

DATA SCIENCE WITH MACHINE LEARNING: REGRESSION

This lecture is
based on course by E. Fox and C. Guestrin, Univ of Washington

22/12 2020

WFAiS UJ, Informatyka Stosowana
I stopień studiów

What is Data Science?

2

Is mainly about extracting knowledge from data (terms “data mining” or “Knowledge Discovery in Databases” are highly related). It can be about analyzing trends, building predictive models, ... etc.

Is an agglomerate of **data collection, data modeling and analysis**, a decision making, and everything you need to know to accomplish your goals. Eventually, it boils down to the following fields/skills:

- Computer science:

Algorithms, programming (patterns, languages etc.), understanding hardware & operating systems, high-performance computing'

- Mathematical aspects:

Linear algebra, differential equations for optimization problems, statistics

- Few others:

Machine learning, domain knowledge, and data visualization & communication skills

Data Science and Machine Learning?

3

Machine learning algorithms are algorithms that learn (often predictive) models from data. I.e., instead of formulating "rules" manually, a machine learning algorithm will learn the model for you.

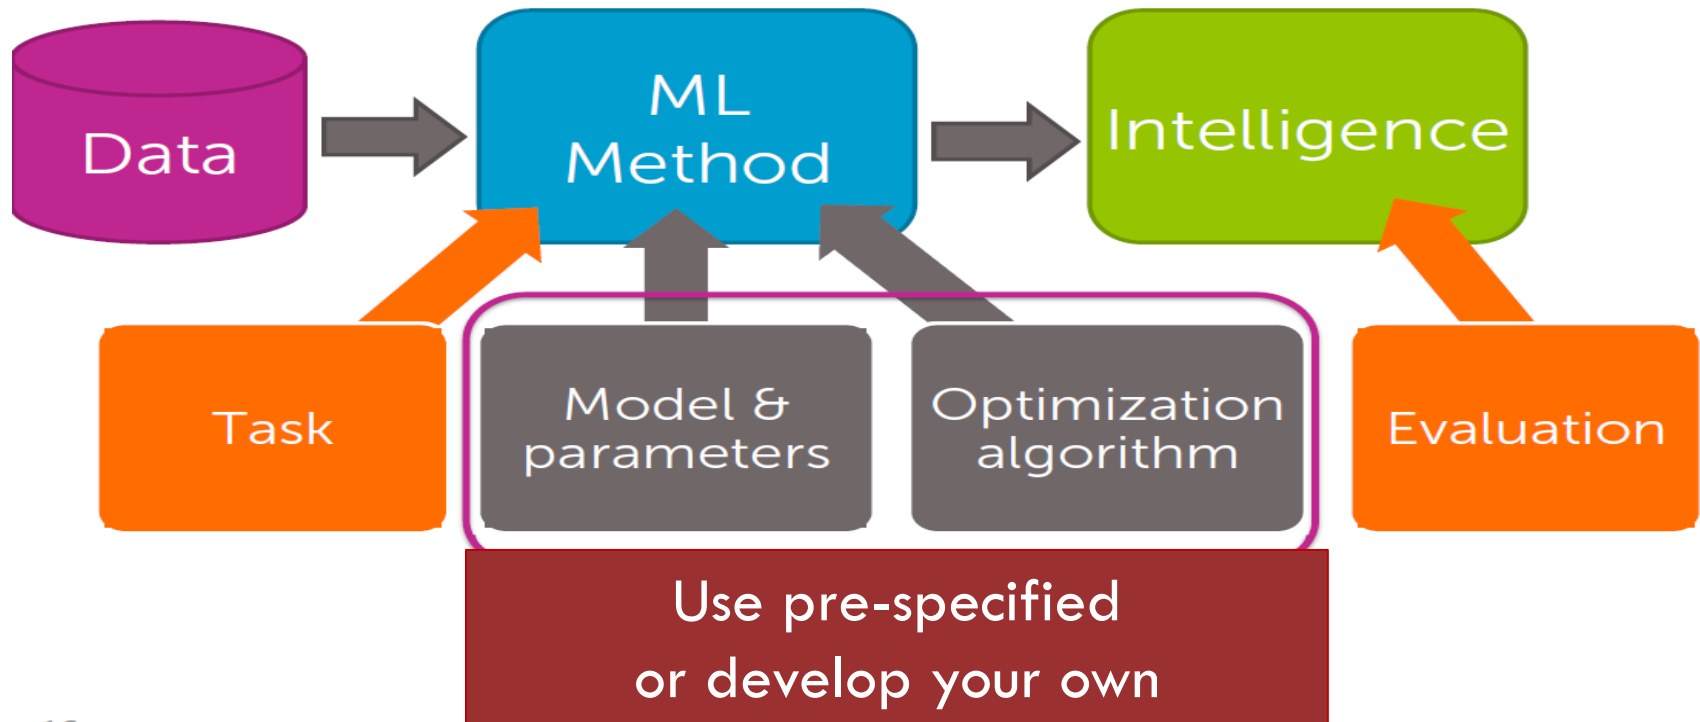
Machine learning - at its core - is about the use and development of these learning algorithms. **Data science** is more about the extraction of knowledge from data to answer particular question or solve particular problems.

Machine learning is often a big part of a "data science" project, e.g., it is often heavily used for exploratory analysis and discovery (clustering algorithms) and building predictive models (supervised learning algorithms). However, in **data science**, you often also worry about the collection, wrangling, and cleaning of your data (i.e., data engineering), and eventually, you want to draw conclusions from your data that helps you solve a particular problem.

Deploying intelligence module

4

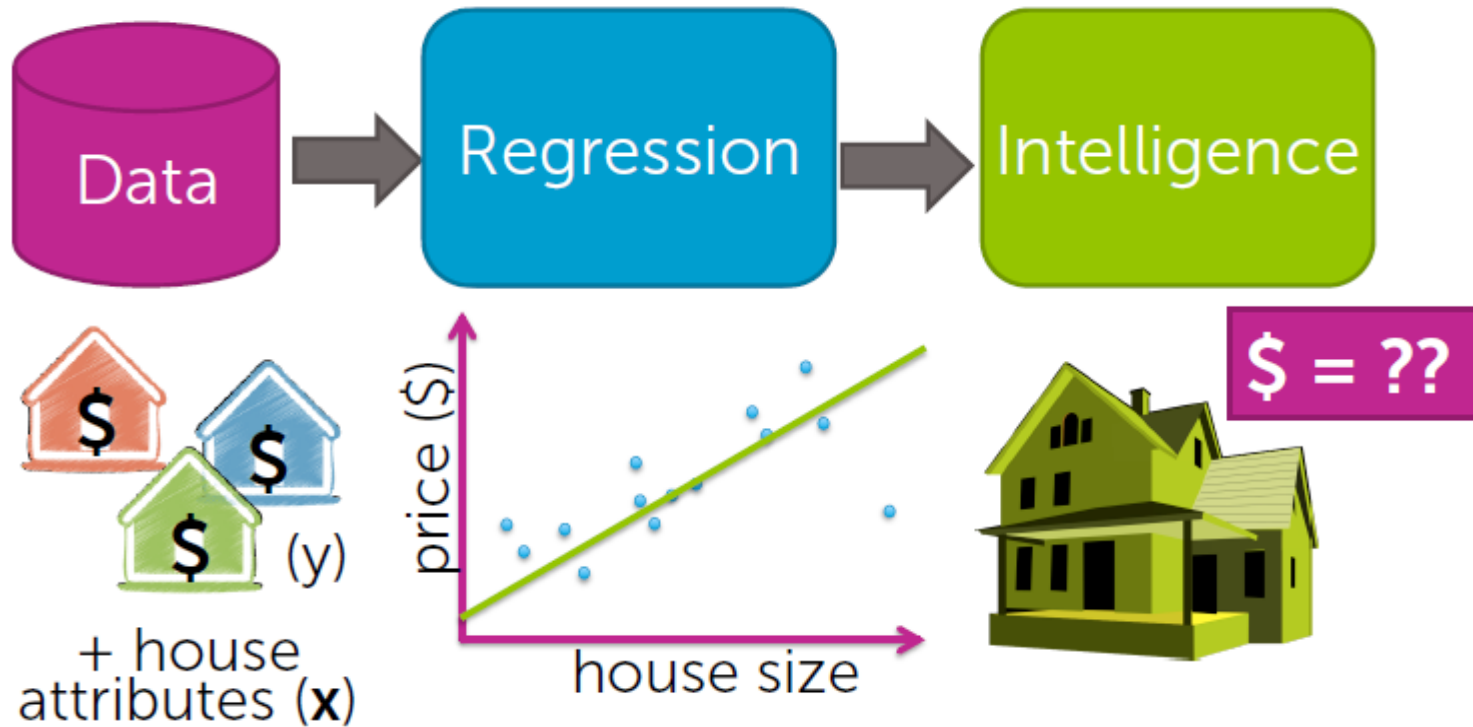
Case studied are about building, evaluating, deploying intelligence in data analysis.



Case study

5

Predicting house prices



Prediction: Predicting house prices

6

Models

- Linear regression
- Regularization: Ridge (L2), Lasso (L1)

Algorithms


- Gradient descent
- Coordinate descent

Concepts

- Loss functions, bias-variance tradeoff, cross-validation, sparsity, overfitting, model selection

Data

7




Input vs output

- y is quantity of interest
- assume y can be predicted from x

Model: assume functional relationship

8



„Essentially, all models are wrong but some are useful.“
George Box, 1987.

Regression model:

$$y_i = f(x_i) + \epsilon_i$$

$$E[\epsilon_i] = 0 \quad \leftarrow \text{equally likely that error is + or -}$$





↑ expected value

⇓
 y_i is equally likely to be above or below $f(x_i)$

Task 1:

9


Which model to fit?



Task 2:

10

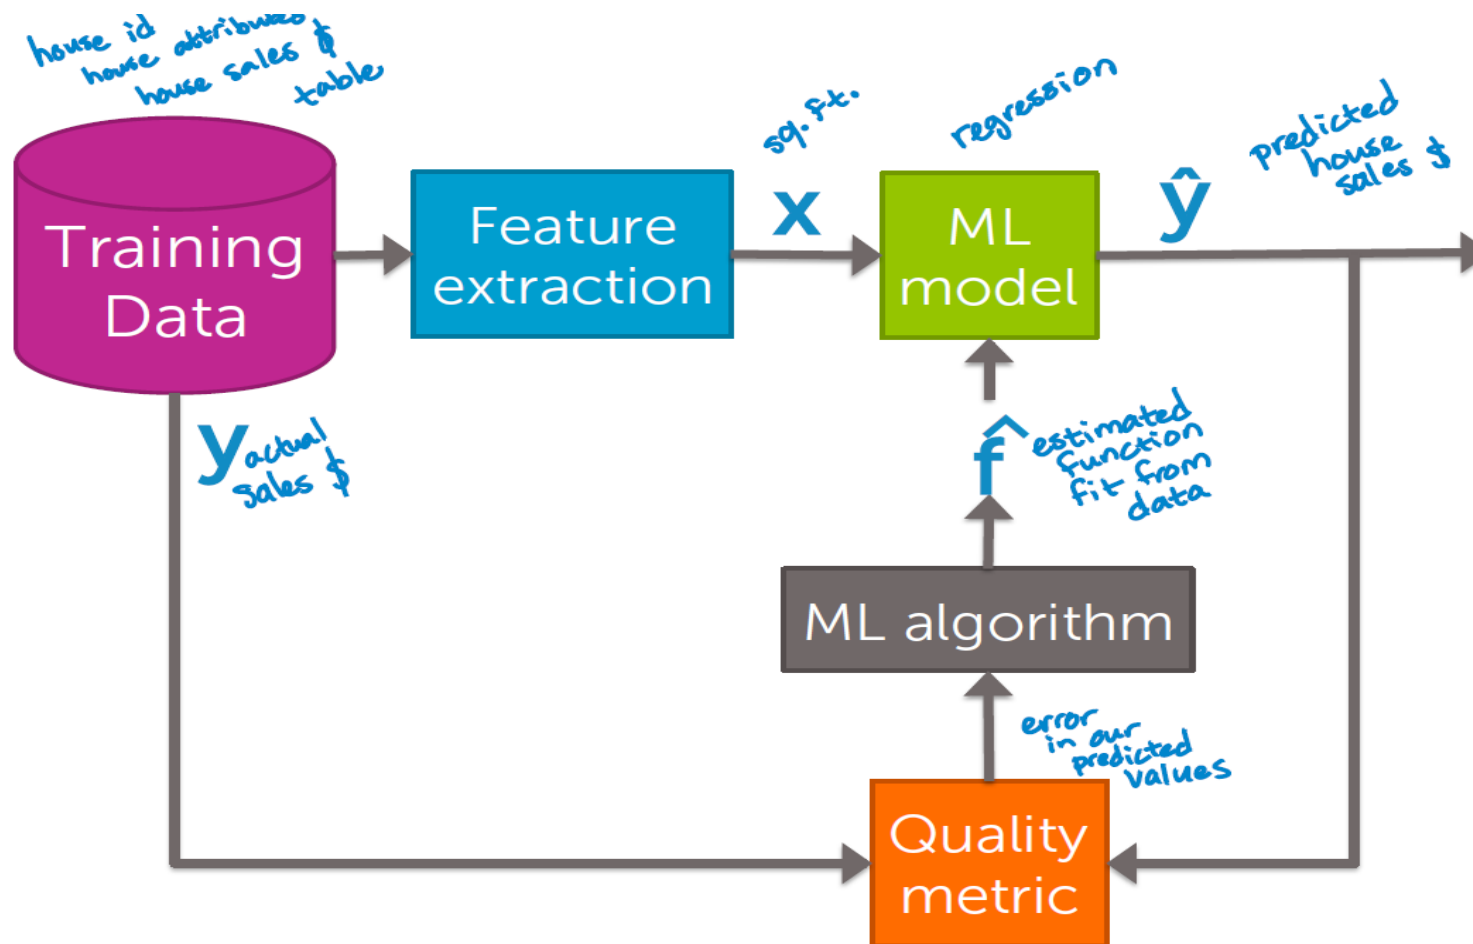
For a given model $f(x)$ estimate function $\hat{f}(x)$ from data



Assume model $f(x)$ is a quadratic function

How it works: baseline flow chart


11



SIMPLE LINEAR REGRESSION


Simple linear regression model

13




The cost of using a given line

14



Find „best” line


15



$$RSS(w_0, w_1) = \sum_{i=1}^N (y_i - [w_0 + w_1 x_i])^2$$

Interpreting the coefficients

16




$$\hat{y} = \hat{w}_0 \text{ when } X=0$$

not very meaningful

Interpreting the coefficients

17




Magnitude of fit parameters depend on the units of both features and observations

$$\begin{aligned} & \hat{\$}_{1001 \text{ sq.ft.}} - \hat{\$}_{1000 \text{ sq.ft.}} \\ &= \hat{w}_0 + \hat{w}_1 \cdot 1001 \text{ sq.ft.} \\ & \quad - (\hat{w}_0 + \hat{w}_1 \cdot 1000 \text{ sq.ft.}) \\ &= \hat{w}_1 \\ & \text{predicted change in the output} \\ & \quad \text{per unit change in input} \end{aligned}$$

ML algorithm: minimizing the cost

18

3D plot of RSS with tangent plane at minimum



Minimize function
over all possible w_0, w_1

$$\min_{w_0, w_1} \sum_{i=1}^N (y_i - [w_0 + w_1 x_i])^2$$

RSS(w_0, w_1) is a function
of 2 variables = $q(w_0, w_1)$

Convergence criteria

19

For convex functions,
optimum occurs when

$$\frac{dg(w)}{dw} = 0$$

In practice, stop when

$$\left| \frac{dg(w)}{dw} \right| < \epsilon$$

threshold to be set

That will be „good enough”
value of ϵ depends on the data we are looking at


Algorithm:

while not converged
 $w^{(t+1)} \leftarrow w^{(t)} - \eta \left. \frac{dg}{dw} \right|_{w^{(t)}}$

Moving to multiple dimensions

20

3D plot of RSS with tangent plane at minimum



$$\nabla g(w) = \begin{bmatrix} \frac{\partial g}{\partial w_0} \\ \frac{\partial g}{\partial w_1} \\ \vdots \\ \frac{\partial g}{\partial w_p} \end{bmatrix}$$

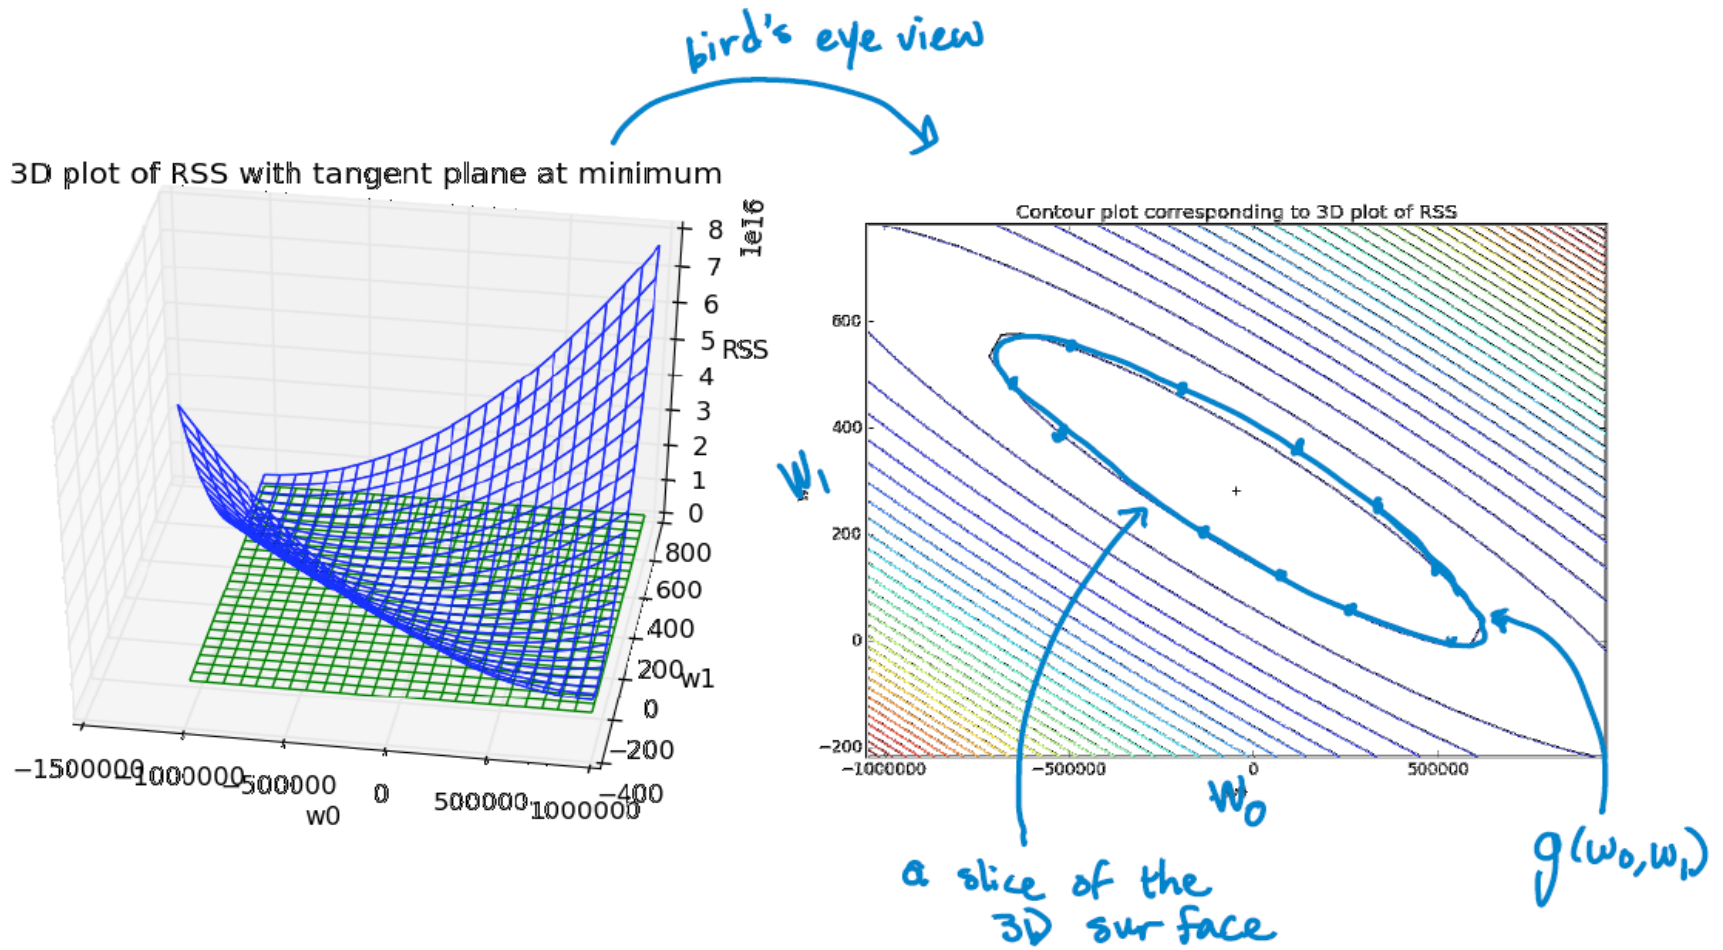
gradient $[w_0, w_1, \dots, w_p]$

$(p+1)$ -dimensional vector

partial derivative is like a derivate with respect to w_i treating all other variables as constants


Contour plots

21



Gradient descent

22



Algorithm:

while not converged

$$w^{(t+1)} \leftarrow w^{(t)} - \eta \nabla g(w^{(t)})$$

$$\begin{bmatrix} \vdots \\ \vdots \\ \vdots \end{bmatrix} \leftarrow \begin{bmatrix} \vdots \\ \vdots \\ \vdots \end{bmatrix} - \eta \begin{bmatrix} \vdots \\ \vdots \\ \vdots \end{bmatrix}$$

Convergence:
 $\|\nabla g(w)\| < \epsilon$

Compute the gradient

23

$$\text{RSS}(w_0, w_1) = \sum_{i=1}^N (y_i - [w_0 + w_1 x_i])^2$$

Taking the derivative w.r.t. w_0

$$\begin{aligned} & \sum_{i=1}^N 2 (y_i - [w_0 + w_1 x_i])' \cdot (-1) \\ &= -2 \sum_{i=1}^N (y_i - [w_0 + w_1 x_i]) \end{aligned}$$

Putting it together:

$$\nabla \text{RSS}(w_0, w_1) = \begin{bmatrix} -2 \sum_{i=1}^N [y_i - (w_0 + w_1 x_i)] \\ -2 \sum_{i=1}^N [y_i - (w_0 + w_1 x_i)] x_i \end{bmatrix}$$

Taking the derivative w.r.t. w_1


$$\begin{aligned} & \sum_{i=1}^N 2 (y_i - [w_0 + w_1 x_i])' \cdot (-x_i) \\ &= -2 \sum_{i=1}^N (y_i - [w_0 + w_1 x_i]) \underline{x_i} \end{aligned}$$

Approach 1: set gradient to 0

$$\nabla \text{RSS}(w_0, w_1) = \begin{bmatrix} -2 \sum_{i=1}^N [y_i - (w_0 + w_1 x_i)] \\ -2 \sum_{i=1}^N [y_i - (w_0 + w_1 x_i)] x_i \end{bmatrix}$$

This method is called „Closed form solution“

3D plot of RSS with tangent plane at minimum



top term: $\hat{w}_0 = \frac{\sum_{i=1}^N y_i}{N} - \hat{w}_1 \frac{\sum_{i=1}^N x_i}{N}$

average house sales price $\leftarrow \frac{\sum y_i}{N}$
 estimate of the slope $\leftarrow \hat{w}_1$
 average sq-ft. $\leftarrow \frac{\sum x_i}{N}$

bottom term: $\sum y_i x_i - \hat{w}_0 \sum x_i - \hat{w}_1 \sum x_i^2 = 0$

$\hat{w}_1 = \frac{\sum y_i x_i - \frac{\sum y_i \sum x_i}{N}}{\sum x_i^2 - \frac{\sum x_i \sum x_i}{N}}$

Note:

$$\sum_{i=1}^N y_i$$

$$\sum_{i=1}^N x_i$$


$$\sum_{i=1}^N y_i x_i$$

$$\sum_{i=1}^N x_i^2$$

Approach 2: gradient descent

25

$$\nabla \text{RSS}(w_0, w_1) = \begin{bmatrix} -2 \sum_{i=1}^N [y_i - \hat{y}_i(w_0, w_1)] \\ -2 \sum_{i=1}^N [y_i - \hat{y}_i(w_0, w_1)] x_i \end{bmatrix}$$



while not converged $(-2) \cdot (-\eta)$

$$\begin{bmatrix} w_0^{(t+1)} \\ w_1^{(t+1)} \end{bmatrix} \leftarrow \begin{bmatrix} w_0^{(t)} \\ w_1^{(t)} \end{bmatrix} + 2\eta \begin{bmatrix} \sum_{i=1}^N [y_i - \hat{y}_i(w_0^{(t)}, w_1^{(t)})] \\ \sum_{i=1}^N [y_i - \hat{y}_i(w_0^{(t)}, w_1^{(t)})] x_i \end{bmatrix}$$

If overall, under predicting \hat{y}_i , then $\sum [y_i - \hat{y}_i]$ is positive

→ w_0 is going to increase

similar intuition for w_1 , but multiply by x_i


Comparing the approaches

26

- For most ML problems, cannot solve $\text{gradient} = 0$
- Even if solving $\text{gradient} = 0$ is feasible, gradient descent can be more efficient
- Gradient descent relies on choosing stepsize and convergence criteria

Asymmetric cost functions

27



We can weight differently positive and negative errors in RSS calculations.

What if cost of listing house too high has **bigger cost**?


Too high \rightarrow no offers ($\$=0$)

Too low \rightarrow offers for lower \$

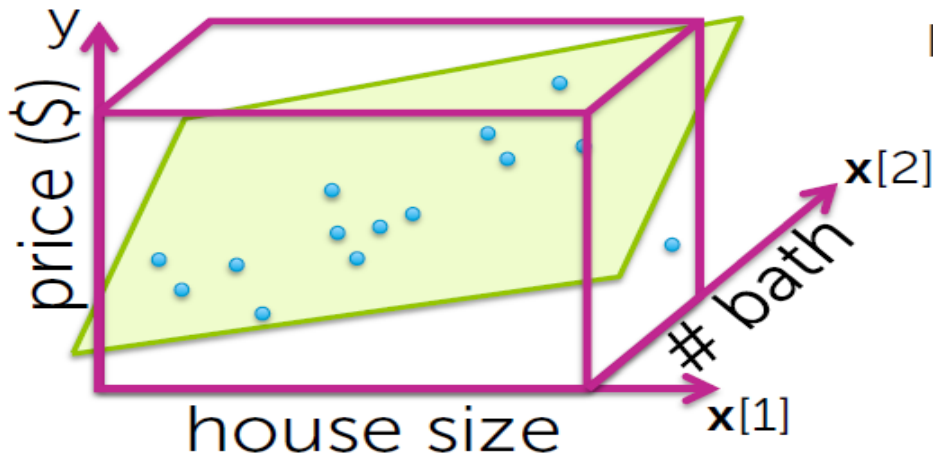
MULTIPLE REGRESSION

Multiple regression

29



Fit more complex relationships than just a line



Incorporate more inputs

- Square feet
- # bathrooms
- # bedrooms
- Lot size
- Year built
- ...

Polynomial regression

30

Model:

$$y_i = w_0 + w_1 x_i + w_2 x_i^2 + \dots + w_p x_i^p + \epsilon_i$$

treat as different **features**



feature 1 = 1 (constant) *parameter 1 = w_0*

feature 2 = x *parameter 2 = w_1*

feature 3 = x^2 *parameter 3 = w_2*

...

...

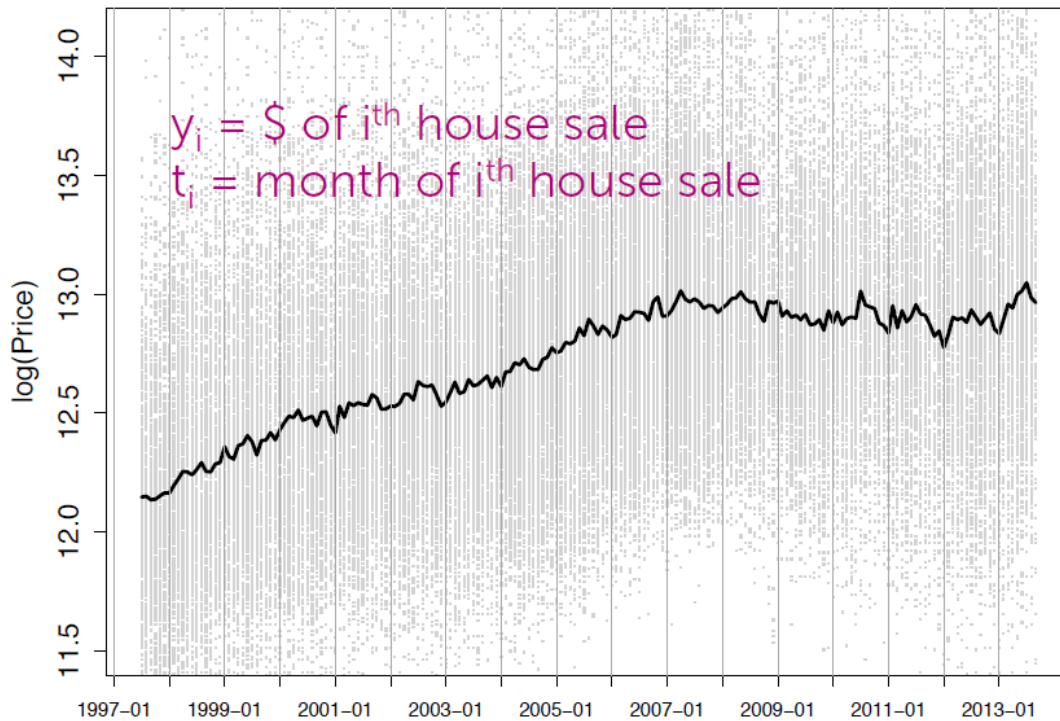
feature $p+1 = x^p$

parameter $p+1 = w_p$

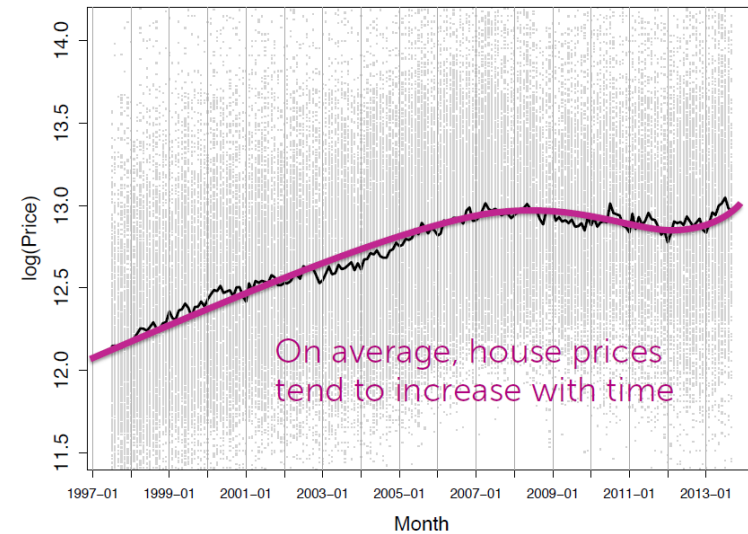
Other functional forms of one input

31

□ Trends in time series



Month ← House sales recorded monthly

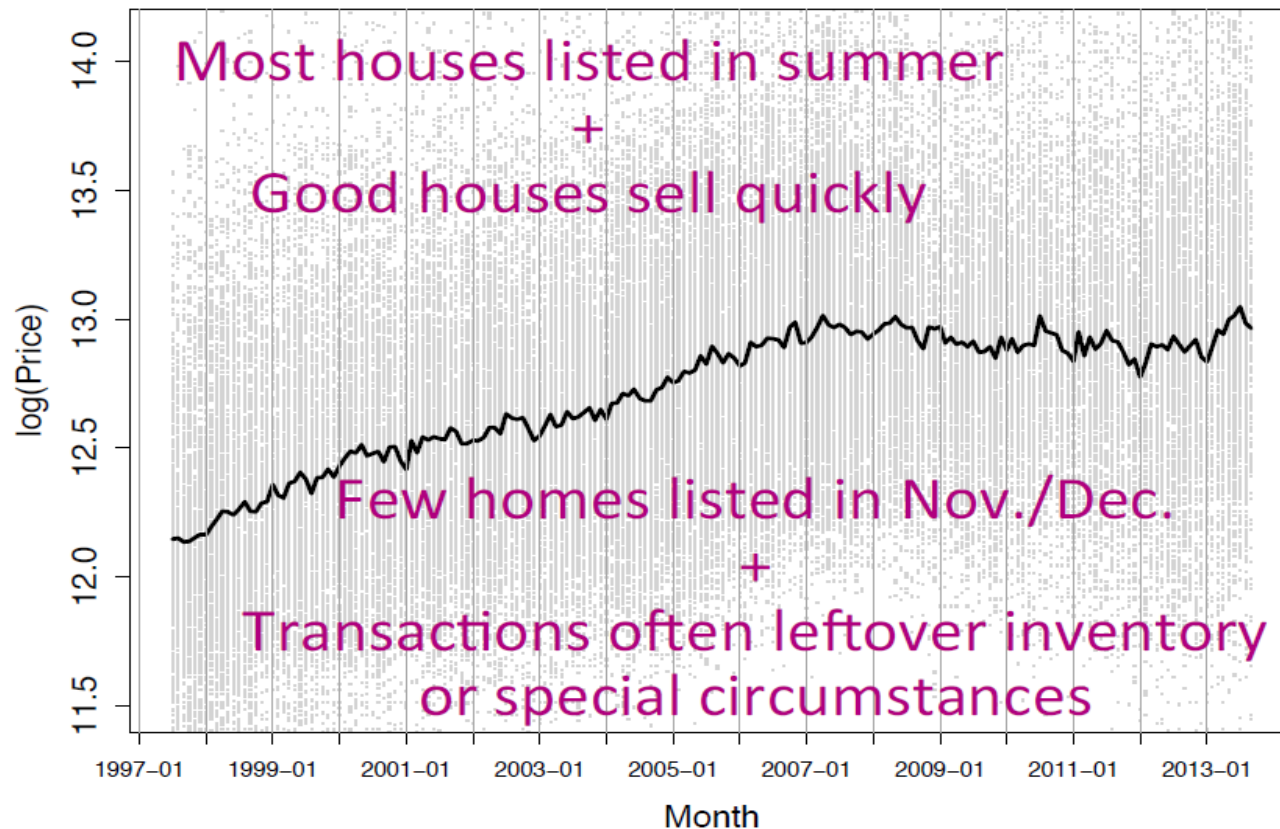


This trend can be modeled with polynomial function.

Other functional forms of one input

32

□ Seasonality



Example of detrending

33


Model:

$$y_i = w_0 + w_1 t_i + w_2 \sin(2\pi t_i / 12 - \Phi) + \varepsilon_i$$

Linear trend

Unknown phase/shift

Seasonal component =
Sinusoid with period 12
(resets annually)



Trigonometric identity: $\sin(a - b) = \sin(a)\cos(b) - \cos(a)\sin(b)$

$$\rightarrow \sin(2\pi t_i / 12 - \Phi) = \sin(2\pi t_i / 12)\cos(\Phi) - \cos(2\pi t_i / 12)\sin(\Phi)$$

Example of detrending

34

Equivalently,


$$y_i = w_0 + w_1 t_i + w_2 \sin(2\pi t_i / 12) + w_3 \cos(2\pi t_i / 12) + \varepsilon_i$$

feature 1 = 1 (constant)

feature 2 = t

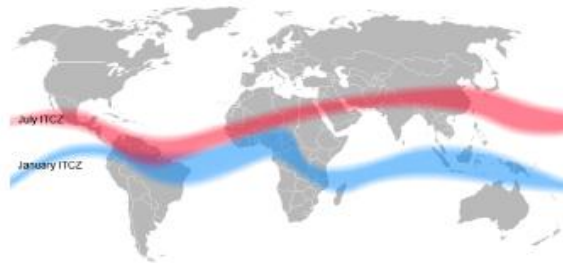
feature 3 = $\sin(2\pi t/12)$

feature 4 = $\cos(2\pi t/12)$

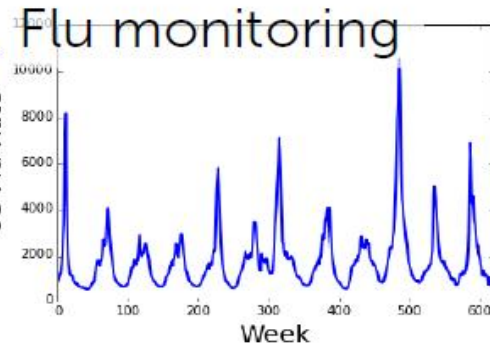


Other examples of seasonality


35



Weather modeling
(e.g., temperature, rainfall)



Demand forecasting
(e.g., jacket purchases)



Generic basic expansion

36

Model:

$$y_i = w_0 h_0(x_i) + w_1 h_1(x_i) + \dots + w_D h_D(x_i) + \varepsilon_i$$
$$= \sum_{j=0}^D w_j h_j(x_i) + \varepsilon_i$$

feature 1 = $h_0(x)$...often 1 (constant)

feature 2 = $h_1(x)$... e.g., x

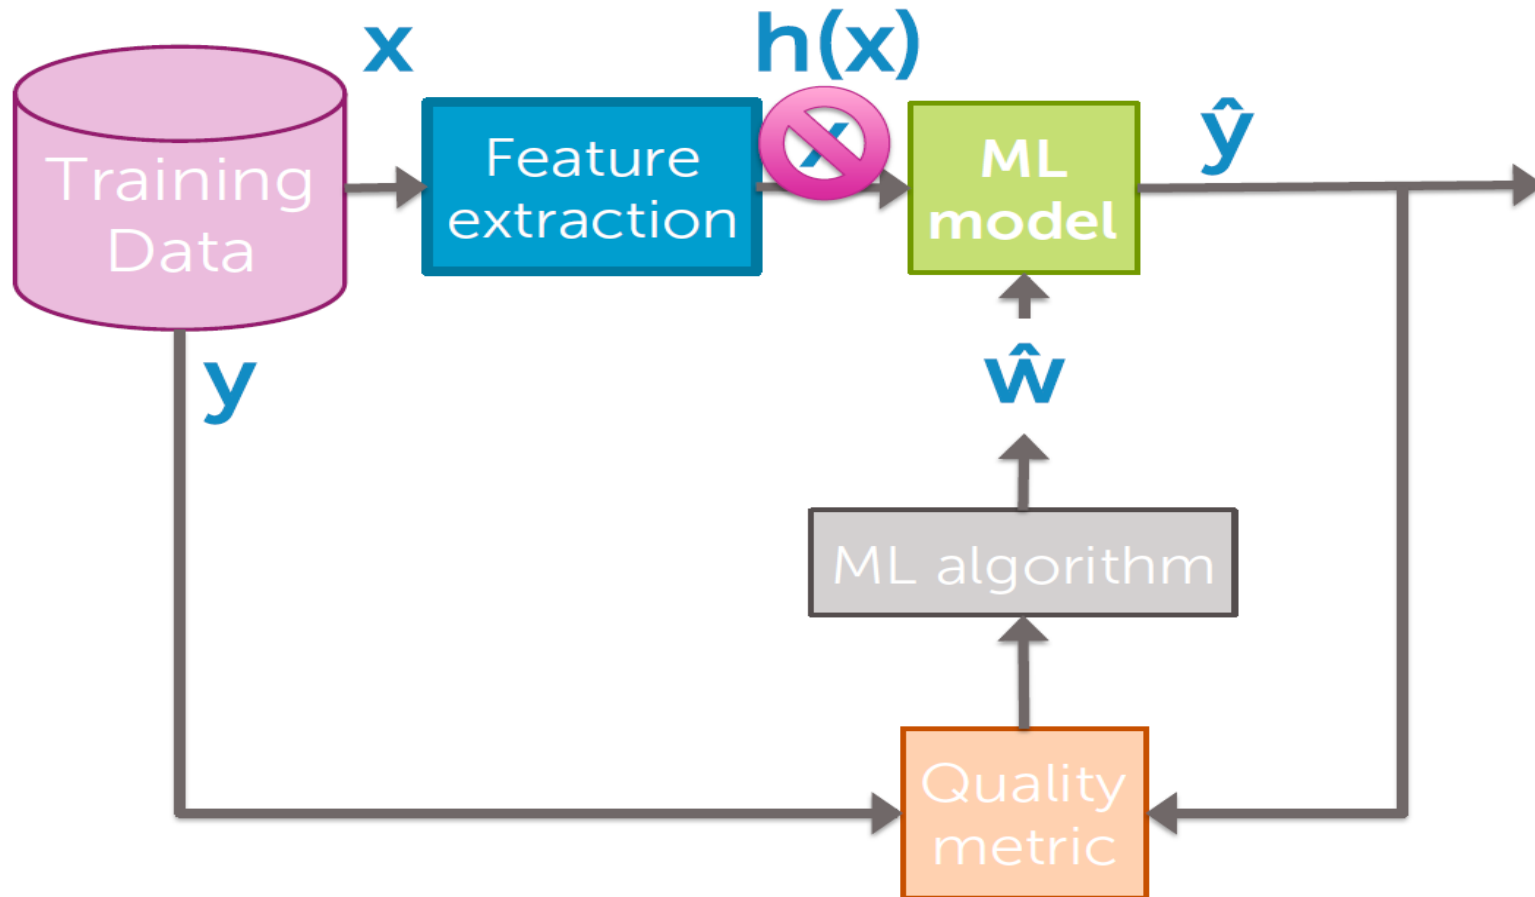
feature 3 = $h_2(x)$... e.g., x^2 or $\sin(2\pi x/12)$

...

feature $D+1$ = $h_D(x)$... e.g., x^p

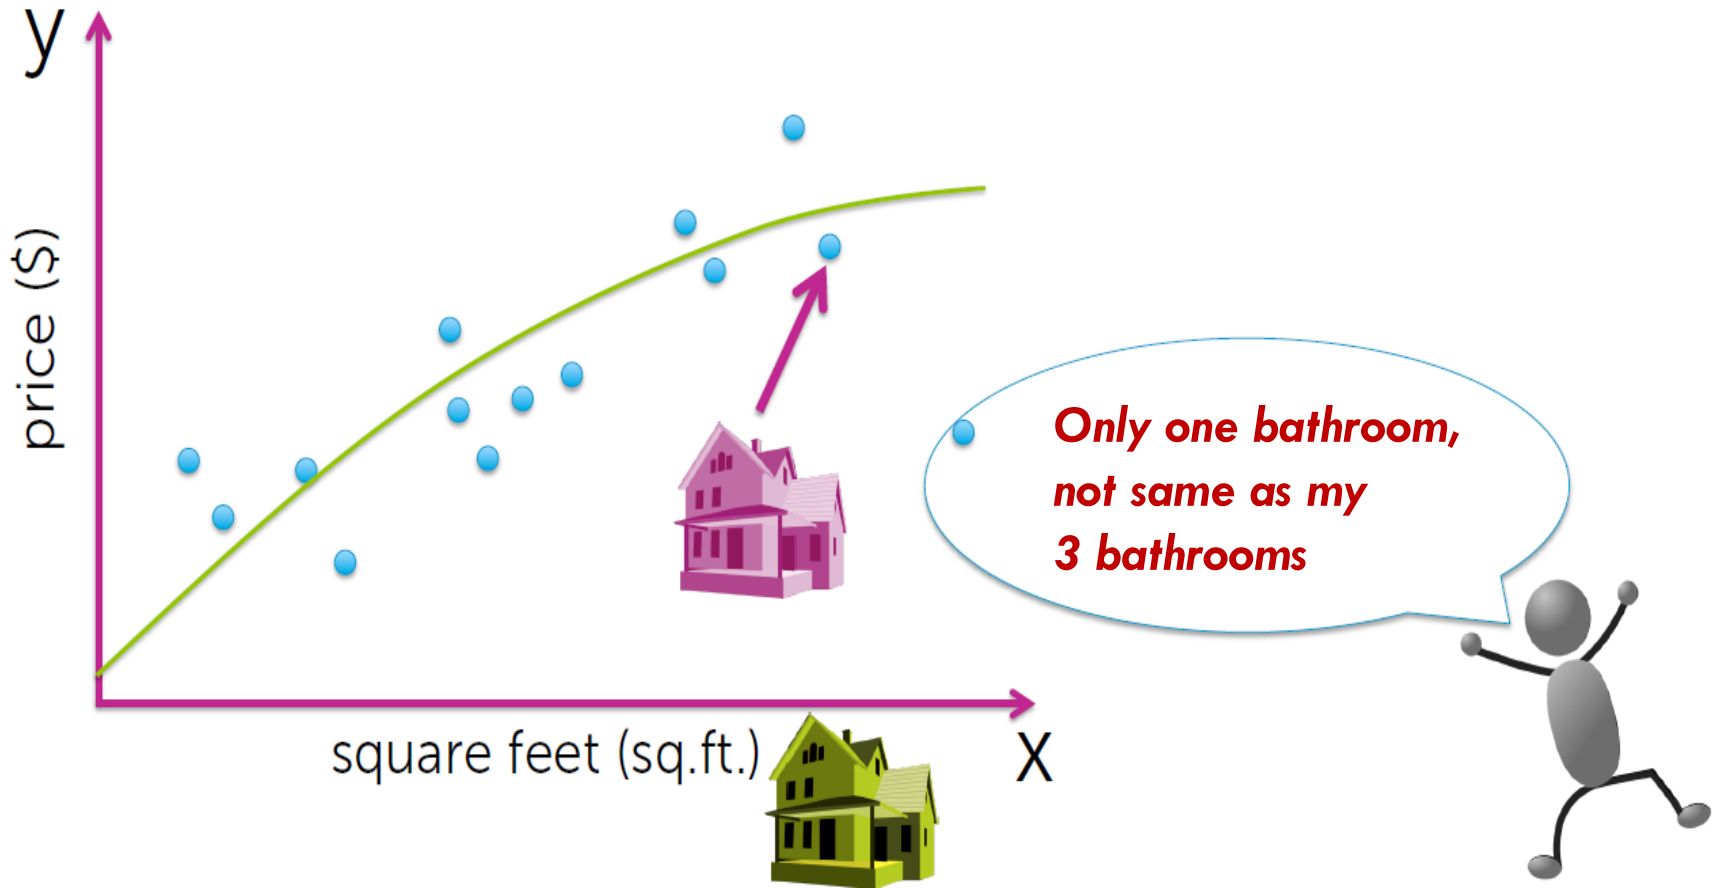
More realistic flow chart

37



Incorporating multiple inputs


38



— ...

Incorporating multiple inputs

39



Many possible inputs

- Square feet
- # bathrooms
- # bedrooms
- Lot size
- Year built
- ...

General notation

40

Output: y ← scalar

Inputs: $\mathbf{x} = (\mathbf{x}[1], \mathbf{x}[2], \dots, \mathbf{x}[d])$

↖ d-dim vector

Notational conventions:

$\mathbf{x}[j]$ = j^{th} input (*scalar*)

$h_j(\mathbf{x})$ = j^{th} feature (*scalar*)

\mathbf{x}_i = input of i^{th} data point (*vector*)

$\mathbf{x}_i[j]$ = j^{th} input of i^{th} data point (*scalar*)

Simple hyperplane

41

Model:

Noise term



$$y_i = w_0 + w_1 \mathbf{x}_i[1] + \dots + w_d \mathbf{x}_i[d] + \epsilon_i$$

feature 1 = 1

feature 2 = $\mathbf{x}[1]$... e.g., sq. ft.

feature 3 = $\mathbf{x}[2]$... e.g., #bath

...

feature $d+1 = \mathbf{x}[d]$... e.g., lot size

More generally: D-dimensional curve

42

Model:

$$y_i = w_0 h_0(\mathbf{x}_i) + w_1 h_1(\mathbf{x}_i) + \dots + w_D h_D(\mathbf{x}_i) + \varepsilon_i$$
$$= \sum_{j=0}^D w_j h_j(\mathbf{x}_i) + \varepsilon_i$$

More on notation

observations (\mathbf{x}_i, y_i) : N

inputs $\mathbf{x}[j]$: d

features $h_j(\mathbf{x})$: D

feature 1 = $h_0(\mathbf{x})$... e.g., 1

feature 2 = $h_1(\mathbf{x})$... e.g., $\mathbf{x}[1]$ = sq. ft.

feature 3 = $h_2(\mathbf{x})$... e.g., $\mathbf{x}[2]$ = #bath


or, $\log(\mathbf{x}[7])$ $\mathbf{x}[2]$ = $\log(\#bed) \times \#bath$

...

feature $D+1$ = $h_D(\mathbf{x})$... some other function of $\mathbf{x}[1], \dots, \mathbf{x}[d]$

Fitting in D-dimensions

43



Rewriting in vector notation

44

For observation i

$$y_i = \sum_{j=0}^D w_j h_j(\mathbf{x}_i) + \varepsilon_i$$

Diagram illustrating the vector notation for the observation equation:

$$y_i = \begin{bmatrix} w_0 & w_1 & w_2 & \dots & w_D \end{bmatrix} \begin{bmatrix} h_0(x_i) \\ h_1(x_i) \\ h_2(x_i) \\ \vdots \\ h_D(x_i) \end{bmatrix} + \varepsilon_i$$

The weight vector w^T is shown as a row of blue boxes with labels $w_0, w_1, w_2, \dots, w_D$ below them. The feature vector $h(x_i)$ is shown as a column of green boxes with labels $h_0(x_i), h_1(x_i), h_2(x_i), \dots, h_D(x_i)$ to its right. The error term ε_i is in a grey box.

The expanded form is shown as:

$$y_i = \begin{bmatrix} h_0(x_i) & h_1(x_i) & \dots & h_D(x_i) \end{bmatrix} \begin{bmatrix} w_0 \\ w_1 \\ \vdots \\ w_D \end{bmatrix} + \varepsilon_i$$

The expanded form is also written as:

$$y_i = h_0(x_i)w_0 + h_1(x_i)w_1 + \dots + h_D(x_i)w_D + \varepsilon_i$$


The final simplified vector notation is:

$$y_i = w^T h(x_i) + \varepsilon_i$$

Rewriting in matrix notation

45

For all observations together




Here is our ML algorithm

\Rightarrow
$$Y = HW + E$$

Cost function in D-dimension

47

RSS in vector notation



$$\text{RSS}(\underline{w}) = \sum_{i=1}^N (y_i - \underbrace{h^T(x_i) w}_{\hat{y}_i(w)})^2$$

$\hat{y}_i =$

$h^T(x_i)$							
$h_0(x_i) \ h_1(x_i) \ \dots \ h_p(x_i)$							

	w
	w_0
	w_1
	w_2
	\vdots
	w_D


Cost function in D-dimension

48

RSS in matrix notation

$$\begin{aligned} \text{RSS}(\mathbf{w}) &= \sum_{i=1}^N (y_i - h(\mathbf{x}_i)^T \mathbf{w})^2 \\ &= (\mathbf{y} - \mathbf{H}\mathbf{w})^T (\mathbf{y} - \mathbf{H}\mathbf{w}) \end{aligned}$$

Why? (part 1)



$$\hat{\mathbf{y}} = \mathbf{H}\mathbf{w}$$

$$(\mathbf{y} - \mathbf{H}\mathbf{w}) = (\mathbf{y} - \hat{\mathbf{y}}) = \begin{bmatrix} \text{residual}_1 \\ \text{residual}_2 \\ \vdots \\ \text{residual}_N \end{bmatrix}$$

residual i
 $= y_i - \hat{y}_i$

Regression model for D-dimension

49

Gradient of RSS

$$\begin{aligned}\nabla \text{RSS}(\mathbf{w}) &= \nabla [(\mathbf{y} - \mathbf{H}\mathbf{w})^\top (\mathbf{y} - \mathbf{H}\mathbf{w})] \\ &= -2\mathbf{H}^\top (\mathbf{y} - \mathbf{H}\mathbf{w})\end{aligned}$$

Why? By analogy to 1D case:

$$\begin{aligned}\frac{d}{dw} (y-hw)(y-hw) &= \frac{d}{dw} (y-hw)^2 = 2 \cdot (y-hw)' (-h) \\ &= -2h(y-hw)\end{aligned}$$


↑
Scalars

Regression model for D-dimension

50

Approach 1: set gradient to zero

3D plot of RSS with tangent plane at minimum



Closed form solution

$$\nabla \text{RSS}(\mathbf{w}) = -2\mathbf{H}^T(\mathbf{y} - \mathbf{H}\mathbf{w}) = 0$$

Solve for \mathbf{w} :

$$-\cancel{2}\mathbf{H}^T\mathbf{y} + \cancel{2}\mathbf{H}^T\mathbf{H}\hat{\mathbf{w}} = 0$$

$$\mathbf{H}^T\mathbf{H}\hat{\mathbf{w}} = \mathbf{H}^T\mathbf{y}$$

$$\underbrace{(\mathbf{H}^T\mathbf{H})^{-1}}_{\mathbf{I}} \mathbf{H}^T\mathbf{H}\hat{\mathbf{w}} = (\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T\mathbf{y}$$

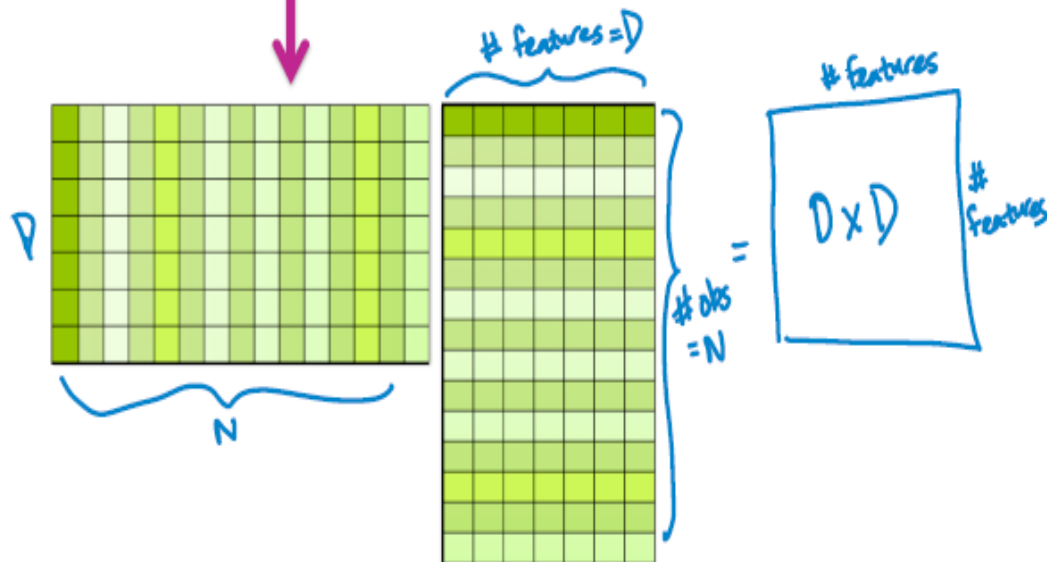
$$\hat{\mathbf{w}} = (\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T\mathbf{y}$$

$\bullet A^{-1}A = I$
 $\bullet I \cdot v = v$
 $I \cdot v = v$

Closed-form solution

51

$$\hat{W} = (\underbrace{H^T H}_{D \times D})^{-1} H^T y$$



This matrix might not be invertible.

Invertible if:
In most cases is $N > D$

really,
of linearly
ind. observations

Complexity of inverse:


$$O(D^3)$$

This might not be CPU feasible.

Regression model for D-dimension

52

Approach 2: gradient descent



We initialise our solution somewhere and then ...

while not converged

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} - \eta \underbrace{\nabla \text{RSS}(\mathbf{w}^{(t)})}_{-2\mathbf{H}^T(\mathbf{y} - \mathbf{H}\mathbf{w})}$$

$$\leftarrow \mathbf{w}^{(t)} + 2\eta \mathbf{H}^T (\mathbf{y} - \underbrace{\mathbf{H}\mathbf{w}^{(t)}}_{\hat{\mathbf{y}}(\mathbf{w}^{(t)})})$$

Gradient descent

53

$$\text{RSS}(\mathbf{w}) = \sum_{i=1}^N (y_i - \mathbf{h}(\mathbf{x}_i)^T \mathbf{w})^2$$
$$= \sum_{i=1}^N (y_i - w_0 h_0(x_i) - w_1 h_1(x_i) - \dots - w_D h_D(x_i))^2$$

Partial with respect to w_j

$$\sum_{i=1}^N 2 (y_i - w_0 h_0(x_i) - w_1 h_1(x_i) - \dots - w_D h_D(x_i)) \cdot (-h_j(x_i))$$
$$= -2 \sum_{i=1}^N h_j(x_i) (y_i - \mathbf{h}(\mathbf{x}_i)^T \mathbf{w})$$


Update to j^{th} feature weight:

$$w_j^{(t+1)} \leftarrow w_j^{(t)} - \eta \left(-2 \sum_{i=1}^N h_j(x_i) (y_i - \underbrace{\mathbf{h}^T(\mathbf{x}_i) \mathbf{w}^{(t)}}_{\hat{y}_i(\mathbf{w}^{(t)})}) \right)$$

Summary of gradient descent

54

Extremely useful algorithm in several applications



```
init  $\mathbf{w}^{(1)} = \mathbf{0}$  (or randomly, or smartly),  $t = 1$   
while  $\|\nabla \text{RSS}(\mathbf{w}^{(t)})\| > \epsilon$   
    for  $j = 0, \dots, D$   
         $\text{partial}[j] = -2 \sum_{i=1}^N h_j(\mathbf{x}_i)(y_i - \hat{y}_i(\mathbf{w}^{(t)}))$   
         $\mathbf{w}_j^{(t+1)} \leftarrow \mathbf{w}_j^{(t)} - \eta \text{partial}[j]$   
     $t \leftarrow t + 1$ 
```

Handwritten notes:
- A blue arrow points from the word "tolerance" to the ϵ in the while loop.
- A blue arrow points from the expression $\sqrt{\text{partial}[0]^2 + \dots + \text{partial}[D]^2}$ to the norm symbol in the while loop.

ACCESSING PERFORMANCE

Measuring loss

“Remember that all models are wrong; the practical question is how wrong do they have to be to not be useful.” George Box, 1987.

56

Loss function:

Cost of using \hat{w} at x
when y is true

$$L(y, f_{\hat{w}}(\mathbf{x}))$$

actual value $\hat{f}(\mathbf{x}) =$ predicted value \hat{y}

Examples:

(assuming loss for underpredicting = overpredicting)

Absolute error: $L(y, f_{\hat{w}}(\mathbf{x})) = |y - f_{\hat{w}}(\mathbf{x})|$

Squared error: $L(y, f_{\hat{w}}(\mathbf{x})) = (y - f_{\hat{w}}(\mathbf{x}))^2$


Symmetric loss functions



Accessing the loss

57

Use training data



Compute training error

58

1. Define a loss function $L(y, f_{\hat{\mathbf{w}}}(\mathbf{x}))$
 - E.g., squared error, absolute error, ...

2. Training error

= avg. loss on houses in training set


$$= \frac{1}{N} \sum_{i=1}^N L(y_i, f_{\hat{\mathbf{w}}}(\mathbf{x}_i))$$

 fit using training data

Training error

59

Use **squared error loss** $(y - f_{\hat{w}}(\mathbf{x}))^2$




Convention is to take average here

$$\begin{aligned} \text{Training error } (\hat{w}) &= 1/N * \\ & [(\$_{\text{train } 1} - f_{\hat{w}}(\text{sq.ft.}_{\text{train } 1}))^2 \\ & + (\$_{\text{train } 2} - f_{\hat{w}}(\text{sq.ft.}_{\text{train } 2}))^2 \\ & + (\$_{\text{train } 3} - f_{\hat{w}}(\text{sq.ft.}_{\text{train } 3}))^2 \\ & + \dots \text{ include all} \\ & \text{training houses}] \end{aligned}$$

Training error vs. model complexity

60




Decrease as you increase your model complexity. Very intuitive why it is like that.

Is training error a good measure?

61

Issue: Training error is overly optimistic
because \hat{w} was fit to training data




Is there something particularly wrong about having x_t square feet ???

Small training error \neq good predictions
unless training data includes everything you might ever see

Generalisation (true) error

62


Really want estimate of loss
over all possible (🏠, \$) pairs



Lots of houses
in neighborhood,
but not in dataset

Generalisation error vs model complexity

63



However ... in contrast to the training error, in practice we cannot really compute true generalisation error. We don't have data on all possible houses in the area.

Can't compute!

Forming a test set

64

Hold out some (🏠, \$) that are *not* used for fitting the model



We want to approximate generalisation error.

Test set: proxy for „everything you might see“

Training set



Test set



Compute test error

65

Test error

= avg. loss on houses in test set

$$= \frac{1}{N_{test}} \sum_{i \text{ in test set}} L(y_i, f_{\hat{w}}(\mathbf{x}_i))$$



test points




fit using training data

**has never seen
test data!**

Training, true and test error vs. model complexity. Notion of overfitting.

66



Training/test splits

67



↑
Too few → \hat{w} poorly estimated



↑
Too few → test error bad approximation of generalization error



Typically, just enough test points to form a reasonable estimate of generalization error

If this leaves too few for training, other methods like **cross validation** (will see later...)

Three sources of errors

68


In forming predictions, there are 3 sources of error:

1. Noise
2. Bias
3. Variance

Data are inherently noisy

69

There is some true relationship between sq.ft and value of the house, specific to the given house.




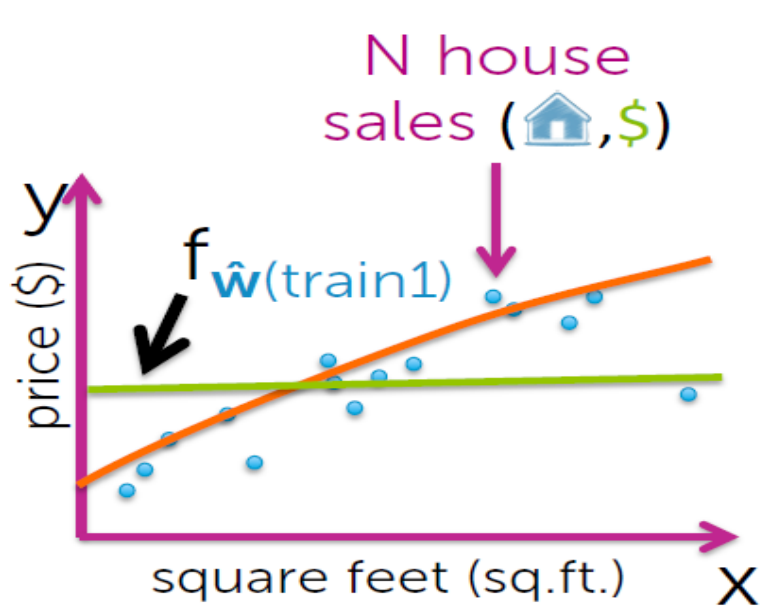
We cannot reduce it by choosing better model or procedure, It is beyond our control.

Bias contribution

70

This contribution we can control.


Assume we fit a constant function



Bias contribution

71


Over all possible size N training sets, what do I expect my fit to be?



Bias contribution

72


$\text{Bias}(\mathbf{x}) = f_{\mathbf{w}(\text{true})}(\mathbf{x}) - f_{\bar{\mathbf{w}}}(\mathbf{x})$ ← Is our approach flexible enough to capture $f_{\mathbf{w}(\text{true})}$?
If not, error in predictions.



Variance contribution

73


How much do specific fits vary from the expected fit?



Variance contribution

74

How much do specific fits vary from the expected fit?





Variance of high complexity models

75

Assume we fit a high-order polynomial

For each train remove
few random houses




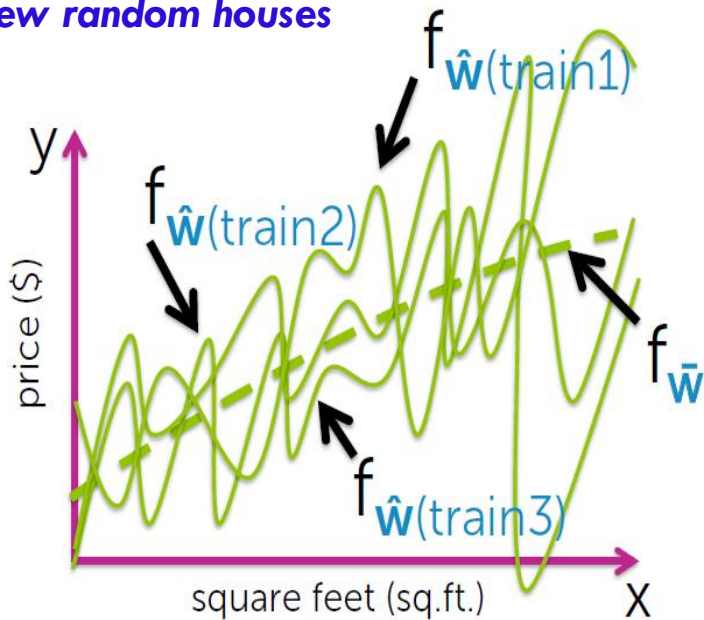
high
complexity
→
high variance

Bias of high complexity models

76

Assume we fit a high-order polynomial


*For each train remove
few random houses*



**High complexity models are very flexible,
pick better average trends.**

Bias –variance tradeoff

77



MSE = mean square error


**Machine Learning
is all about this tradeoff**

But....

Just like with
generalization error,
we cannot compute
bias and variance

Errors vs amount of data

78



The regression/ML workflow

79

1. Model selection

Often, need to choose tuning parameters λ controlling model complexity (e.g. degree of polynomial)

2. Model assessment

Having selected a model, assess the generalization error

Hypothetical implementation

80

Training set

Test set

1. Model selection

For each considered model complexity λ :

- i. Estimate parameters $\hat{\mathbf{w}}_\lambda$ on **training data**
- ii. Assess performance of $\hat{\mathbf{w}}_\lambda$ on **test data**
- iii. Choose λ^* to be λ with **lowest test error**


2. Model assessment

Compute test error of $\hat{\mathbf{w}}_{\lambda^*}$ (fitted model for selected complexity λ^*) to approx. generalization error

Overly optimistic!

Practical implementation

81



Typical splits

82



80%

10%

10%

50%

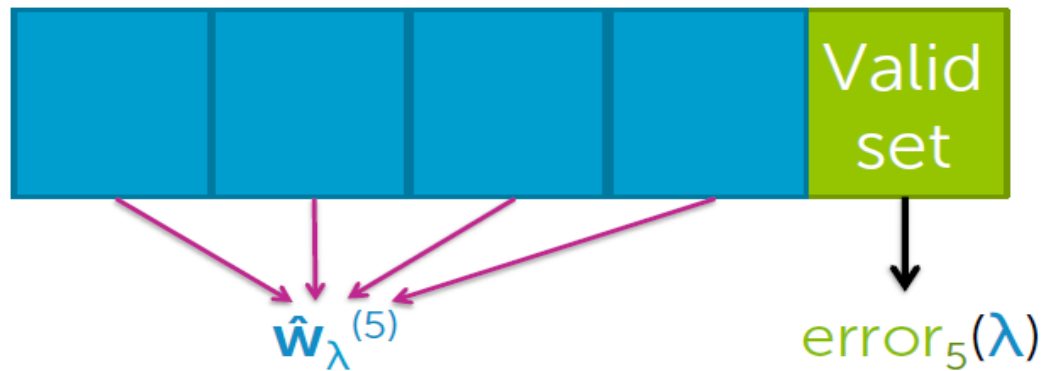
25%

25%

K-fold cross validation

83

K-fold cross validation



For $k=1, \dots, K$

1. Estimate $\hat{w}_\lambda^{(k)}$ on the training blocks
2. Compute error on validation block: $\text{error}_k(\lambda)$

Compute average error: $\text{CV}(\lambda) = \frac{1}{K} \sum_{k=1}^K \text{error}_k(\lambda)$

What value of K

84

Formally, the best approximation occurs for validation sets of size 1 ($K=N$)

leave-one-out
cross validation

Computationally intensive

- requires computing N fits of model per λ

Typically, $K=5$ or 10

5-fold CV

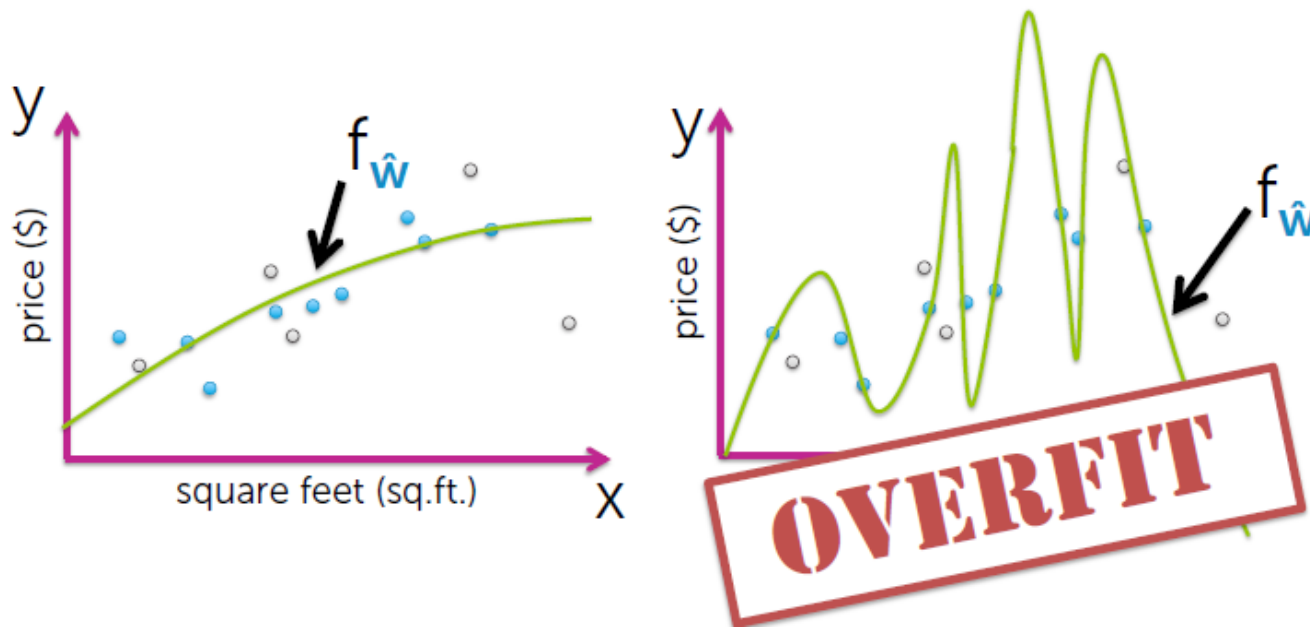
10-fold CV

RIDGE REGRESSION

Flexibility of high-order polynomials

86

$$y_i = w_0 + w_1 x_i + w_2 x_i^2 + \dots + w_p x_i^p + \varepsilon_i$$



Symptoms for overfitting: often associated with very large value of estimated parameters \hat{w}

How does # of observations influence overfitting?


87

Few observations (N small)

→ rapidly overfit as model complexity increases


Many observations (N very large)

→ harder to overfit



Lets improve quality metric blok

88

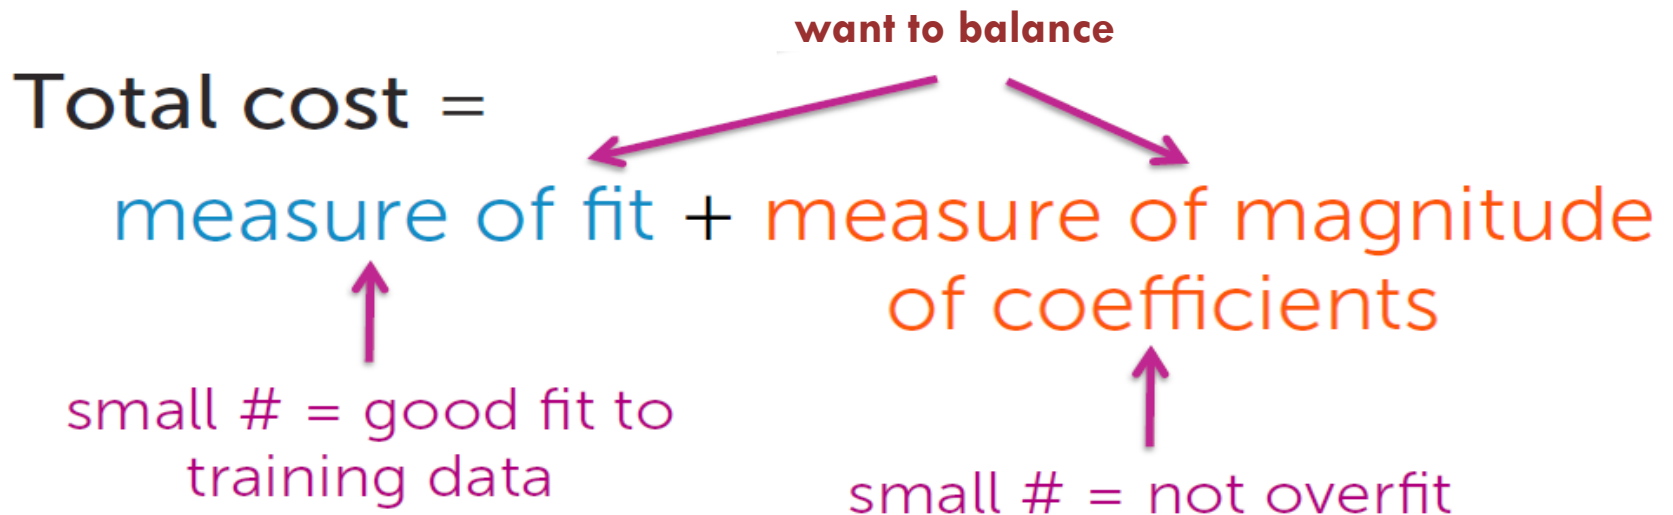


Desire total cost format

89

Want to balance:

- i. How well function fits data
- ii. Magnitude of coefficients



Measure of magnitude of regression coefficients

90

What summary # is indicative of size of regression coefficients?

- Sum? $w_0 = 1,527,301$ $w_1 = -1,605,253$
 $w_0 + w_1 = \text{small } \#$

← **But ... the coefficients are very large**

- Sum of absolute value?
 $|w_0| + |w_1| + \dots + |w_D| = \sum_{j=0}^D |w_j| \triangleq \|w\|_1$ L_1 norm ... discuss more in next module
- Sum of squares (L_2 norm)
 $w_0^2 + w_1^2 + \dots + w_D^2 = \sum_{j=0}^D w_j^2 \triangleq \|w\|_2^2$ L_2 norm ... focus of this module

Consider specific total cost

91

Total cost =

$$\underbrace{\text{measure of fit}}_{\text{RSS}(\mathbf{w})} + \underbrace{\text{measure of magnitude of coefficients}}_{\|\mathbf{w}\|_2^2}$$

Consider resulting objectives

92

What if $\hat{\mathbf{w}}$ selected to minimize

$$\text{RSS}(\mathbf{w}) + \lambda \|\mathbf{w}\|_2^2$$

λ tuning parameter = balance of fit and magnitude

Ridge regression
(a.k.a L_2 regularization)

If $\lambda=0$:

reduces to minimizing $\text{RSS}(\mathbf{w})$, as before (old solution) $\rightarrow \hat{\mathbf{w}}^{\text{LS}}$ ← least squares

If $\lambda=\infty$:

For solutions where $\hat{\mathbf{w}} \neq 0$, then total cost is ∞

If $\hat{\mathbf{w}}=0$, then total cost = $\text{RSS}(0)$ \rightarrow solution is $\hat{\mathbf{w}}=0$

If λ in between: Then $0 \leq \|\hat{\mathbf{w}}\|_2^2 \leq \|\hat{\mathbf{w}}^{\text{LS}}\|_2^2$

Ridge regression: bias-variance tradeoff

93

Large λ :

high bias, low variance

(e.g., $\hat{\mathbf{w}} = 0$ for $\lambda = \infty$)

In essence, λ
controls model
complexity

Small λ :

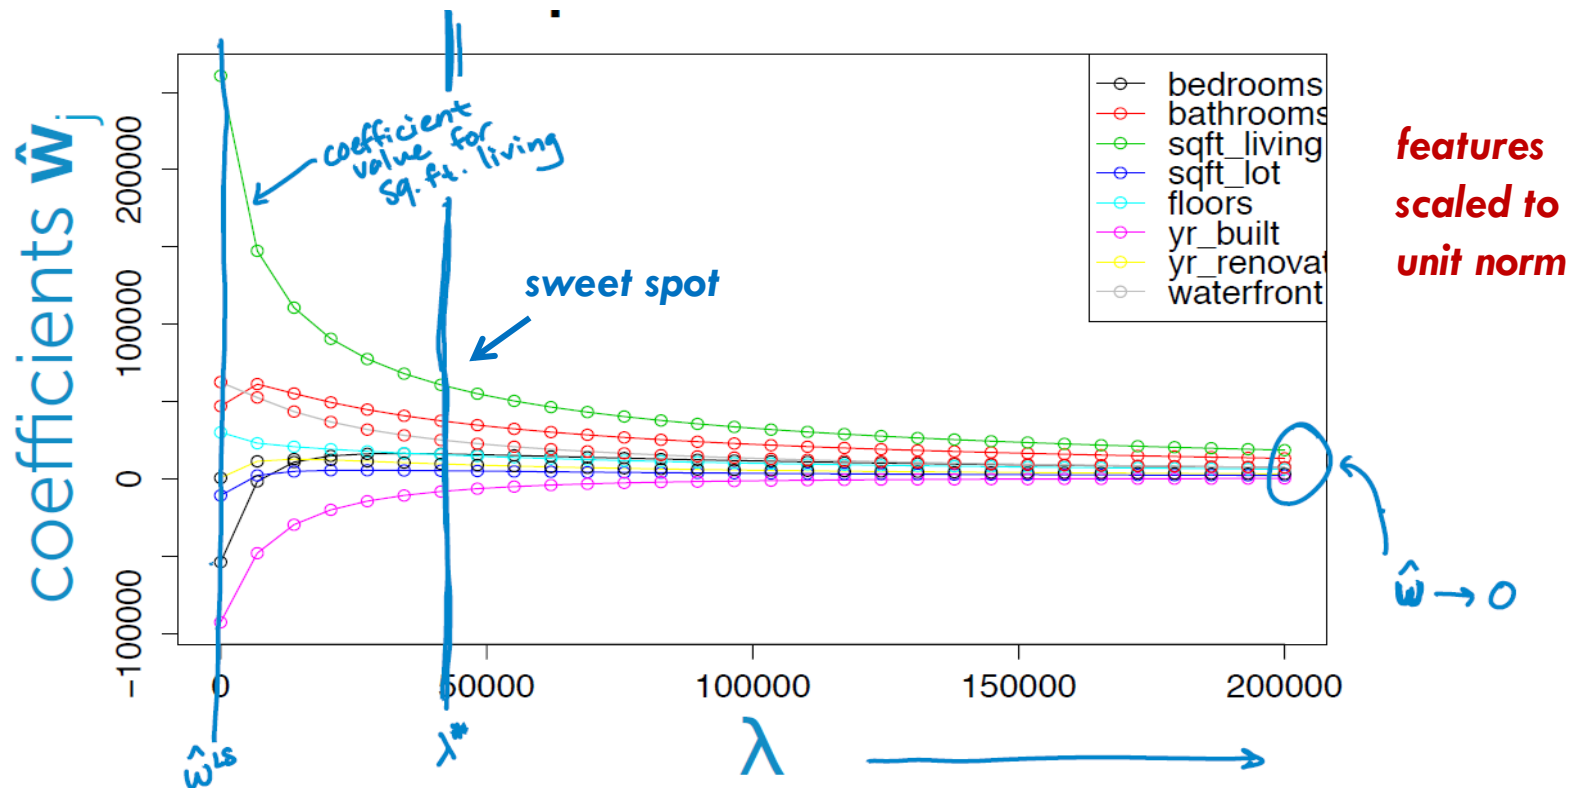
low bias, high variance

(e.g., standard least squares (RSS) fit of
high-order polynomial for $\lambda = 0$)

Ridge regression: coefficients path


94

What happens if we refit our high-order polynomial, but now using ridge regression?



Flow chart

95



Model for all N observations together

$$\mathbf{y} = \mathbf{H} \mathbf{w} + \boldsymbol{\epsilon}$$

Gradient of ridge regression cost

97

$$\begin{aligned}\nabla [\text{RSS}(\mathbf{w}) + \lambda \|\mathbf{w}\|_2^2] &= \nabla [(\mathbf{y} - \mathbf{H}\mathbf{w})^\top (\mathbf{y} - \mathbf{H}\mathbf{w}) + \lambda \mathbf{w}^\top \mathbf{w}] \\ &= \underbrace{\nabla [(\mathbf{y} - \mathbf{H}\mathbf{w})^\top (\mathbf{y} - \mathbf{H}\mathbf{w})]}_{-2\mathbf{H}^\top (\mathbf{y} - \mathbf{H}\mathbf{w})} + \lambda \underbrace{\nabla [\mathbf{w}^\top \mathbf{w}]}_{2\mathbf{w}}\end{aligned}$$


Why? By analogy to 1d case...

$\mathbf{w}^\top \mathbf{w}$ analogous to w^2 and derivative of $w^2 = 2w$

Ridge regression: closed-form solution

98

3D plot of RSS with tangent plane at minimum



$$\nabla \text{cost}(\mathbf{w}) = -2\mathbf{H}^T(\mathbf{y} - \mathbf{H}\mathbf{w}) + 2\lambda\mathbf{I}\mathbf{w} = 0$$

$$\text{Solve for } \mathbf{w}: -\mathbf{H}^T\mathbf{y} + \mathbf{H}^T\mathbf{H}\hat{\mathbf{w}} + \lambda\mathbf{I}\hat{\mathbf{w}} = 0$$

$$\mathbf{H}^T\mathbf{H}\hat{\mathbf{w}} + \lambda\mathbf{I}\hat{\mathbf{w}} = \mathbf{H}^T\mathbf{y}$$


$$(\mathbf{H}^T\mathbf{H} + \lambda\mathbf{I})\hat{\mathbf{w}} = \mathbf{H}^T\mathbf{y}$$

$$\hat{\mathbf{w}}^{\text{ridge}} = (\mathbf{H}^T\mathbf{H} + \lambda\mathbf{I})^{-1}\mathbf{H}^T\mathbf{y}$$

Ridge regression: gradient descent

99

$$\nabla \text{cost}(\mathbf{w}) = -2\mathbf{H}^T(\mathbf{y}-\mathbf{H}\mathbf{w}) + 2\lambda\mathbf{w}$$



Update to j^{th} feature weight:

$$w_j^{(t+1)} \leftarrow w_j^{(t)} - \eta *$$

Same as before (from RSS term) \rightarrow $-2 \sum_{i=1}^N (\mathbf{x}_i)(y_i - \hat{y}_i(\mathbf{w}^{(t)}))$

$$+ 2\lambda w_j^{(t)}$$

new term, comes from the j^{th} component of $2\lambda\mathbf{w}$

Summary of ridge regression algorithm

100

init $\mathbf{w}^{(1)} = 0$ (or randomly, or smartly), $t = 1$

while $\|\nabla \text{RSS}(\mathbf{w}^{(t)})\| > \varepsilon$

for $j = 0, \dots, D$

partial[j] = $-2 \sum_{i=1}^N \mathbf{x}_i (y_i - \hat{y}_i(\mathbf{w}^{(t)}))$

$w_j^{(t+1)} \leftarrow (1 - 2\eta\lambda)w_j^{(t)} - \eta \text{partial}[j]$

$t \leftarrow t + 1$

How to handle the intercept

101

Recall multiple regression model

Model:

$$\begin{aligned} y_i &= w_0 h_0(\mathbf{x}_i) + w_1 h_1(\mathbf{x}_i) + \dots + w_D h_D(\mathbf{x}_i) + \varepsilon_i \\ &= \sum_{j=0}^D w_j h_j(\mathbf{x}_i) + \varepsilon_i \end{aligned}$$

feature 1 = $h_0(\mathbf{x})$...often 1 (constant)

feature 2 = $h_1(\mathbf{x})$... e.g., $\mathbf{x}[1]$

feature 3 = $h_2(\mathbf{x})$... e.g., $\mathbf{x}[2]$

...

feature D+1 = $h_D(\mathbf{x})$... e.g., $\mathbf{x}[d]$

Do we penalize intercept?

102

Standard ridge regression cost:

$$\text{RSS}(\mathbf{w}) + \lambda \|\mathbf{w}\|_2^2$$

 strength of penalty

Encourages intercept w_0 to also be small

Do we want a small intercept?

Conceptually, not indicative of overfitting...

Do we penalize intercept?

103

□ **Option 1: don't penalize intercept**

Modified ridge regression cost:

$$\text{RSS}(\mathbf{w}_0, \mathbf{w}_{\text{rest}}) + \lambda \|\mathbf{w}_{\text{rest}}\|_2^2$$

□ **Option 2: Center data first**

If data are first **centered about 0**, then favoring small intercept not so worrisome

Step 1: Transform y to have 0 mean

Step 2: Run ridge regression as normal
(closed-form or gradient algorithms)

FEATURES SELECTION & LASSO REGRESSION

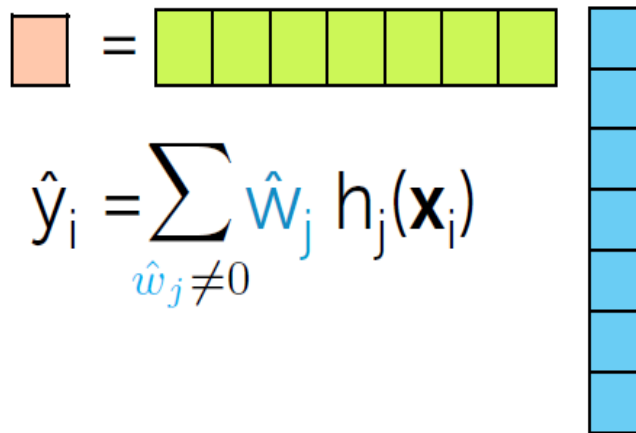
Why features selection?

105

Efficiency:

- If $\text{size}(\mathbf{w}) = 100\text{B}$, each prediction is expensive
- If $\hat{\mathbf{w}}$ **sparse**, computation only depends on # of non-zeros

 many zeros


$$\hat{y}_i = \sum_{\hat{w}_j \neq 0} \hat{w}_j h_j(\mathbf{x}_i)$$

Interpretability:

- Which features are relevant for prediction?


Housing application



- Lot size
- Single Family
- Year built
- Last sold price
- Last sale price/sqft
- Finished sqft
- Unfinished sqft
- Finished basement sqft
- # floors
- Flooring types
- Parking type
- Parking amount
- Cooling
- Heating
- Exterior materials
- Roof type
- Structure style
- Dishwasher
- Garbage disposal
- Microwave
- Range / Oven
- Refrigerator
- Washer
- Dryer
- Laundry location
- Heating type
- Jetted Tub
- Deck
- Fenced Yard
- Lawn
- Garden
- Sprinkler System
- ⋮

Find best model of size: 0

107




start with no features:
 $y_i = \epsilon_i$
↑
noise

- # bedrooms
- # bathrooms
- sq.ft. living
- sq.ft. lot
- floors
- year built
- year renovated
- waterfront

Find best model of size: 1

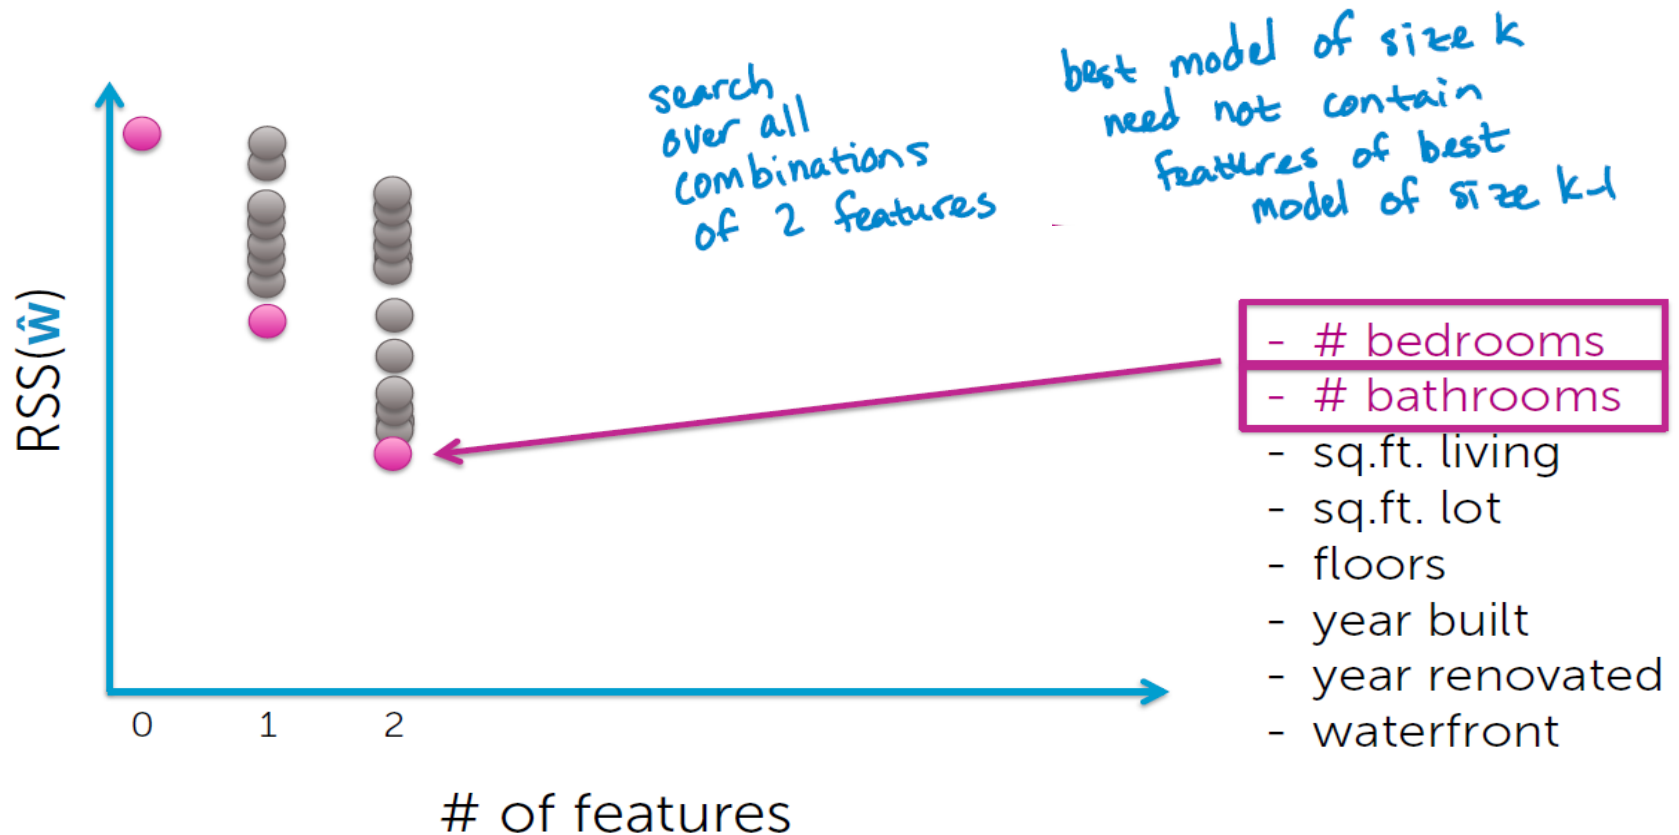
108



Find best model of size: 2


109

Note: not necessarily nested!



Find best model of size: N

110



- # bedrooms
- # bathrooms
- sq.ft. living
- sq.ft. lot
- floors
- year built
- year renovated
- waterfront

Choosing model complexity

111

Option 1: Assess on validation set

Option 2: Cross validation

Option 3+: Other metrics for penalizing model complexity like BIC...

Complexity of „all subsets”

112

How many models were evaluated?

- each indexed by features included

$$y_i = \varepsilon_i$$

$$y_i = w_0 h_0(\mathbf{x}_i) + \varepsilon_i$$

$$y_i = w_1 h_1(\mathbf{x}_i) + \varepsilon_i$$

⋮

$$y_i = w_0 h_0(\mathbf{x}_i) + w_1 h_1(\mathbf{x}_i) + \varepsilon_i$$

⋮

$$y_i = w_0 h_0(\mathbf{x}_i) + w_1 h_1(\mathbf{x}_i) + \dots + w_D h_D(\mathbf{x}_i) + \varepsilon_i$$

feature 0 ← 0 if "no"
1 if "yes"
feature 1 ... feature D

[0 0 0 ... 0 0 0]

[1 0 0 ... 0 0 0]

[0 1 0 ... 0 0 0]

⋮

[1 1 0 ... 0 0 0]

⋮

[1 1 1 ... 1 1 1]

2 2 2 ··· 2

2^{D+1}

$$2^8 = 256$$

$$2^{30} = 1,073,741,824$$

$$2^{1000} = 1.071509 \times 10^{301}$$

$$2^{100B} = \text{HUGE!!!!!!}$$

Typically,
computationally
infeasible

Greedy algorithm


113

Forward stepwise algorithm

1. Pick a dictionary of features $\{h_0(\mathbf{x}), \dots, h_D(\mathbf{x})\}$
 - e.g., polynomials for linear regression
2. Greedy heuristic:
 - i. Start with empty set of features $F_0 = \emptyset$
(or simple set, like just $h_0(\mathbf{x})=1 \rightarrow y_i = w_0 + \epsilon_i$)
 - ii. Fit model using current feature set F_t to get $\hat{\mathbf{w}}^{(t)}$
 - iii. Select next best feature $h_{j^*}(\mathbf{x})$
 - e.g., $h_j(\mathbf{x})$ resulting in lowest training error when learning with $F_t + \{h_j(\mathbf{x})\}$
 - iv. Set $F_{t+1} \leftarrow F_t + \{h_{j^*}(\mathbf{x})\}$
 - v. Recurse

Visualizing greedy algorithm


114



- # bedrooms
- # bathrooms
- sq.ft. living
- sq.ft. lot
- floors
- year built
- year renovated
- waterfront

Visualizing greedy algorithm


115



- # bedrooms
- # bathrooms
- sq.ft. living
- sq.ft. lot
- floors
- year built
- year renovated
- waterfront


Visualizing greedy algorithm

116



Visualizing greedy algorithm

117



Complexity of forward stepwise

118

How many models were evaluated?

- 1st step, D models
- 2nd step, $D-1$ models (add 1 feature out of $D-1$ possible)
- 3rd step, $D-2$ models (add 1 feature out of $D-2$ possible)
- ...

How many steps?

- Depends
- At most D steps (to full model)

$$O(D^2) \ll 2^D$$

for large D

Other greedy algorithms

119

Instead of starting from simple model and always growing...

Backward stepwise:

Start with full model and iteratively remove features least useful to fit

Combining forward and backward steps:

In forward algorithm, insert steps to remove features no longer as important

Lots of other variants, too.

Using regularisation for features selection

120

Instead of searching over a **discrete** set of solutions, can we use **regularization**?

- Start with full model (all possible features)
- “Shrink” some coefficients *exactly to 0*
 - i.e., knock out certain features
- Non-zero coefficients indicate “selected” features

Thresholding ridge coefficients?

121


Why don't we just set small ridge coefficients to 0?



Thresholding ridge coefficients?

122


Selected features for a given threshold value



Thresholding ridge coefficients?

123

Let's look at two related features...




Nothing measuring bathrooms was included!

Thresholding ridge coefficients?

124

If only one of the features had been included...




Remember:

this is linear model. If we assume that #showers = #bathrooms and remove one of them from the model, coefficients will sum up.

Thresholding ridge coefficients?

125

Would have included bathrooms in selected model



Can regularization lead directly to sparsity?

Try this cost instead of ridge ...

126

Total cost =

measure of fit + λ measure of magnitude of coefficients
RSS(\mathbf{w})

$$\|\mathbf{w}\|_1 = |w_0| + \dots + |w_D|$$

Lasso regression
(a.k.a. L_1 regularized regression)


Leads to
sparse
solutions!

Lasso regression

127

Just like ridge regression, solution is governed by a continuous parameter λ

$$\text{RSS}(\mathbf{w}) + \lambda \|\mathbf{w}\|_1$$

 tuning parameter = balance of fit and sparsity


If $\lambda=0$: $\hat{\mathbf{w}}^{\text{lasso}} = \hat{\mathbf{w}}^{\text{LS}}$ (unregularized solution)

If $\lambda=\infty$: $\hat{\mathbf{w}}^{\text{lasso}} = \mathbf{0}$

If λ in between: $\mathbf{0} \leq \|\hat{\mathbf{w}}^{\text{lasso}}\|_1 \leq \|\hat{\mathbf{w}}^{\text{LS}}\|_1$


Coefficient path: ridge

128



Coefficient path: lasso

129







NONPARAMETRIC REGRESSION

Fit globaly vs fit locally


131

Parametric models



Below ...

$f(x)$ is not really a polynomial function



What alternative do we have?

132

If we:

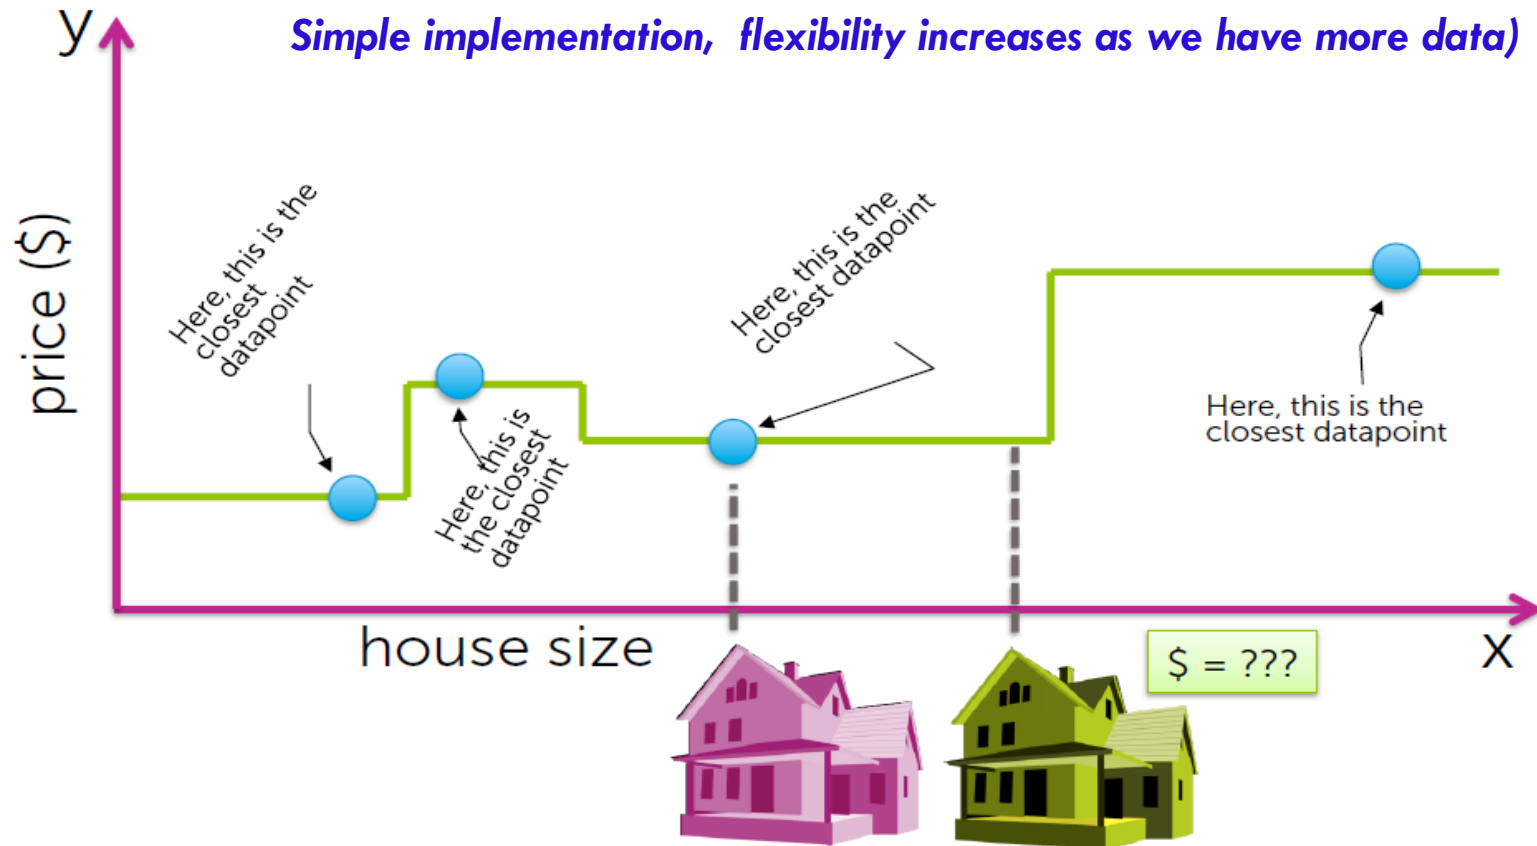
- Want to allow flexibility in $f(\mathbf{x})$ having **local structure**
- Don't want to infer "structural breaks"

What's a simple option we have?

- Assuming we have **plenty of data...**

Nearest Neighbor & Kernel Regression (nonparametric approach)


133



Fit locally to each data point

134

Predicted value = "closest" y_i



What people do naturally...

135

Real estate agent assesses value by finding sale of most similar house



\$ = ???



\$ = 850k

1-NN regression more formally

136

Dataset of (🏠, \$) pairs: $(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_N, y_N)$

Query point: \mathbf{x}_q ← 🏠 \$?
big lime green house

1. Find "closest" \mathbf{x}_i in dataset


$$x_{NN} \leftarrow \min_i \text{distance}(x_i, x_q)$$

big pink house

2. Predict


$$\hat{y}_q = y_{NN}$$

sales price of big pink house



Visualizing 1-NN in multiple dimensions

137



Voronoi tessellation (or diagram):

- Divide space into N regions, each containing 1 datapoint
- Defined such that any x in region is "closest" to region's datapoint

x_j closer to x_i
than any other
 x_j for $j \neq i$.

Don't explicitly form!

Distance metrics: Notion of „closest”

138

In 1D, just Euclidean distance:

$$\text{distance}(x_j, x_q) = |x_j - x_q|$$

In multiple dimensions:

- can define many interesting distance functions
- most straightforwardly, might want to weight different dimensions differently

Weighting housing inputs

139

Some inputs are more relevant than others



bedrooms
bathrooms
sq.ft. living
sq.ft. lot
floors
year built
year renovated
waterfront



Scaled Euclidean distance

140

Formally, this is achieved via

distance($\mathbf{x}_j, \mathbf{x}_q$) =

$$\sqrt{a_1(\mathbf{x}_j[1] - \mathbf{x}_q[1])^2 + \dots + a_d(\mathbf{x}_j[d] - \mathbf{x}_q[d])^2}$$

weight on each input
(defining relative importance)


Other example distance metrics:

- Mahalanobis, rank-based, correlation-based, cosine similarity, Manhattan, Hamming, ...

Different distance metrics

141

Euclidean distance





Manhattan distance



Performing 1-NN search


142

- Query house: 
- Dataset: 
- **Specify:** Distance metric
- **Output:** Most similar house




1-NN algorithm



143

Initialize **Dist2NN** = ∞ ,  = \emptyset ← closest house



For $i=1,2,\dots,N$

 Compute: $\delta = \text{distance}(\img alt="house icon" data-bbox="535 445 590 515"/>_i, \img alt="house icon" data-bbox="620 445 675 515"/>_q)$ ← query house 

 If $\delta < \text{Dist2NN}$




 set  _i

 set **Dist2NN** = δ

Return **most similar house**  ← closest house to query house 

1-NN in practice

144




1-NN sensitive to noise in the data

Get more „comps”

145

More reliable estimate if you base estimate off of a larger set of comparable homes



K-NN regression more formally

146

Dataset of (🏠, \$) pairs: $(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_N, y_N)$

Query point: \mathbf{x}_q

1. Find k closest \mathbf{x}_i in dataset



$(x_{NN1}, x_{NN2}, \dots, x_{NNk})$ such that for any x_i not in nearest neighbor set,
 $\text{distance}(x_i, x_q) \geq \text{distance}(x_{NNk}, x_q)$

2. Predict

$$\begin{aligned}\hat{y}_q &= \frac{1}{k} (y_{NN1} + y_{NN2} + \dots + y_{NNk}) \\ &= \frac{1}{k} \sum_{j=1}^k y_{NNj}\end{aligned}$$

K-NN more formally

147

- Query house: 
- Dataset: 
- **Specify:** Distance metric
- **Output:** Most similar houses



K-NN algorithm

148

Initialize **Dist2kNN** = $\text{sort}(\delta_1, \dots, \delta_k)$ ← list of sorted distances
= $\text{sort}(\text{house}_1, \dots, \text{house}_k)$ ← list of sorted houses

For $i = k+1, \dots, N$
Compute: $\delta = \text{distance}(\text{house}_i, \text{house}_q)$ ← query house

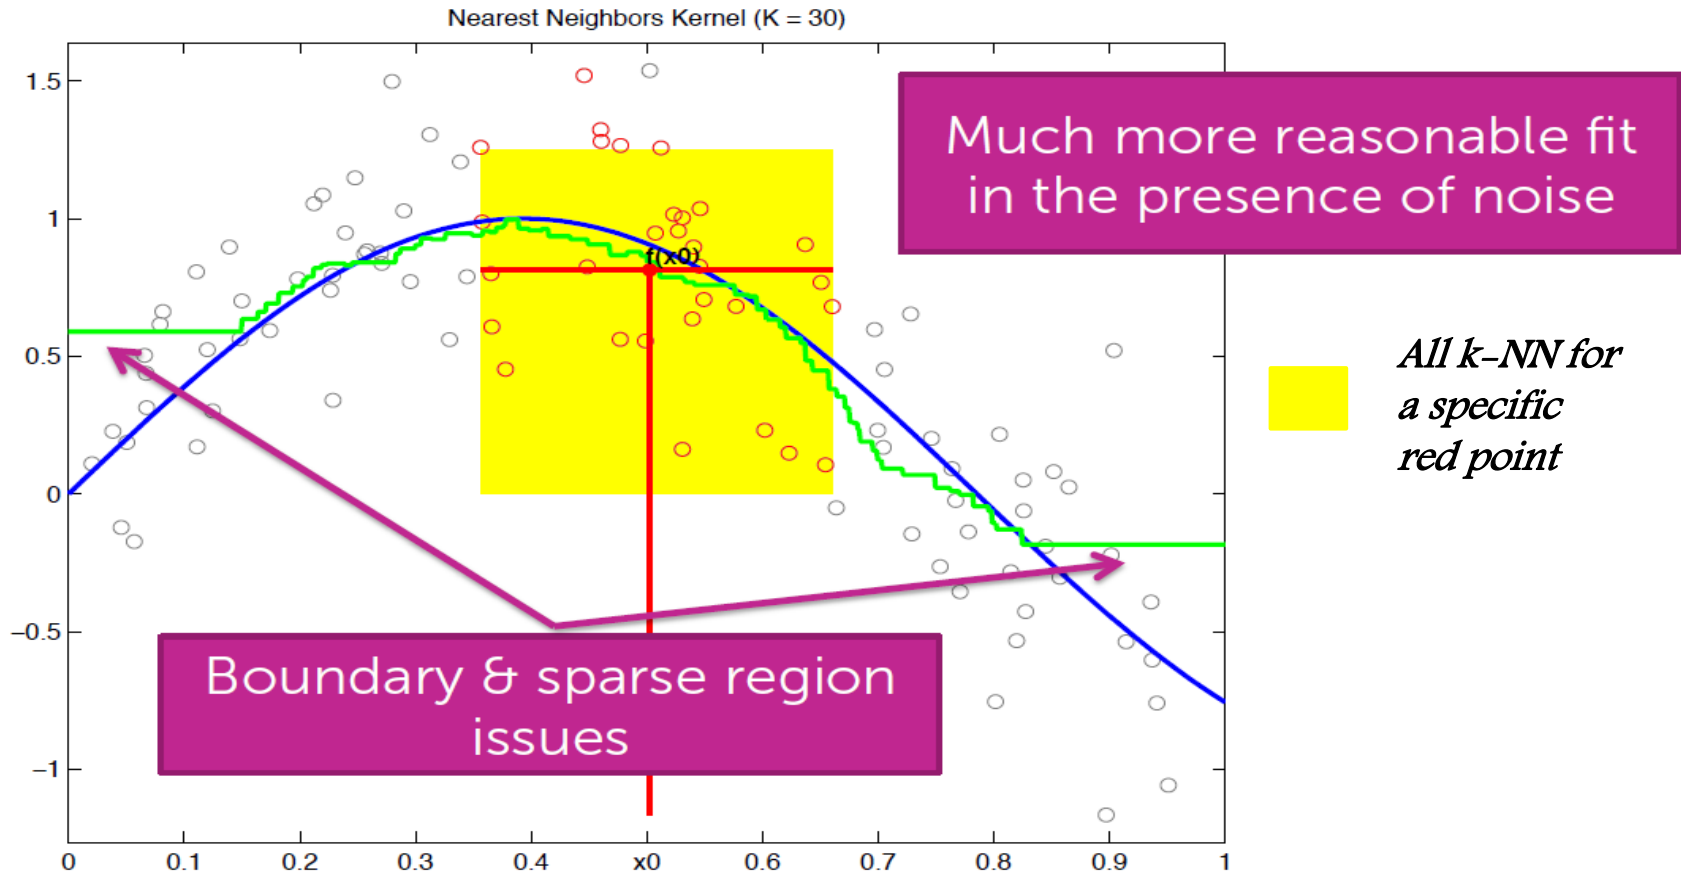
If $\delta < \text{Dist2kNN}[k]$
find j such that $\delta > \text{Dist2kNN}[j-1]$ but $\delta < \text{Dist2kNN}[j]$
remove furthest house and shift queue:
 $[j+1:k] = [j:k]$
Dist2kNN[$j+1:k$] = **Dist2kNN**[$j:k-1$]

insert new NN
set **Dist2kNN**[j] = δ and $\text{house}_j = \text{house}_i$

Return **k most similar houses** ← closest houses to query house


K-NN in practice

149



K-NN in practice

150



Issues with discontinuities

151

Overall predictive accuracy might be okay, but...

For example, in housing application:

- If you are a buyer or seller, this matters
- Can be a jump in estimated value of house going just from 2640 sq.ft. to 2641 sq.ft.
- Don't really believe this type of fit

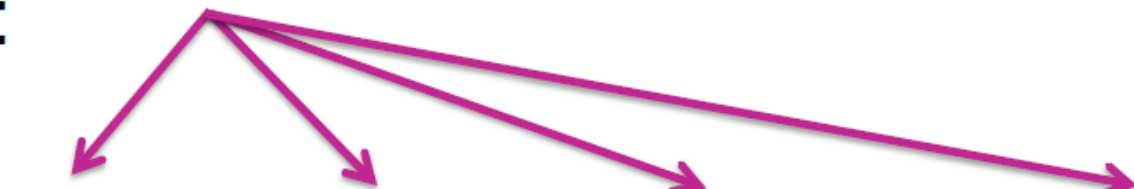
Weighted k-NN

152

Weigh more similar houses more than those less similar in list of k-NN

Predict:

weights on NN


$$\hat{y}_q = \frac{C_{qNN1}Y_{NN1} + C_{qNN2}Y_{NN2} + C_{qNN3}Y_{NN3} + \dots + C_{qNNk}Y_{NNk}}{\sum_{j=1}^k C_{qNNj}}$$

How to define weights

153

Want weight c_{qNNj} to be small when
distance($\mathbf{x}_{NNj}, \mathbf{x}_q$) large

and c_{qNNj} to be large when
distance($\mathbf{x}_{NNj}, \mathbf{x}_q$) small

Simple method:


$$c_{qNNj} = \frac{1}{\text{distance}(x_j, x_q)}$$

Kernel weights for $d=1$

154

Define: $c_{qNNj} = \text{Kernel}_\lambda(|x_{NNj} - x_q|)$

simple isotropic case



Gaussian kernel:

$$\text{Kernel}_\lambda(|x_i - x_q|) = \exp(-(x_i - x_q)^2 / \lambda)$$

Note: never exactly 0!

Kernel drives how the weights will decay, if at all, as a function of the distance.

Kernel regression

Nadaraya-Watson
kernel weighted average

155

Instead of just weighting NN, weight *all* points


Predict:

weight on each datapoint

$$\hat{y}_q = \frac{\sum_{i=1}^N c_{qi} y_i}{\sum_{i=1}^N c_{qi}} = \frac{\sum_{i=1}^N \text{Kernel}_\lambda(\text{distance}(\mathbf{x}_i, \mathbf{x}_q)) * y_i}{\sum_{i=1}^N \text{Kernel}_\lambda(\text{distance}(\mathbf{x}_i, \mathbf{x}_q))}$$

Kernel regression in practice


156



Choice of bandwidth λ

157

Often, choice of kernel matters much less than choice of λ



Choosing λ (or k on k -NN)

158

How to choose? Same story as always...

Cross Validation


Contrasting with global average

159

A globally constant fit weights all points equally

$$\hat{y}_q = \frac{1}{N} \sum_{i=1}^N y_i = \frac{\sum_{i=1}^N c y_i}{\sum_{i=1}^N c}$$

equal weight on each datapoint




Contrasting with global average

160

Kernel regression leads to **locally constant fit**

- slowly add in some points and
and let others gradually die off

$$\hat{y}_q = \frac{\sum_{i=1}^N \text{Kernel}_\lambda(\text{distance}(\mathbf{x}_i, \mathbf{x}_q)) * y_i}{\sum_{i=1}^N \text{Kernel}_\lambda(\text{distance}(\mathbf{x}_i, \mathbf{x}_q))}$$



Local linear regression

161

So far, discussed fitting **constant function locally** at each point

→ “locally weighted averages”

Can instead fit a **line or polynomial locally** at each point

→ “locally weighted linear regression”

Local regression rules of thumb

162

- Local linear fit reduces bias at boundaries with minimum increase in variance
- Local quadratic fit doesn't help at boundaries and increases variance, but does help capture curvature in the interior
- With sufficient data, local polynomials of odd degree dominate those of even degree

Recommended default choice:
local linear regression

Nonparametric approaches

163

k-NN and kernel regression are examples of **nonparametric** regression

General goals of nonparametrics:

- Flexibility
- Make few assumptions about $f(\mathbf{x})$
- **Complexity can grow with the number of observations N**

Lots of other choices:


- Splines, trees, locally weighted structured regression models...

Limiting behaviour of NN

164

Noiseless setting ($\epsilon_i=0$)

In the limit of getting an infinite amount of noiseless data, the **MSE of 1-NN fit goes to 0**




Limiting behaviour of NN


165

Noiseless setting ($\epsilon_i=0$)

In the limit of getting an infinite amount of noiseless data, the **MSE of 1-NN fit goes to 0**



1-NN fit




Quadratic fit

Not true for parametric models!

Error vs amount of data

166




Limiting behaviour of NN

167

Noisy data setting

In the limit of getting an infinite amount of data, the **MSE of NN fit goes to 0** if **k grows, too**



Issues: NN and kernel methods

168

NN and kernel methods work well when the data cover the space, but...

- the more dimensions d you have, the more points N you need to cover the space
- need $N = O(\exp(d))$ data points for good performance

This is where parametric models become useful...


Issues: Complexity of NN search

169

Naïve approach: Brute force search

- Given a query point \mathbf{x}_q
- Scan through each point $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$
- $O(N)$ distance computations per 1-NN query!
- $O(N \log k)$ per k-NN query!

What if N is huge???
(and many queries)



Will talk more about efficient methods in
Clustering & Retrieval course

Summarising

170

Models

- Linear regression
- Regularization: Ridge (L2), Lasso (L1)
- Nearest neighbor and kernel regression

Algorithms

- Gradient descent
- Coordinate descent

Concepts

- Loss functions, bias-variance tradeoff, cross-validation, sparsity, overfitting, model selection, feature selection