w_{0}+w_{1} x_{1}^{2}+\ldots+w_{k} x_{k}^{2}+\varepsilon
$$

## Non-Linear basis function

- What type of functions can we use?
- A few common examples:
- Polynomial: $\phi_{\mathrm{j}}(\mathrm{x})=\mathrm{x}^{\mathrm{j}}$ for $\mathrm{j}=0 \ldots \mathrm{n}$
- Gaussian: $\quad \phi_{j}(x)=\frac{\left(x-\mu_{j}\right)}{2 \sigma_{j}^{2}}$
- Sigmoid: $\quad \phi_{j}(x)=\frac{1}{1+\exp \left(-s_{j} x\right)}$

Any function of the input values can be used. The solution for the parameters of the regression remains the same.

## General linear regression problem

- Using our new notations for the basis function linear regression can be written as

$$
y=\sum_{j=0}^{n} w_{j} \phi_{j}(x)
$$

- Where $\phi_{j}(x)$ can be either $x_{j}$ for multivariate regression or one of the non linear basis we defined
- Once again we can use 'least squares' to find the optimal solution.


## LMS for the general linear regression problem

Our goal is to minimize the following loss function:

$$
y=\sum_{j=0}^{n} w_{j} \phi_{j}(x)
$$

$J(\mathrm{w})=\sum_{i}\left(y^{i}-\sum_{j} w_{j} \phi_{j}\left(x^{i}\right)\right)^{2}$
$w$ - vector of dimension $\mathrm{k}+1$
$\phi\left(x^{i}\right)$ - vector of dimension $\mathrm{k}+1$
Moving to vector notations we get. $\longrightarrow y^{i}$-a scaler

$$
J(\mathrm{w})=\sum_{i}\left(y^{i}-\mathrm{w}^{\mathrm{T}} \phi\left(\widehat{\left.\left.x^{i}\right)\right)^{2}}\right.\right.
$$

We take the derivative w.r.t $\mathbf{w}$

$$
\frac{\partial}{\partial w} \sum_{i}\left(y^{i}-\mathrm{w}^{\mathrm{T}} \phi\left(x^{i}\right)\right)^{2}=2 \sum_{i}\left(y^{i}-\mathrm{w}^{\mathrm{T}} \phi\left(x^{i}\right)\right) \phi\left(x^{i}\right)^{\mathrm{T}}
$$

Equating to 0 we get $2 \sum_{i}\left(y^{i}-\mathrm{w}^{\mathrm{T}} \phi\left(x^{i}\right)\right) \phi\left(x^{i}\right)^{\mathrm{T}}=0 \Rightarrow$

$$
\sum_{i} y^{i} \phi\left(x^{i}\right)^{\mathrm{T}}=\mathrm{w}^{\mathrm{T}}\left[\sum_{i} \phi\left(x^{i}\right) \phi\left(x^{i}\right)^{\mathrm{T}}\right]
$$

## LMS for general linear regression problem

We take the derivative w.r.t w

$$
J(\mathrm{w})=\sum_{i}\left(y^{i}-\mathrm{w}^{\mathrm{T}} \phi\left(x^{i}\right)\right)^{2}
$$

$\frac{\partial}{\partial w} \sum_{i}\left(y^{i}-\mathrm{w}^{\mathrm{T}} \phi\left(x^{i}\right)\right)^{2}=2 \sum_{i}\left(y^{i}-\mathrm{w}^{\mathrm{T}} \phi\left(x^{i}\right)\right) \phi\left(x^{i}\right)^{\mathrm{T}}$
Equating to 0 we get

$$
\begin{aligned}
& 2 \sum_{i}\left(y^{i}-\mathrm{w}^{\mathrm{T}} \phi\left(x^{i}\right)\right) \phi\left(x^{i}\right)^{\mathrm{T}}=0 \Rightarrow \\
& \sum_{i} y^{i} \phi\left(x^{i}\right)^{\mathrm{T}}=\mathrm{w}^{\mathrm{T}}\left[\sum_{i} \phi\left(x^{i}\right) \phi\left(x^{i}\right)^{\mathrm{T}}\right]
\end{aligned}
$$

Define:

$$
\Phi=\left(\begin{array}{cccc}
\phi_{0}\left(x^{1}\right) & \phi_{1}\left(x^{1}\right) & \cdots & \phi_{m}\left(x^{1}\right) \\
\phi_{0}\left(x^{2}\right) & \phi_{1}\left(x^{2}\right) & \cdots & \phi_{m}\left(x^{2}\right) \\
\vdots & \vdots & \cdots & \vdots \\
\phi_{0}\left(x^{n}\right) & \phi_{1}\left(x^{n}\right) & \cdots & \phi_{m}\left(x^{n}\right)
\end{array}\right)
$$

Then deriving w we get:

$$
\mathrm{w}=\left(\Phi^{T} \Phi\right)^{-1} \Phi^{T} \mathrm{y}
$$

## LMS for general linear regression problem

$$
J(\mathrm{w})=\sum_{i}\left(y^{i}-\mathrm{w}^{\mathrm{T}} \phi\left(x^{i}\right)\right)^{2}
$$



This solution is also known as 'psuedo inverse'

## Example: Polynomial regression



## A probabilistic interpretation

Our least squares minimization solution can also be motivated by a probabilistic in interpretation of the regression problem: $y=\mathrm{w}^{\mathrm{T}} \phi(x)+\varepsilon$

The MLE for w in this model is the same as the solution we derived for least squares criteria:

$$
\mathrm{w}=\left(\Phi^{T} \Phi\right)^{-1} \Phi^{T} \mathrm{y}
$$

## Other types of linear regression

- Linear regression is a useful model for many problems
- However, the parameters we learn for this model are global; they are the same regardless of the value of the input $x$
- Extension to linear regression adjust their parameters based on the region of the input we are dealing with


## Splines

- Instead of fitting one function for the entire region, fit a set of piecewise (usually cubic) polynomials satisfying continuity and smoothness constraints.
- Results in smooth and flexible functions without too many parameters
- Need to define the regions in advance (usually uniform)

\[

\]

## Splines

- The polynomials are not independent
- For cubic splines we require that they agree in the border point on the value, the values of the first derivative and the value of the second derivative
- How many free parameters do we actually have?

\[

\]

## Splines

- Splines sometimes contain additional requirements for the first and last polynomial (for example, having them start at 0)
- Once Splines are fitted to the data they can be used to predict new values in the same way as regular linear regression, though they are limited to the support regions for which they have been defined
- Note the range of functions that can be displayed with relatively small number of

8.099551e-09 YER124C

$3.583023 \mathrm{e}-08$ YLR194C




1.690041e-10 EGT2

$2.716782 \mathrm{e}-08$ YHR143W
9.716482e-08 YNLO58C
 polynomials (in the example I am using 5)


## Locally weighted models

- Splines rely on a fixed region for each polynomial and the weight of all points within the region is the same.
- An alternative option is to set the region based on the density of the input data and have points closer to the point we are trying to estimate have a higher weight

R-I plot raw data


## Weighted regression

- For a point x we use weight function $\Omega_{\mathrm{x}}$ centered at x to assign weight to points in x's vicinity
- Next we solve the following weighted regression problem

$$
\min _{\mathrm{w}} \sum_{i} \Omega_{x}\left(x^{i}\right)\left(y^{i}-\mathrm{w}^{\mathrm{T}} \phi\left(x^{i}\right)\right)^{2}
$$

- The solution is the same as our general solution (the weight is given for every input)



## Determining the weights

- There are a number of ways to determine the weights
- One options is to use a Gaussian centered at $x$, such that

$$
\Omega_{x}\left(x^{i}\right)=\frac{1}{\sqrt{2 \pi} \sigma} e^{-\frac{\left(x-x^{i}\right)^{2}}{2 \sigma^{2}}}
$$

$\sigma^{2}$ is a parameter that should be selected by the user

More on these weights when we discuss kernels

## Important points

- Linear regression
- basic model
- as a function of the input
- Solving linear regression
- Error in linear regression
- Advanced regression models

