

Optimization and (Under/over)fitting

EECS 442 – Prof. David Fouhey

Winter 2019, University of Michigan

http://web.eecs.umich.edu/~fouhey/teaching/EECS442_W19/

Administrivia

- We're grading HW2 and will try to get it back ASAP
- Midterm next week
 - Review sessions next week (go to whichever you want)
 - We will post priorities from slides
 - There's a post on piazza looking for topics you want reviewed

Previously


Regularized Least Squares

Add **regularization** to objective that prefers some solutions:

Before: $\arg \min_{\mathbf{w}} \|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2 \longrightarrow \text{Loss}$

After: $\arg \min_{\mathbf{w}} \|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2 + \lambda \|\mathbf{w}\|_2^2$

Loss Trade-off Regularization



Want model “smaller”: pay a penalty for \mathbf{w} with big norm

Intuitive Objective: accurate model (low loss) but not too complex (low regularization). λ controls how much of each.

Nearest Neighbors

Known Images

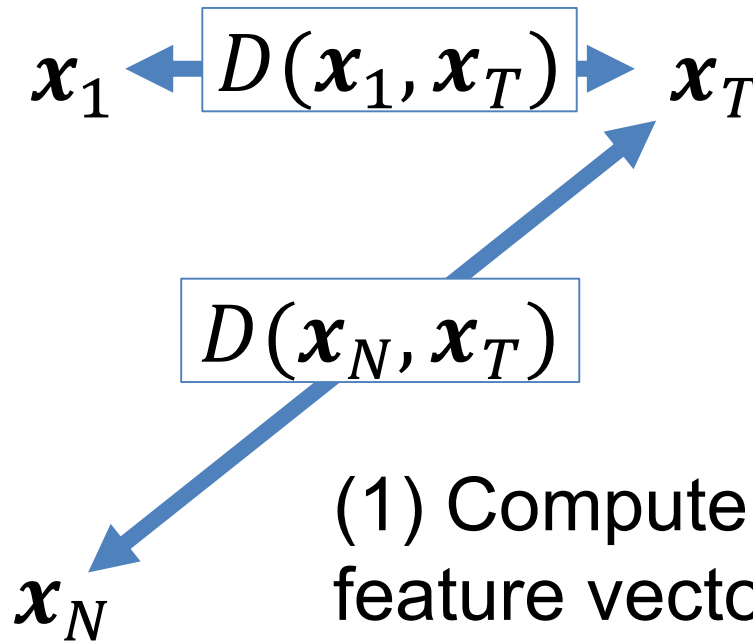
Labels



...



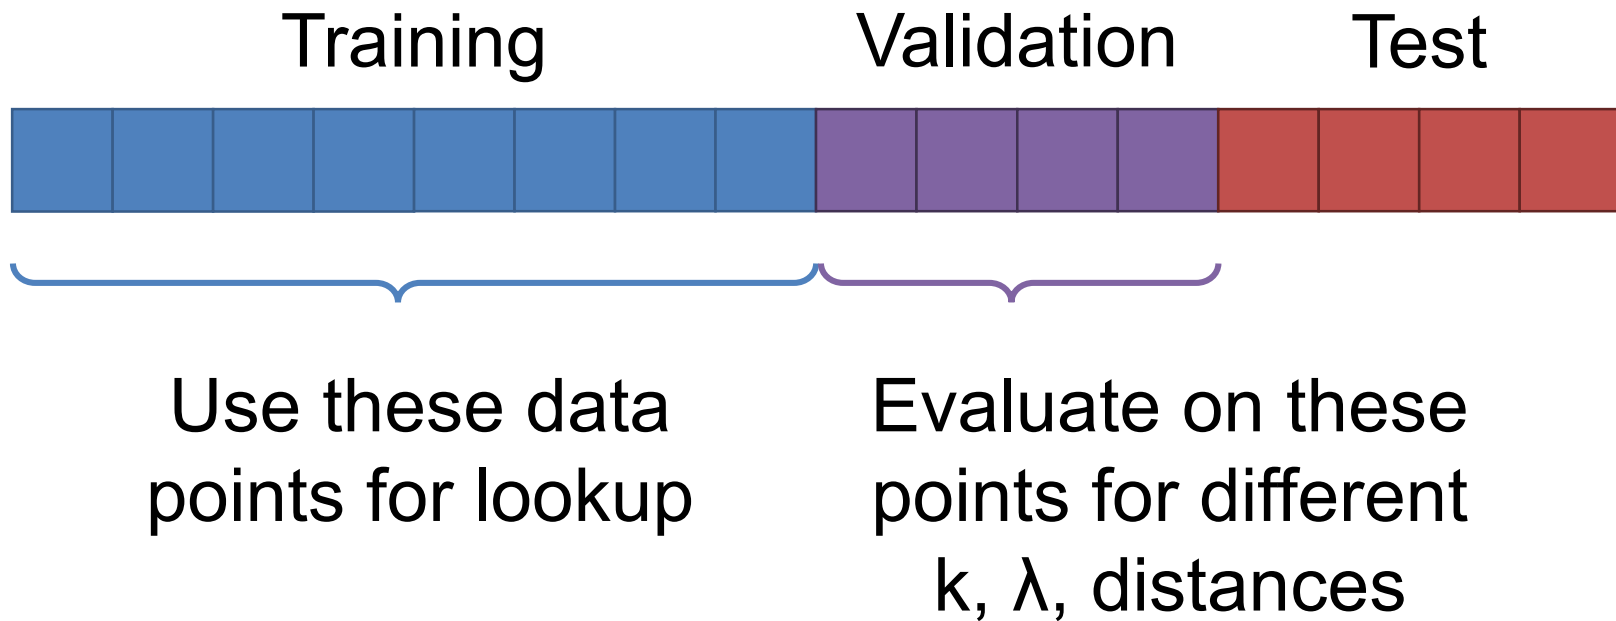
Test Image



- (1) Compute distance between feature vectors
- (2) find nearest
- (3) use label.

Picking Parameters

What distance? What value for k / λ ?



Linear Models

Example Setup: 3 classes



Model – one weight per class: w_0, w_1, w_2



$w_0^T x$ big if cat

Want: $w_1^T x$ big if dog

$w_2^T x$ big if hippo

Stack together: $W_{3 \times F}$ where x is in \mathbb{R}^F

Linear Models



Objective 1: Multiclass SVM

Inference (\mathbf{x}, y) : $\arg \max_k (\mathbf{W}\mathbf{x})_k$

(Take the class whose weight vector gives the highest score)

Training (\mathbf{x}_i, y_i) :

$$\arg \min_W \lambda \|\mathbf{W}\|_2^2 + \sum_{i=1}^n \sum_{j \neq y_i} \max(0, (\mathbf{W}\mathbf{x}_i)_j - (\mathbf{W}\mathbf{x}_i)_{y_i} + m)$$

Regularization

Over all data points


For every class j that's NOT the correct one (y_i)

Pay no penalty if prediction for class y_i is bigger than j by m ("margin"). Otherwise, pay proportional to the score of the wrong class.

Objective 1:

Called: Support Vector Machine

Lots of great theory as to why this is a sensible thing to do. See



Useful book (Free too!):


The Elements of Statistical Learning

Hastie, Tibshirani, Friedman

<https://web.stanford.edu/~hastie/ElemStatLearn/>

Objective 2: Making Probabilities

Converting Scores to “Probability Distribution”



$$\text{Generally } P(\text{class } j): \frac{\exp((Wx)_j)}{\sum_k \exp((Wx)_k)}$$

Called softmax function

Objective 2: Softmax

Inference (x): $\arg \max_k (Wx)_k$ (Take the class whose weight vector gives the highest score)

$$P(\text{class } j): \frac{\exp((Wx)_j)}{\sum_k \exp((Wx)_k)}$$

Why can we skip the exp/sum exp thing to make a decision?

Objective 2: Softmax

Inference (\mathbf{x}): $\arg \max_k (W\mathbf{x})_k$

(Take the class whose weight vector gives the highest score)

Training (\mathbf{x}_i, y_i):

$$\arg \min_W \lambda \|W\|_2^2 + \sum_{i=1}^n -\log \left(\frac{\exp((W\mathbf{x})_{y_i})}{\sum_k \exp((W\mathbf{x})_k)} \right)$$

Regularization

Over all data points

P(correct class)

Pay penalty for not making correct class likely.
“Negative log-likelihood”

Today

- Optimizing things
- Time permitting:
 - (Under/over)fitting or (Bias/Variance)
 - (Not on the exam)

How Do We Optimize Things?

Goal: find the \mathbf{w} minimizing some loss function L .

$$\arg \min_{\mathbf{w} \in \mathbb{R}^N} L(\mathbf{w})$$

Works for lots of different L s:


$$L(\mathbf{W}) = \lambda \|\mathbf{W}\|_2^2 + \sum_{i=1}^n -\log \left(\frac{\exp((W\mathbf{x})_{y_i})}{\sum_k \exp((W\mathbf{x})_k)} \right)$$

$$L(\mathbf{w}) = \lambda \|\mathbf{w}\|_2^2 + \sum_{i=1}^n (y_i - \mathbf{w}^T \mathbf{x}_i)^2$$

$$L(\mathbf{w}) = C \|\mathbf{w}\|_2^2 + \sum_{i=1}^n \max(0, 1 - y_i \mathbf{w}^T \mathbf{x}_i)$$

Sample Function to Optimize

$$f(x,y) = (x+2y-7)^2 + (2x+y-5)^2$$




Sample Function to Optimize

- I'll switch back and forth between this 2D function (called the *Booth Function*) and other more-learning-focused functions
- Beauty of optimization is that it's all the same in principle.
- But don't draw too many conclusions: 2D space has qualitative differences from 1000D space.

See intro of: Dauphin et al. *Identifying and attacking the saddle point problem in high-dimensional non-convex optimization* NIPS 2014

<https://ganguli-gang.stanford.edu/pdf/14.SaddlePoint.NIPS.pdf>



A Caveat



- Each point in the picture is a function evaluation
- Here it takes microseconds – so we can easily see the answer
- Functions we want to optimize may take hours to evaluate

A Few Caveat

Model in your head: moving around a landscape with a teleportation device



Landscape diagram: Karpathy and Fei-Fei

Option #1A – *Grid Search*

```
#systematically try things
best, bestScore = None, Inf
for dim1Value in dim1Values:
    ....
    for dimNValue in dimNValues:
        w = [dim1Value, ..., dimNValue]
        if L(w) < bestScore:
            best, bestScore = w, L(w)
return best
```


Option #1A – “Grid Search”

Pros:

1. Super simple
2. Only requires being able to evaluate model

Cons:

1. Scales horribly to high dimensional spaces

Complexity: $\text{samplesPerDim}^{\text{numberOfDims}}$

Option #1B – Random Search

#Do random stuff RANSAC Style

best, bestScore = None, Inf

for iter in range(numIters):

\mathbf{w} = random(N,1) #sample


 score = $L(\mathbf{w})$ #evaluate

 if score < bestScore:

 best, bestScore = \mathbf{w} , score

return best

Option #1B – Random Search



Option #1B – Random Search

Pros:

1. Super simple
2. Only requires being able to sample model and evaluate it

Cons:

1. Slow and stupid – throwing darts at high dimensional dart board
2. Might miss something

$$P(\text{all correct}) =$$

$$\epsilon^N$$

Good parameters



All parameters



When Do You Use Options 1A/1B?


Use these when

- Number of dimensions small, space bounded
- Objective is impossible to analyze (e.g., test accuracy if we use this distance function)

Random search is arguably more effective; grid search makes it easy to systematically test something (people love certainty)


Option 2 – Use The Gradient

Arrows:
gradient



Option 2 – Use The Gradient

Arrows:
**gradient
direction**
(scaled to unit
length)



Option 2 – Use The Gradient

Want: $\arg \min_{\mathbf{w}} L(\mathbf{w})$

What's the geometric interpretation of:

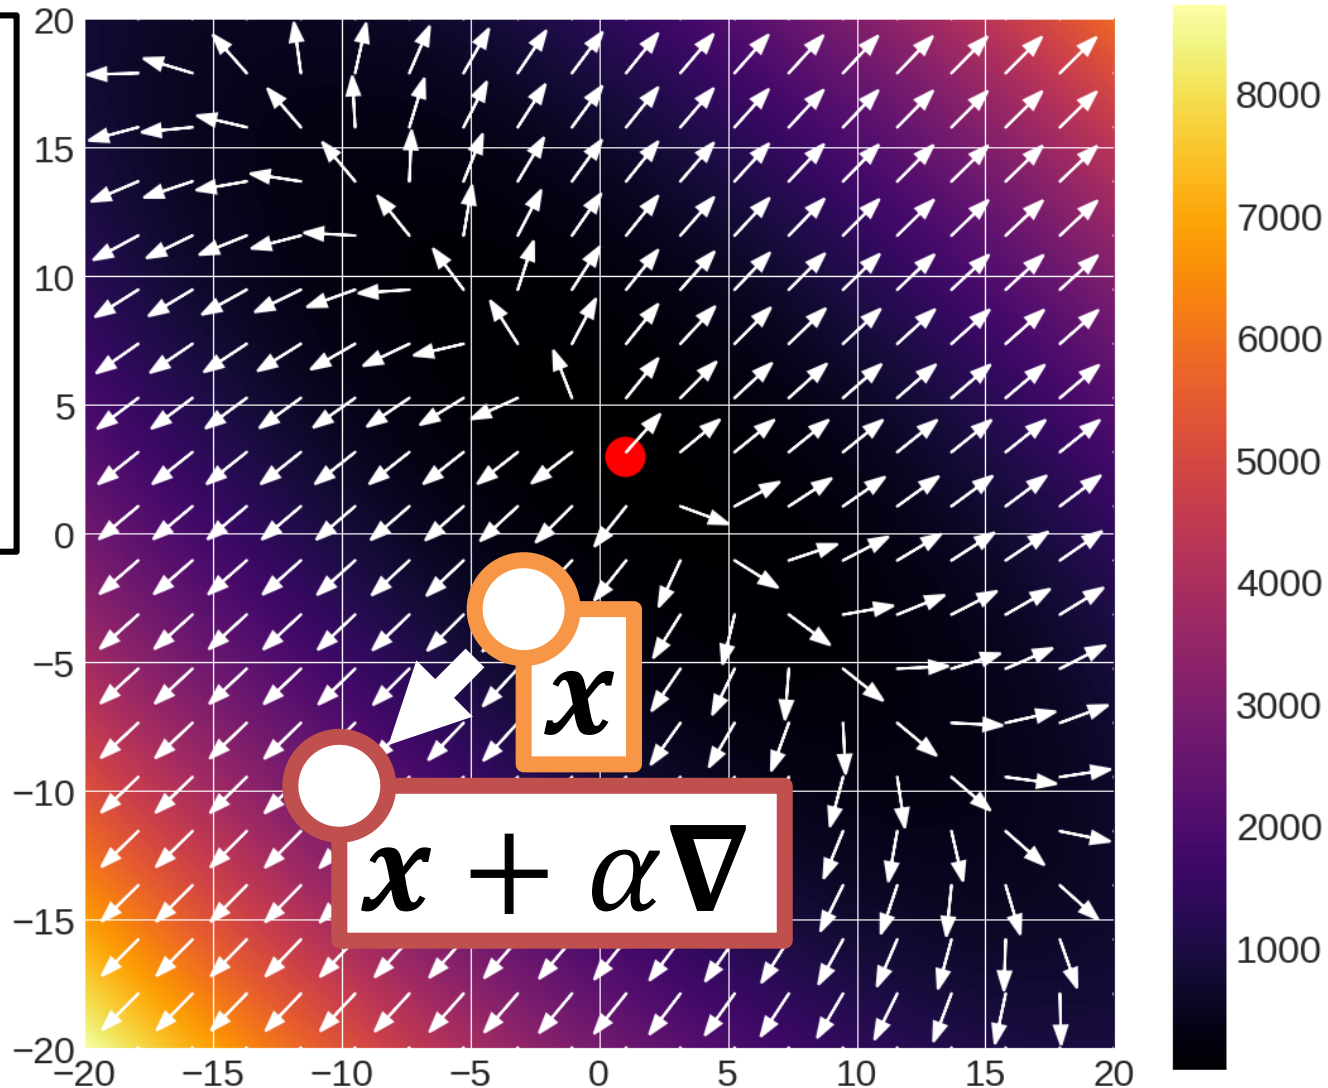
$$\nabla_{\mathbf{w}} L(\mathbf{w}) = \begin{bmatrix} \partial L / \partial x_1 \\ \vdots \\ \partial L / \partial x_N \end{bmatrix}$$

Which is bigger (for small α)?

$$L(\mathbf{x}) \begin{array}{c} \leq? \\ >? \end{array} L(\mathbf{x} + \alpha \nabla_{\mathbf{w}} L(\mathbf{w}))$$

Option 2 – Use The Gradient

Arrows:
gradient
direction
(scaled to unit
length)



Option 2 – Use The Gradient



Method: at each step, move in direction of negative gradient

```
w0 = initialize() #initialize  
for iter in range(numIters):  
    g =  $\nabla_{\mathbf{w}}L(\mathbf{w})$  #eval gradient  
    w = w + -stepsize(iter)*g #update w  
return w
```

Gradient Descent

Given starting point (blue)

$$w_{i+1} = w_i + -9.8 \times 10^{-2} \times \text{gradient}$$



Computing the Gradient

How Do You Compute The Gradient?

Numerical Method:

$$\nabla_{\mathbf{w}}L(\mathbf{w}) = \begin{bmatrix} \frac{\partial L(\mathbf{w})}{\partial x_1} \\ \vdots \\ \frac{\partial L(\mathbf{w})}{\partial x_n} \end{bmatrix}$$

How do you compute this?

$$\frac{\partial f(x)}{\partial x} = \lim_{\epsilon \rightarrow 0} \frac{f(x + \epsilon) - f(x)}{\epsilon}$$

In practice, use:

$$\frac{f(x + \epsilon) - f(x - \epsilon)}{2\epsilon}$$

Computing the Gradient

How Do You Compute The Gradient?

Numerical Method:

$$\nabla_{\mathbf{w}}L(\mathbf{w}) = \begin{bmatrix} \frac{\partial L(\mathbf{w})}{\partial x_1} \\ \vdots \\ \frac{\partial L(\mathbf{w})}{\partial x_n} \end{bmatrix}$$

$$\text{Use: } \frac{f(x + \epsilon) - f(x - \epsilon)}{2\epsilon}$$

How many function evaluations per dimension?

Computing the Gradient

How Do You Compute The Gradient?

Analytical Method:

$$\nabla_{\mathbf{w}}L(\mathbf{w}) = \begin{bmatrix} \frac{\partial L(\mathbf{w})}{\partial x_1} \\ \vdots \\ \frac{\partial L(\mathbf{w})}{\partial x_n} \end{bmatrix}$$

Use calculus!

Computing the Gradient

$$L(\mathbf{w}) = \lambda \|\mathbf{w}\|_2^2 + \sum_{i=1}^n (y_i - \mathbf{w}^T \mathbf{x}_i)^2$$

$$\downarrow \quad \frac{\partial}{\partial \mathbf{w}} \quad \downarrow$$

$$\nabla_{\mathbf{w}} L(\mathbf{w}) = 2\lambda \mathbf{w} + \sum_{i=1}^n -(2(y_i - \mathbf{w}^T \mathbf{x}_i) \mathbf{x}_i)$$

Note: if you look at other derivations, things are written either $(y - \mathbf{w}^T \mathbf{x})$ or $(\mathbf{w}^T \mathbf{x} - y)$; the gradients will differ by a minus.

Interpreting the Gradient

Recall:

$$\mathbf{w} = \mathbf{w} + -\nabla_{\mathbf{w}}L(\mathbf{w}) \text{ \#update } \mathbf{w}$$

$$\nabla_{\mathbf{w}}L(\mathbf{w}) = 2\lambda\mathbf{w} + \sum_{i=1}^n -(2(y_i - \mathbf{w}^T \mathbf{x}_i)\mathbf{x}_i)$$

Push \mathbf{w} towards 0

$$-\nabla_{\mathbf{w}}L(\mathbf{w}) = \underbrace{-2\lambda\mathbf{w}}_{\text{Push } \mathbf{w} \text{ towards } 0} + \sum_{i=1}^n \underbrace{(2(y_i - \mathbf{w}^T \mathbf{x}_i)\mathbf{x}_i)}_{\alpha}$$

If $y_i > \mathbf{w}^T \mathbf{x}_i$ (too low): then $\mathbf{w} = \mathbf{w} + \alpha \mathbf{x}_i$ for some α

Before: $\mathbf{w}^T \mathbf{x}$

After: $(\mathbf{w} + \alpha \mathbf{x})^T \mathbf{x} = \mathbf{w}^T \mathbf{x} + \alpha \mathbf{x}^T \mathbf{x}$

Quick annoying detail: subgradients


What is the derivative of $|x|$?

Derivatives/Gradients

Defined everywhere but 0

$$\frac{\partial}{\partial x} f(x) = \text{sign}(x) \quad x \neq 0$$

undefined $x = 0$



Oh no! A discontinuity!

Let's Compute Another Gradient

Computing The Gradient

Multiclass Support Vector Machine

$$\arg \min_W \lambda \|W\|_2^2 + \sum_{i=1}^n \sum_{j \neq y_i} \max(0, (Wx_i)_j - (Wx_i)_{y_i} + m)$$

Notation:

$W \rightarrow$ rows w_i (i.e., per-class scorer)

$(Wx_i)_j \rightarrow w_j^T x_i$

$$\arg \min_W \lambda \sum_{j=1}^K \|w_j\|_2^2 + \sum_{i=1}^n \sum_{j \neq y_i} \max(0, w_j^T x_i - w_{y_i}^T x_i + m)$$

Computing The Gradient

$$\arg \min_W \lambda \sum_{j=1}^K \|w_j\|_2^2 + \sum_{i=1}^n \underbrace{\sum_{j \neq y_i} \max(0, w_j^T x_i - w_{y_i}^T x_i + m)}$$

$$\frac{\partial}{\partial w_j} : \quad \begin{array}{l} w_j^T x_i - w_{y_i}^T x_i + m \leq 0: \quad 0 \\ w_j^T x_i - w_{y_i}^T x_i + m > 0: \quad x_i \end{array}$$

$$\rightarrow 1(w_j^T x_i - w_{y_i}^T x_i + m > 0)x_i$$

Computing The Gradient

$$\arg \min_W \lambda \sum_{j=1}^K \|w_j\|_2^2 + \sum_{i=1}^n \underbrace{\sum_{j \neq y_i} \max(0, w_j^T x_i - w_{y_i}^T x_i + m)}$$

$$\frac{\partial}{\partial w_{y_i}} : \sum_{j \neq y_i} 1(w_j^T x_i - w_{y_i}^T x_i + m > 0)(-x_i)$$

Interpreting The Gradient

$$-\frac{\partial}{\partial w_j} : \underbrace{-1(w_j^T x_i - w_{y_i}^T x_i + m > 0)}_{\text{Want incorrect class's scoring vector to score that point lower.}} \underbrace{x_i}_{\text{If we do not predict the correct class by at least a score difference of } m \dots}$$

If we do not predict the correct class by at least a score difference of m ...

Want incorrect class's scoring vector to score that point lower.

Recall:

Before: $w^T x$;

After: $(w - \alpha x)^T x = w^T x - \alpha x^T x$

Computing The Gradient

- Numerical: foolproof but slow
- Analytical: can mess things up 😊
- In practice: do analytical, but check with numerical (called a gradient check)

Implementing Gradient Descent

Loss is a function that we can evaluate over data

All Data

$$-\nabla_{\mathbf{w}}L(\mathbf{w}) = -2\lambda\mathbf{w} + \sum_{i=1}^n (2(y_i - \mathbf{w}^T \mathbf{x}_i)\mathbf{x}_i)$$

Subset B

$$-\nabla_{\mathbf{w}}L_B(\mathbf{w}) = -2\lambda\mathbf{w} + \sum_{i \in B} (2(y_i - \mathbf{w}^T \mathbf{x}_i)\mathbf{x}_i)$$

Implementing Gradient Descent

Option 1: Vanilla Gradient Descent

Compute gradient of L over all data points

for iter in range(numIters):

\mathbf{g} = gradient(data, L)

$\mathbf{w} = \mathbf{w} + -\textit{stepsize}(\textit{iter}) * \mathbf{g}$ #update w

Implementing Gradient Descent

Option 2: *Stochastic* Gradient Descent

Compute gradient of L over 1 random sample

```
for iter in range(numIters):
```

```
    index = randint(0,#data)
```

```
    g = gradient(data[index],L)
```

```
    w = w + -stepsize(iter)*g #update w
```


Implementing Gradient Descent

Option 3: *Minibatch* Gradient Descent

Compute gradient of L over subset of B samples

```
for iter in range(numIters):  
    subset = choose_samples(#data,B)  
    g = gradient(data[subset],L)  
    w = w + -stepsize(iter)*g #update w
```


Typical batch sizes: ~ 100

Gradient Descent Details

Step size (also called **learning rate / lr**)
critical parameter


1×10^{-2}

falls short




10×10^{-2}

converges




12×10^{-2}

diverges




Gradient Descent Details

11×10^{-2} :oscillates
(Raw gradients)



Gradient Descent Details

Typical solution: start with initial rate l_r , multiply by f every N iterations




$\text{init_lr} = 10^{-1}$

$f = 0.1$

$N = 10\text{K}$


Gradient Descent Details

11×10^{-2} :oscillates
(Raw gradients)




Solution:
Average gradients

With exponentially decaying weights, called “momentum”




Gradient Descent Details

11×10^{-2} :oscillates
(Raw gradients)




11×10^{-2}
(0.25 momentum)




Gradient Descent Details

Multiple Minima
→
Gradient Descent
Finds **local**
minimum




Gradient Descent Details

Guess the minimum!




Gradient Descent Details



Gradient Descent Details

Guess the minimum!




Gradient Descent Details

Dynamics are fairly complex

Many important functions are **convex**: any local minimum is a global minimum

Many important functions are not.



In practice

- **Conventional wisdom:** minibatch stochastic gradient descent (SGD) + momentum (package implements it for you) + decaying learning rate
- The above is typically what is meant by “SGD”
- Other update rules exist; benefits in general not clear (sometimes better, sometimes worse)

Optimizing Everything

$$L(\mathbf{W}) = \lambda \|\mathbf{W}\|_2^2 + \sum_{i=1}^n -\log \left(\frac{\exp((W\mathbf{x})_{y_i})}{\sum_k \exp((W\mathbf{x})_k)} \right)$$

$$L(\mathbf{w}) = \lambda \|\mathbf{w}\|_2^2 + \sum_{i=1}^n (y_i - \mathbf{w}^T \mathbf{x}_i)^2$$

- Optimize \mathbf{w} on training set with SGD to maximize training accuracy
- Optimize λ with random/grid search to maximize validation accuracy
- Note: Optimizing λ on training sets it to 0

(Over/Under)fitting and Complexity

Let's fit a polynomial: given x , predict y

Note: can do non-linear regression with copies of x


$$\begin{bmatrix} y_1 \\ \vdots \\ y_N \end{bmatrix} = \begin{bmatrix} x_1^F & \cdots & x_1^2 & x_1 & 1 \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ x_N^F & \cdots & x_N^2 & x_N & 1 \end{bmatrix} \begin{bmatrix} w_F \\ \vdots \\ w_2 \\ w_1 \\ w_0 \end{bmatrix}$$

Matrix of all polynomial degrees 

Weights: one per polynomial degree 



(Over/Under)fitting and Complexity

Model: $1.5x^2 + 2.3x + 2 + N(0,0.5)$




Underfitting

Model: $1.5x^2 + 2.3x + 2 + N(0,0.5)$



Underfitting





Model doesn't have the parameters to fit the data.

Bias (statistics): Error intrinsic to the model.



Overfitting

Model: $1.5x^2 + 2.3x + 2 + N(0,0.5)$




Overfitting

Model has high ***variance***: remove **one point**, and model changes dramatically



(Continuous) Model Complexity

$$\arg \min_W \lambda \|W\|_2^2 + \sum_{i=1}^n \underbrace{-\log \left(\frac{\exp((Wx)_{y_i})}{\sum_k \exp((Wx)_k)} \right)}_{\downarrow}$$


Regularization: penalty
for complex model

Pay penalty for negative log-
likelihood of correct class

Intuitively: big weights = more complex model

Model 1: $0.01 \cdot x_1 + 1.3 \cdot x_2 + -0.02 \cdot x_3 + -2.1x_4 + 10$

Model 2: $37.2 \cdot x_1 + 13.4 \cdot x_2 + 5.6 \cdot x_3 + -6.1x_4 + 30$

Fitting Model

Again, fitting polynomial, but with regularization


$$\arg \min_{\mathbf{w}} \|\mathbf{y} - \mathbf{X}\mathbf{w}\| + \lambda \|\mathbf{w}\|$$

The diagram illustrates the components of the regularization equation. A blue arrow points from the matrix \mathbf{X} in the equation to the matrix representation below. Two red arrows point from the weight vector \mathbf{w} in the equation to the vector representation below.


$$\begin{bmatrix} x_1^F & \cdots & x_1^2 & x_1 & 1 \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ x_N^F & \cdots & x_N^2 & x_N & 1 \end{bmatrix} \quad \begin{bmatrix} w_F \\ \vdots \\ w_0 \end{bmatrix}$$

Adding Regularization

No regularization:
fits all data points



Regularization:
can't fit all data points




In General

Error on new data comes from combination of:


1. Bias: model is oversimplified and can't fit the underlying data
2. Variance: you don't have the ability to estimate your model from limited data
3. Inherent: the data is intrinsically difficult

Bias and variance trade-off. Fixing one hurts the other.


Underfitting and Overfitting



Underfitting and Overfitting




Underfitting




Do poorly on both training and validation data due to bias.

Solution:

1. More features
2. More powerful model
3. Reduce regularization




Overfitting



Do well on training data, but poorly on validation data due to variance

Solution:

1. More data
2. Less powerful model
3. Regularize your model more



Cris Dima rule: first make sure you *can* overfit, then stop overfitting.

Next Class

- Non-linear models (neural nets)