## 10-301/601: Introduction

 to Machine Learning Lecture 8 - Optimization for Machine LearningHenry Chai \& Matt Gormley
9/25/23

- Exam 1 on 2/19 (next Monday!) from 7 PM - 9 PM
- Location \& Seats: You all will be split across multiple (large) rooms.
- Everyone will have an assigned seat
- Please watch Piazza carefully for more details
- If you have exam accommodations through ODR, they will be proctoring your exam on our behalf; you are responsible for submitting the exam proctoring request through your student portal.
- Format of questions:
- Multiple choice
- True / False (with justification)
- Derivations
- Short answers
- Drawing \& Interpreting figures
- Implementing algorithms on paper
- No electronic devices (you won't need them!)
- You are allowed to bring one letter-size sheet of notes; you can put whatever you want on both sides
- Covered material: Lectures 1 - 7
- Foundations
- Probability, Linear Algebra, Geometry, Calculus
- Optimization
- Important Concepts
- Overfitting
- Model selection / Hyperparameter optimization


## Exam 1 Topics

- Decision Trees
- $k$-NN
- Perceptron
- Regression
- Decision Tree and $k$-NN Regression
- Linear Regression
- Review the exam practice problems (released 2/12 on the course website, under Coursework)
- Attend the dedicated exam 1 review OH (in lieu of recitation on $2 / 16$ )
- Review HWs 1-3
- Consider whether you have achieved the "learning objectives" for each lecture / section
- Write your one-page cheat sheet (back and front)
- Solve the easy problems first
- If a problem seems extremely complicated, you might be missing something
- If you make an assumption, write it down
- Don't leave any answer blank
- If you look at a question and don't know the answer: - just start trying things
- consider multiple approaches
- imagine arguing for some answer and see if you like it

1. Assume $\mathcal{D}$ generated as:

$$
\begin{aligned}
\mathbf{x}^{(i)} & \sim p^{*}(\cdot) \\
y^{(i)} & =h^{*}\left(\mathbf{x}^{(i)}\right)
\end{aligned}
$$

2. Choose hypothesis space, $\mathcal{H}$ : all linear functions in $M$-dimensional space

$$
\mathcal{H}=\left\{h_{\boldsymbol{\theta}}: h_{\boldsymbol{\theta}}(\mathbf{x})=\boldsymbol{\theta}^{T} \mathbf{x}, \boldsymbol{\theta} \in \mathbb{R}^{M}\right\}
$$

3. Choose an objective function: mean squared error (MSE)

$$
\begin{aligned}
J(\boldsymbol{\theta}) & =\frac{1}{N} \sum_{i=1}^{N} e_{i}^{2} \\
& =\frac{1}{N} \sum_{i=1}^{N}\left(y^{(i)}-h_{\boldsymbol{\theta}}\left(\mathbf{x}^{(i)}\right)\right)^{2} \\
& \left.=\frac{1}{N} \sum_{i=1}^{N}\left(y^{(i)}-\boldsymbol{\theta}^{T} \mathbf{x}^{(i)}\right)\right)^{2}
\end{aligned}
$$

4. Solve the unconstrained optimization problem via favorite method:

- gradient descent
- closed form
- stochastic gradient descent
- ...

$$
\hat{\boldsymbol{\theta}}=\underset{\boldsymbol{\theta}}{\operatorname{argmin}} J(\boldsymbol{\theta})
$$

5. Test time: given a new $\mathbf{x}$, make prediction $\hat{y}$

$$
\hat{y}=h_{\hat{\boldsymbol{\theta}}}(\mathbf{x})=\hat{\boldsymbol{\theta}}^{T} \mathbf{x}
$$

## Linear Regression by Rand. Guessing

## Optimization Method \#o: <br> Random Guessing

1. Pick a random $\boldsymbol{\theta}$
2. Evaluate $J(\boldsymbol{\theta})$
3. Repeat steps 1 and 2 many times
4. Return $\boldsymbol{\theta}$ that gives smallest J( $\boldsymbol{\theta}$ )

| t | $\theta_{1}$ | $\theta_{2}$ | $J\left(\theta_{1}, \theta_{2}\right)$ |
| :---: | :---: | :---: | :---: |
| 1 | 0.2 | 0.2 | 10.4 |
| 2 | 0.3 | 0.7 | 7.2 |
| 3 | 0.6 | 0.4 | 1.0 |
| 4 | 0.9 | 0.7 | 16.2 |

$\left.\mathrm{J}(\boldsymbol{\theta})=\mathrm{J}\left(\theta_{1}, \theta_{2}\right)=\frac{1}{N} \sum_{i=1}^{N}\left(y^{(i)}-\boldsymbol{\theta}^{T} \mathbf{x}^{(i)}\right)\right)^{2}$


## Gradients



## Gradients



These are the gradients that

## Gradients



These are the gradients that Gradient Ascent would follow.

## (Negative) Gradients



These are the negative gradients that

## (Negative)

 Gradients

These are the negative gradients that

## (Negative) Gradient Pa



Shown are the paths that Gradient Descent would follow if it were making infinitesimally

## Recall: <br> Gradient Descent for Linear Regression

- Gradient descent for linear regression repeatedly takes steps opposite the gradient of the objective function

```
```

Algorithm 1 GD for Linear Regression

```
```

Algorithm 1 GD for Linear Regression
procedure $\operatorname{GDLR}\left(\mathcal{D}, \boldsymbol{\theta}^{(0)}\right)$
procedure $\operatorname{GDLR}\left(\mathcal{D}, \boldsymbol{\theta}^{(0)}\right)$
$\boldsymbol{\theta} \leftarrow \boldsymbol{\theta}^{(0)}$
$\boldsymbol{\theta} \leftarrow \boldsymbol{\theta}^{(0)}$
while not converged do
while not converged do
$\mathbf{g} \leftarrow \sum_{i=1}^{N}\left(\boldsymbol{\theta}^{T} \mathbf{x}^{(i)}-y^{(i)}\right) \mathbf{x}^{(i)} \quad \triangleright$ Compute gradient
$\mathbf{g} \leftarrow \sum_{i=1}^{N}\left(\boldsymbol{\theta}^{T} \mathbf{x}^{(i)}-y^{(i)}\right) \mathbf{x}^{(i)} \quad \triangleright$ Compute gradient
$\boldsymbol{\theta} \leftarrow \boldsymbol{\theta}-\gamma \mathbf{g} \quad \triangleright$ Update parameters
$\boldsymbol{\theta} \leftarrow \boldsymbol{\theta}-\gamma \mathbf{g} \quad \triangleright$ Update parameters
return $\theta$

```
```

            return \(\theta\)
    ```
```

$\triangleright$ Initialize parameters
$\triangleright$ Update parameters

## Gradient Calculation for Linear Regression

Derivative of $J^{(i)}(\boldsymbol{\theta})$ :

$$
\begin{aligned}
\frac{d}{d \theta_{k}} J^{(i)}(\boldsymbol{\theta}) & =\frac{d}{d \theta_{k}} \frac{1}{2}\left(\boldsymbol{\theta}^{T} \mathbf{x}^{(i)}-y^{(i)}\right)^{2} \\
& =\frac{1}{2} \frac{d}{d \theta_{k}}\left(\boldsymbol{\theta}^{T} \mathbf{x}^{(i)}-y^{(i)}\right)^{2} \\
& =\left(\boldsymbol{\theta}^{T} \mathbf{x}^{(i)}-y^{(i)}\right) \frac{d}{d \theta_{k}}\left(\boldsymbol{\theta}^{T} \mathbf{x}^{(i)}-y^{(i)}\right) \\
& =\left(\boldsymbol{\theta}^{T} \mathbf{x}^{(i)}-y^{(i)}\right) \frac{d}{d \theta_{k}}\left(\sum_{j=1}^{K} \theta_{j} x_{j}^{(i)}-y^{(i)}\right) \\
& =\left(\boldsymbol{\theta}^{T} \mathbf{x}^{(i)}-y^{(i)}\right) x_{k}^{(i)}
\end{aligned}
$$

Derivative of $J(\boldsymbol{\theta})$ :

$$
\begin{aligned}
\frac{d}{d \theta_{k}} J(\boldsymbol{\theta}) & =\sum_{i=1}^{N} \frac{d}{d \theta_{k}} J^{(i)}(\boldsymbol{\theta}) \\
& =\sum_{i=1}^{N}\left(\boldsymbol{\theta}^{T} \mathbf{x}^{(i)}-y^{(i)}\right) x_{k}^{(i)}
\end{aligned}
$$

$$
\text { Gradient of } J(\boldsymbol{\theta}) \quad \text { [used by Gradient Descent] }
$$

$$
\nabla_{\boldsymbol{\theta}} J(\boldsymbol{\theta})=\left[\begin{array}{c}
\frac{d}{d \theta_{1}} J(\boldsymbol{\theta}) \\
\frac{\theta_{1}}{d \theta_{2}} J(\boldsymbol{\theta}) \\
\vdots \\
\frac{d}{d \theta_{M}} J(\boldsymbol{\theta})
\end{array}\right]=\left[\begin{array}{c}
\sum_{i=1}^{N}\left(\boldsymbol{\theta}^{T} \mathbf{x}^{(i)}-y^{(i)}\right) x_{1}^{(i)} \\
\sum_{i=1}^{N}\left(\boldsymbol{\theta}^{T} \mathbf{x}^{(i)}-y^{(i)}\right) x_{2}^{(i)} \\
\vdots \\
\sum_{i=1}^{N}\left(\boldsymbol{\theta}^{T} \mathbf{x}^{(i)}-y^{(i)}\right) x_{M}^{(i)}
\end{array}\right]
$$

$$
=\sum_{i=1}^{N}\left(\boldsymbol{\theta}^{T} \mathbf{x}^{(i)}-y^{(i)}\right) \mathbf{x}^{(i)}
$$

## Linear Regression by Gradient Desc.

## Optimization Method \#1:

## Gradient Descent

1. Pick a random $\boldsymbol{\theta}$
2. Repeat:
a. Evaluate gradient $\nabla \mathrm{J}(\boldsymbol{\theta})$
b. Step opposite gradient
3. Return $\boldsymbol{\theta}$ that gives smallest J( $\boldsymbol{\theta}$ )
$\left.J(\boldsymbol{\theta})=J\left(\theta_{1}, \theta_{2}\right)=\frac{1}{N} \sum_{i=1}^{N}\left(y^{(i)}-\boldsymbol{\theta}^{T} \mathbf{x}^{(i)}\right)\right)^{2}$


| $t$ | $\theta_{1}$ | $\theta_{2}$ | $J\left(\theta_{1}, \theta_{2}\right)$ |
| :---: | :---: | :---: | :---: |
| 1 | 0.01 | 0.02 | 25.2 |
| 2 | 0.30 | 0.12 | 8.7 |
| 3 | 0.51 | 0.30 | 1.5 |
| 4 | 0.59 | 0.43 | 0.2 |

## Linear Regression by Gradient Desc.

## Optimization Method \#1:

## Gradient Descent

1. Pick a random $\boldsymbol{\theta}$
2. Repeat:
a. Evaluate gradient $\nabla \mathrm{J}(\boldsymbol{\theta})$
b. Step opposite gradient
3. Return $\boldsymbol{\theta}$ that gives
smallest J( $\boldsymbol{\theta}$ )


| t | $\theta_{1}$ | $\theta_{2}$ | $\mathrm{~J}\left(\theta_{1}, \theta_{2}\right)$ |
| :---: | :---: | :---: | :---: |
| 1 | 0.01 | 0.02 | 25.2 |
| 2 | 0.30 | 0.12 | 8.7 |
| 3 | 0.51 | 0.30 | 1.5 |
| 4 | 0.59 | 0.43 | 0.2 |

## Linear Regression by Gradient Desc.

## Optimization Method \#1:

## Gradient Descent

1. Pick a random $\boldsymbol{\theta}$
2. Repeat:
a. Evaluate gradient $\nabla \mathrm{J}(\boldsymbol{\theta})$
b. Step opposite gradient
3. Return $\boldsymbol{\theta}$ that gives smallest J( $\boldsymbol{\theta}$ )



| t | $\theta_{1}$ | $\theta_{2}$ | $\mathrm{~J}\left(\theta_{1}, \theta_{2}\right)$ |
| :---: | :---: | :---: | :---: |
| 1 | 0.01 | 0.02 | 25.2 |
| 2 | 0.30 | 0.12 | 8.7 |
| 3 | 0.51 | 0.30 | 1.5 |
| 4 | 0.59 | 0.43 | 0.2 |

Linear Regression by Gradient Desc.



Linear Regression by Gradient Desc.

$\left.\mathrm{J}(\boldsymbol{\theta})=\mathrm{J}\left(\theta_{1}, \theta_{2}\right)=\frac{1}{N} \sum_{i=1}^{N}\left(y^{(i)}-\boldsymbol{\theta}^{T} \mathbf{x}^{(i)}\right)\right)^{2}$



| $t$ | $\theta_{1}$ | $\theta_{2}$ | $J\left(\theta_{1}, \theta_{2}\right)$ |
| :---: | :---: | :---: | :---: |
| 1 | 0.01 | 0.02 | 25.2 |
| 2 | 0.30 | 0.12 | 8.7 |
| 3 | 0.51 | 0.30 | 1.5 |
| 4 | 0.59 | 0.43 | 0.2 |

$$
J\left(\theta_{1}, \theta_{2}\right)=\frac{1}{N} \sum_{i=1}^{N}\left(y^{(i)}-\boldsymbol{\theta}^{T} \boldsymbol{x}^{(i)}\right)^{2}
$$


iteration $t$



- A function $f: \mathbb{R}^{D} \rightarrow \mathbb{R}$ is convex if

$$
\begin{aligned}
& \forall \boldsymbol{x}^{(1)} \in \mathbb{R}^{D}, \boldsymbol{x}^{(2)} \in \mathbb{R}^{D} \text { and } 0 \leq c \leq 1 \\
& f\left(c \boldsymbol{x}^{(1)}+(1-c) \boldsymbol{x}^{(2)}\right) \leq c f\left(\boldsymbol{x}^{(1)}\right)+(1-c) f\left(\boldsymbol{x}^{(2)}\right)
\end{aligned}
$$

## Convexity



- A function $f: \mathbb{R}^{D} \rightarrow \mathbb{R}$ is convex if

$$
\begin{aligned}
& \forall \boldsymbol{x}^{(1)} \in \mathbb{R}^{D}, \boldsymbol{x}^{(2)} \in \mathbb{R}^{D} \text { and } 0 \leq c \leq 1 \\
& f\left(c \boldsymbol{x}^{(1)}+(1-c) \boldsymbol{x}^{(2)}\right) \leq c f\left(\boldsymbol{x}^{(1)}\right)+(1-c) f\left(\boldsymbol{x}^{(2)}\right)
\end{aligned}
$$

## Convexity



- A function $f: \mathbb{R}^{D} \rightarrow \mathbb{R}$ is strictly convex if

$$
\begin{aligned}
& \forall \boldsymbol{x}^{(1)} \in \mathbb{R}^{D}, \boldsymbol{x}^{(2)} \in \mathbb{R}^{D} \text { and } 0<c<1 \\
& f\left(c \boldsymbol{x}^{(1)}+(1-c) \boldsymbol{x}^{(2)}\right)<c f\left(\boldsymbol{x}^{(1)}\right)+(1-c) f\left(\boldsymbol{x}^{(2)}\right)
\end{aligned}
$$

## Convexity




## Convexity



## Convexity



## Convexity




Non-convex functions:
A local minimum may or may not be a global minimum...

## Convexity



Strictly convex functions:
There exists a unique global minimum!

Non-convex functions:
A local minimum may or may not be a global minimum...

## Gradient Descent \& Convexity

- Gradient descent is a local optimization algorithm - it will converge to a local minimum (if it converges)
- Works great if the objective function is convex!



## Gradient Descent \& Convexity

- Gradient descent is a local optimization algorithm - it will converge to a local minimum (if it converges)
- Works great if the objective function is convex!



## Gradient Descent \& Convexity

- Gradient descent is a local optimization algorithm - it will converge to a local minimum (if it converges)
- Works great if the objective function is convex!



## Gradient Descent \& Convexity

- Gradient descent is a local optimization algorithm - it will converge to a local minimum (if it converges)
- Works great if the objective function is convex!



## Gradient Descent \& Convexity

- Gradient descent is a local optimization algorithm - it will converge to a local minimum (if it converges)
- Not ideal if the objective function is non-convex...



## Gradient Descent \& Convexity

- Gradient descent is a local optimization algorithm - it will converge to a local minimum (if it converges)
- Not ideal if the objective function is non-convex...



## Gradient Descent \& Convexity

- Gradient descent is a local optimization algorithm - it will converge to a local minimum (if it converges)
- Not ideal if the objective function is non-convex...



## Gradient Descent \& Convexity

- Gradient descent is a local optimization algorithm - it will converge to a local minimum (if it converges)
- Not ideal if the objective function is non-convex...


$$
J\left(\theta_{1}, \theta_{2}\right)=\frac{1}{N} \sum_{i=1}^{N}\left(y^{(i)}-\boldsymbol{\theta}^{T} \boldsymbol{x}^{(i)}\right)^{2}
$$


iteration $t$


| $t$ | $\theta_{1}$ | $\theta_{2}$ | $J\left(\theta_{1}, \theta_{2}\right)$ |
| :---: | :---: | :---: | :---: |
| 1 | 0.01 | 0.02 | 25.2 |
| 2 | 0.30 | 0.12 | 8.7 |
| 3 | 0.51 | 0.30 | 1.5 |
| 4 | 0.59 | 0.43 | 0.2 |

## The mean squared error is convex (but not always strictly convex)




## Okay, fine

 but couldn't we do something simpler?$$
J\left(\theta_{1}, \theta_{2}\right)=\frac{1}{N} \sum_{i=1}^{N}\left(y^{(i)}-\boldsymbol{\theta}^{T} \boldsymbol{x}^{(i)}\right)^{2}
$$





- Idea: find the critical points of the objective function, specifically the ones where $\nabla J(\theta)=\mathbf{0}$ (the vector of all zeros), and check if any of them are local minima


## Closed Form Optimization

- Notation: given training data $\mathcal{D}=\left\{\left(\boldsymbol{x}^{(n)}, y^{(n)}\right)\right\}_{n=1}^{N}$
$X=\left[\begin{array}{cc}1 & x^{(1)^{T}} \\ 1 & x^{(2)^{T}} \\ \vdots & \vdots \\ 1 & x^{(N)^{T}}\end{array}\right]=\left[\begin{array}{cccc}1 & x_{1}^{(1)} & \cdots & x_{D}^{(1)} \\ 1 & x_{1}^{(2)} & \cdots & x_{D}^{(2)} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & x_{1}^{(N)} & \cdots & x_{D}^{(N)}\end{array}\right] \in \mathbb{R}^{N \times D+1}$
is the design matrix
- $\boldsymbol{y}=\left[y^{(1)}, \ldots, y^{(N)}\right]^{T} \in \mathbb{R}^{N}$ is the target vector

$$
J(\boldsymbol{\theta})=\frac{1}{N} \sum_{i=1}^{N} \frac{1}{2}\left(y^{(i)}-\boldsymbol{\theta}^{T} \boldsymbol{x}^{(i)}\right)^{2}
$$

Minimizing the
Mean Squared
Error

$$
\widehat{\boldsymbol{\theta}}=\left(X^{T} X\right)^{-1} X^{T} \boldsymbol{y}
$$

## Closed Form

 Optimization


| $t$ | $\theta_{1}$ | $\theta_{2}$ | $J\left(\theta_{1}, \theta_{2}\right)$ |
| :---: | :---: | :---: | :---: |
| 1 | 0.59 | 0.43 | 0.2 |

$$
\widehat{\boldsymbol{\theta}}=\left(X^{T} X\right)^{-1} X^{T} \boldsymbol{y}
$$

1. Is $X^{T} X$ invertible?

## Closed Form Solution

2. If so, how computationally expensive is inverting $X^{T} X$ ?

# Linear <br> Regression: Uniqueness 

- Consider a 1D linear regression model trained to minimize the mean squared error: how many optimal solutions (i.e., sets of parameters $\theta$ ) are there for the given dataset?


# Linear <br> Regression: Uniqueness 

# Linear <br> Regression: Uniqueness 

## Poll Question 3

- Consider a 1D linear regression model trained to minimize the mean squared error: how many optimal solutions (i.e., sets of parameters $\theta$ ) are there for the given dataset?
A. -1 (TOXIC)
B. 0
C. 1
D. 2
E. $\infty$


## Linear <br> Regression: Uniqueness

- Consider a 2D linear regression model trained to minimize the mean squared error: how many optimal solutions (i.e., sets of parameters $\theta$ ) are there for the given dataset?



## Linear <br> Regression: Uniqueness

- Consider a 2D linear regression model trained to minimize the mean squared error: how many optimal solutions (i.e., sets of parameters $\theta$ ) are there for the given dataset?



## Linear <br> Regression: Uniqueness

- Consider a 2D linear regression model trained to minimize the mean squared error: how many optimal solutions (i.e., sets of parameters $\theta$ ) are there for the given dataset?


$$
\widehat{\boldsymbol{\theta}}=\left(X^{T} X\right)^{-1} X^{T} \boldsymbol{y}
$$

1. Is $X^{T} X$ invertible?

## Closed Form Solution

2. If so, how computationally expensive is inverting $X^{T} X$ ?

$$
\widehat{\boldsymbol{\theta}}=\left(X^{T} X\right)^{-1} X^{T} \boldsymbol{y}
$$

1. Is $X^{T} X$ invertible?

- When $N \gg D+1, X^{T} X$ is (almost always) full rank and therefore, invertible!
- If $X^{T} X$ is not invertible (occurs when one of the features is a linear combination of the others), then there are infinitely many solutions

2. If so, how computationally expensive is inverting $X^{T} X$ ?

- $X^{T} X \in \mathbb{R}^{D+1 \times D+1}$ so inverting $X^{T} X$ takes $O\left(D^{3}\right)$ time...
- Computing $X^{T} X$ takes $O\left(N D^{2}\right)$ time
- Can use gradient descent to (potentially) speed things up when $N$ and $D$ are large!

You should be able to...

- Design k-NN Regression and Decision Tree Regression
- Implement learning for Linear Regression using gradient descent or closed form optimization
- Choose a Linear Regression optimization technique that is appropriate for a particular dataset by analyzing the tradeoff of computational complexity vs. convergence speed
- Identify situations where least squares regression has exactly one solution or infinitely many solutions

