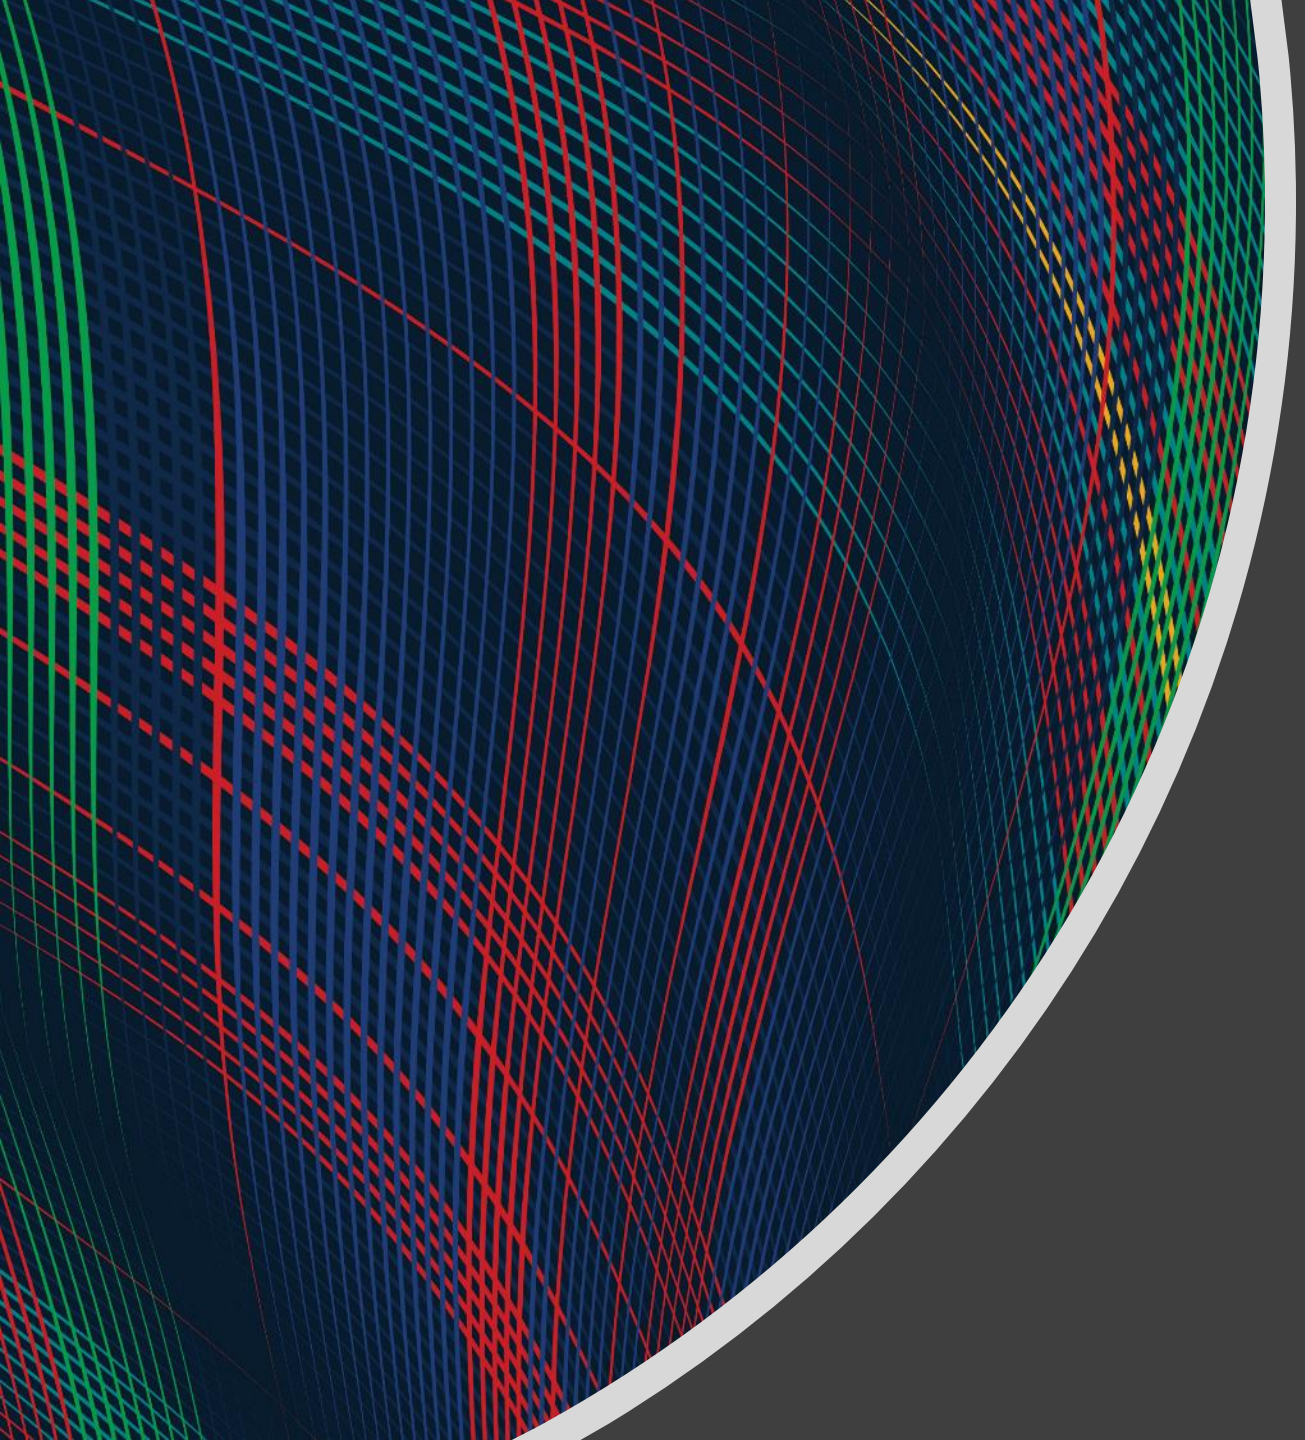

Demystifying AI

Regression

Instructor: Pat Virtue

Regression

Regression: learning a model to predict a numerical output (but not numbers that just represent categories, that would be classification)

Example

Trying to see how much I should sell my car for. Looking up data from car websites, I find the mileage for a set of cars and the selling price for each car.

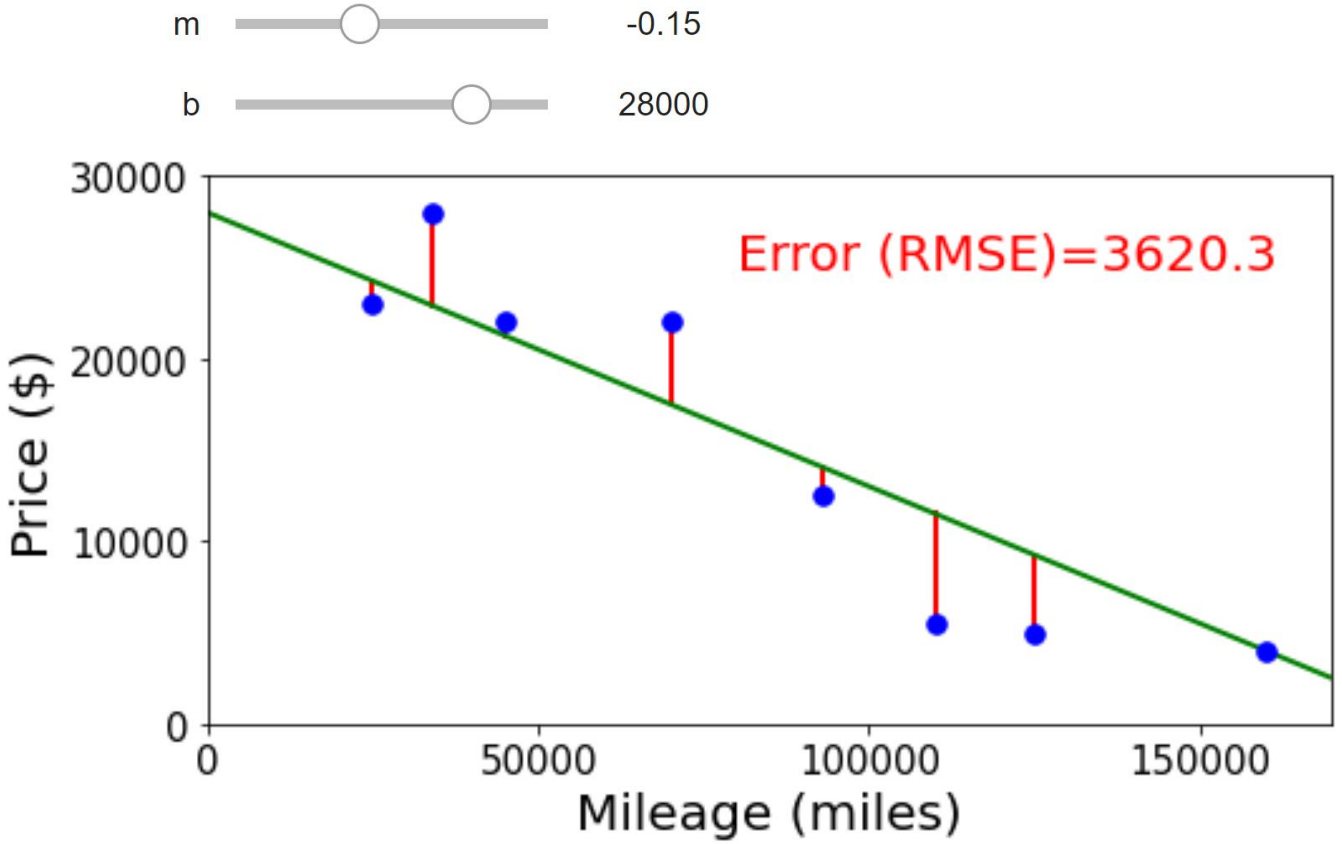

Poll 1

Linear model to predict car price:

Which of the following do need when doing prediction? Assume we have already finished training.

Select ALL that apply.

A. Mileage values from dataset
B. Price values from dataset
C. Final mean squared error from training
D. m parameter (slope)
E. b parameter (intercept)

ML: Linear Regression

Using (training) data to learn a model that we'll later use for prediction

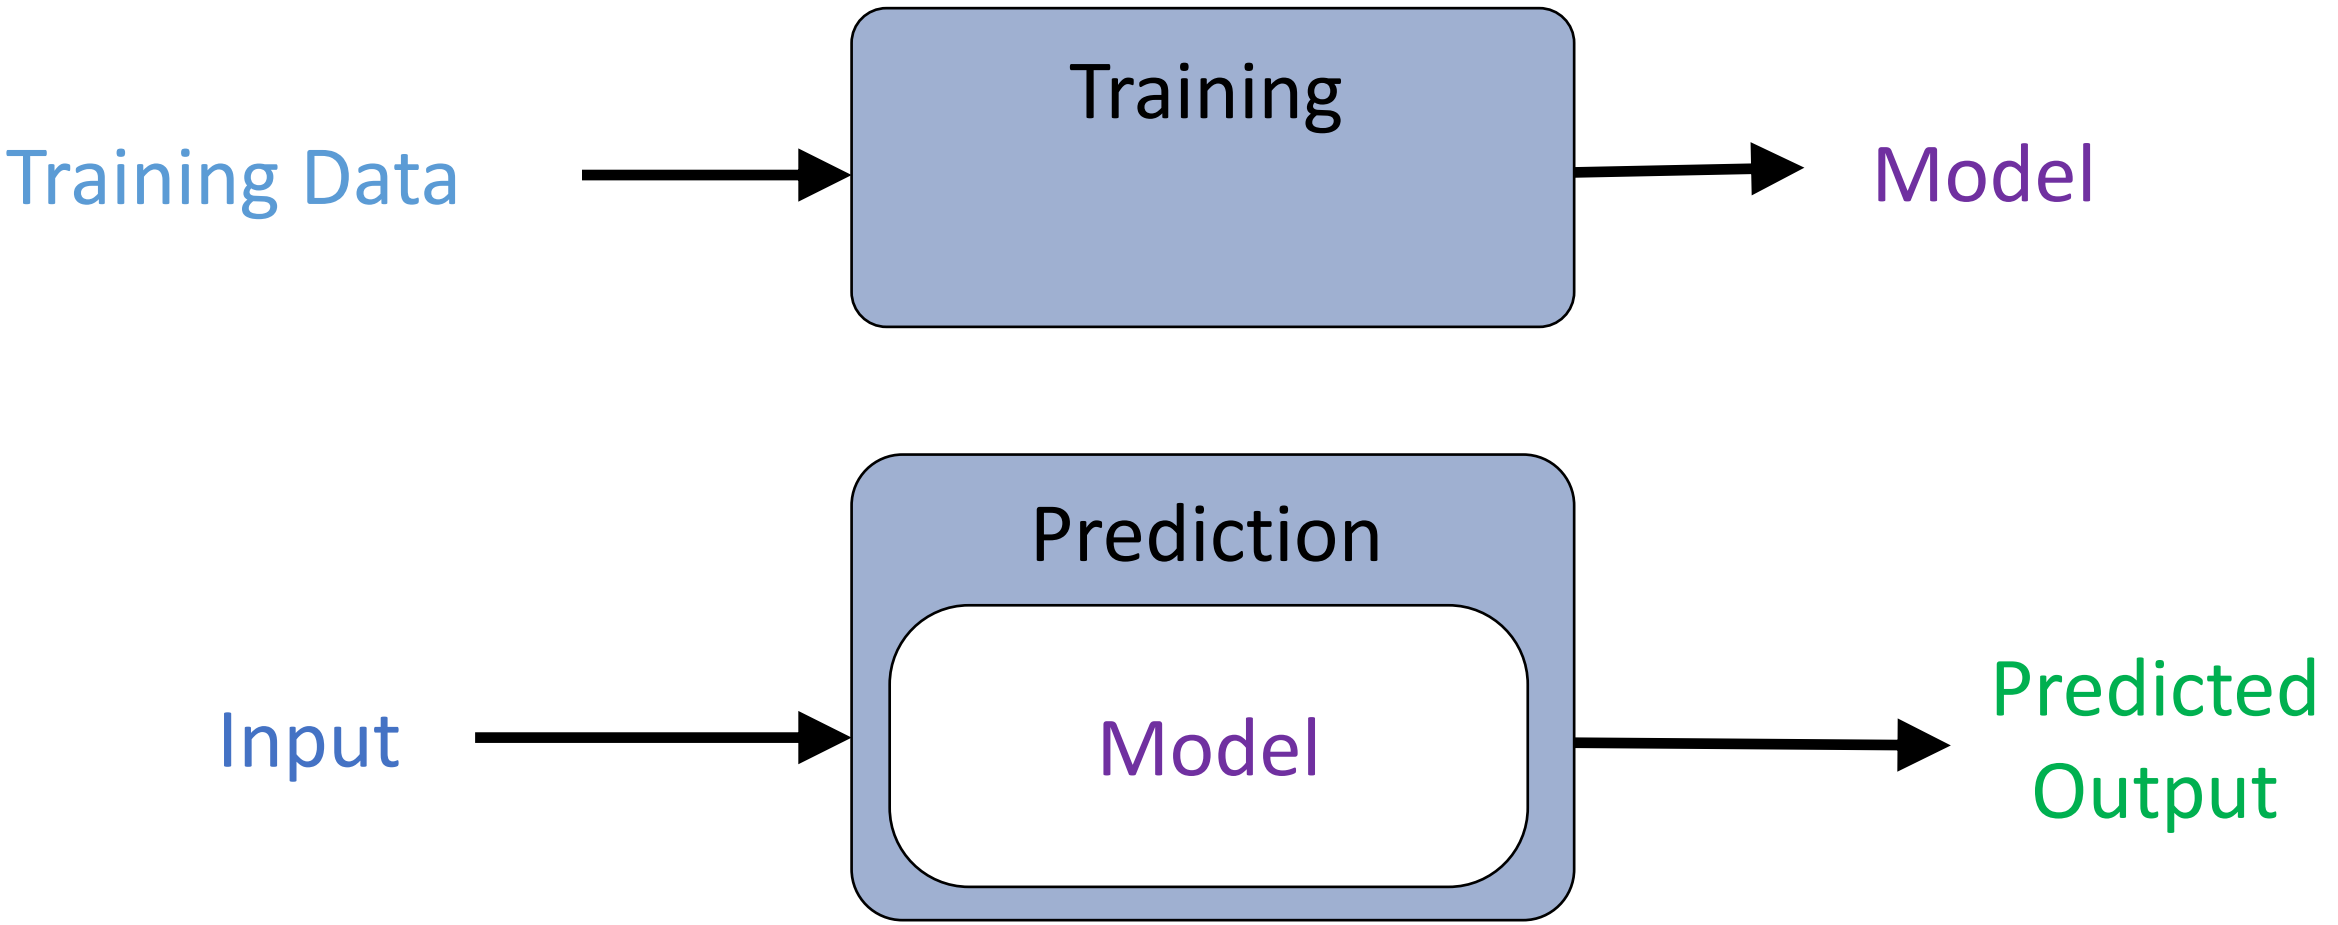

Training: Finding the Best Parameters

Searching for parameters that minimize mean squared error

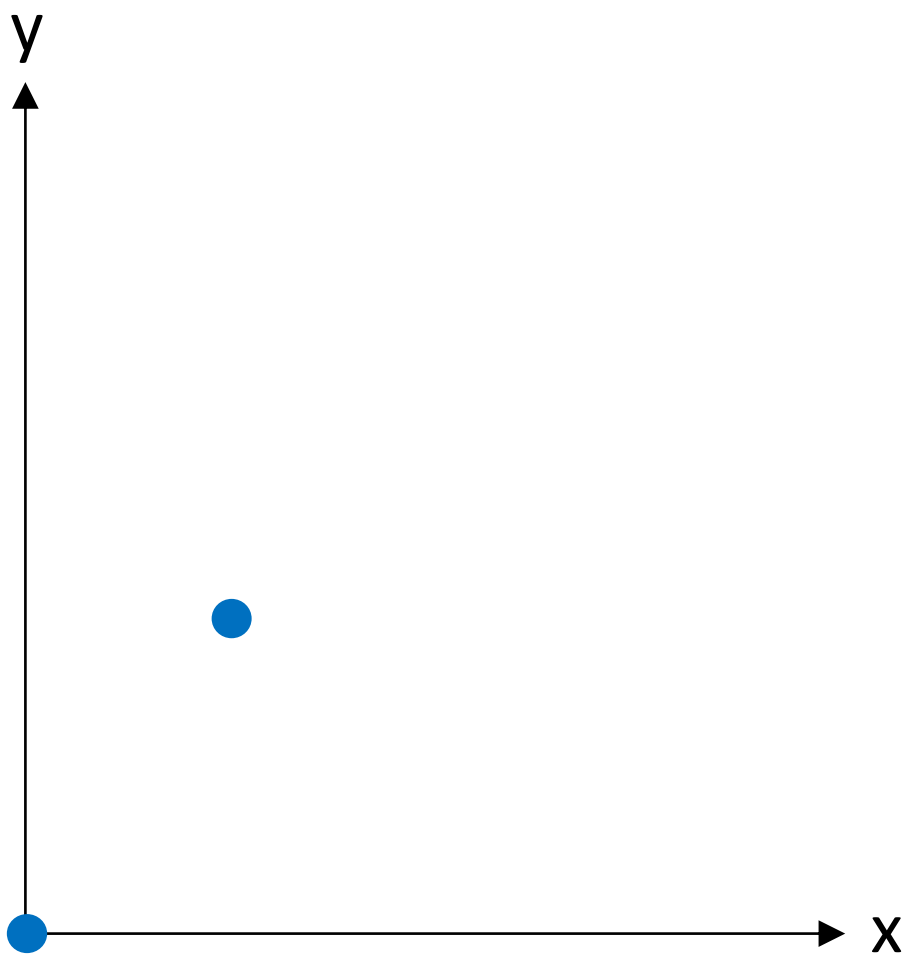

Training a Model

Random search

- Repeatedly choosing random combinations of parameters
- Keep the parameters that have the lowest loss (mean squared error)

Grid search

- Looping over a fixed set of parameter combinations
- Keep the parameters that have the lowest loss (mean squared error)

Implementing Grid Search

Shared Drive Folder / Lec_9_10_Regression / Group N

regression_2_grid_search.ipynb

Visualizing our Parameter Search

Shared Drive Folder / Lec_9_10_Regression / Group N

regression_3_search_visualization.ipynb

Training: Finding the Best Parameters

Optimizing the performance measure

Minimizing mean square error

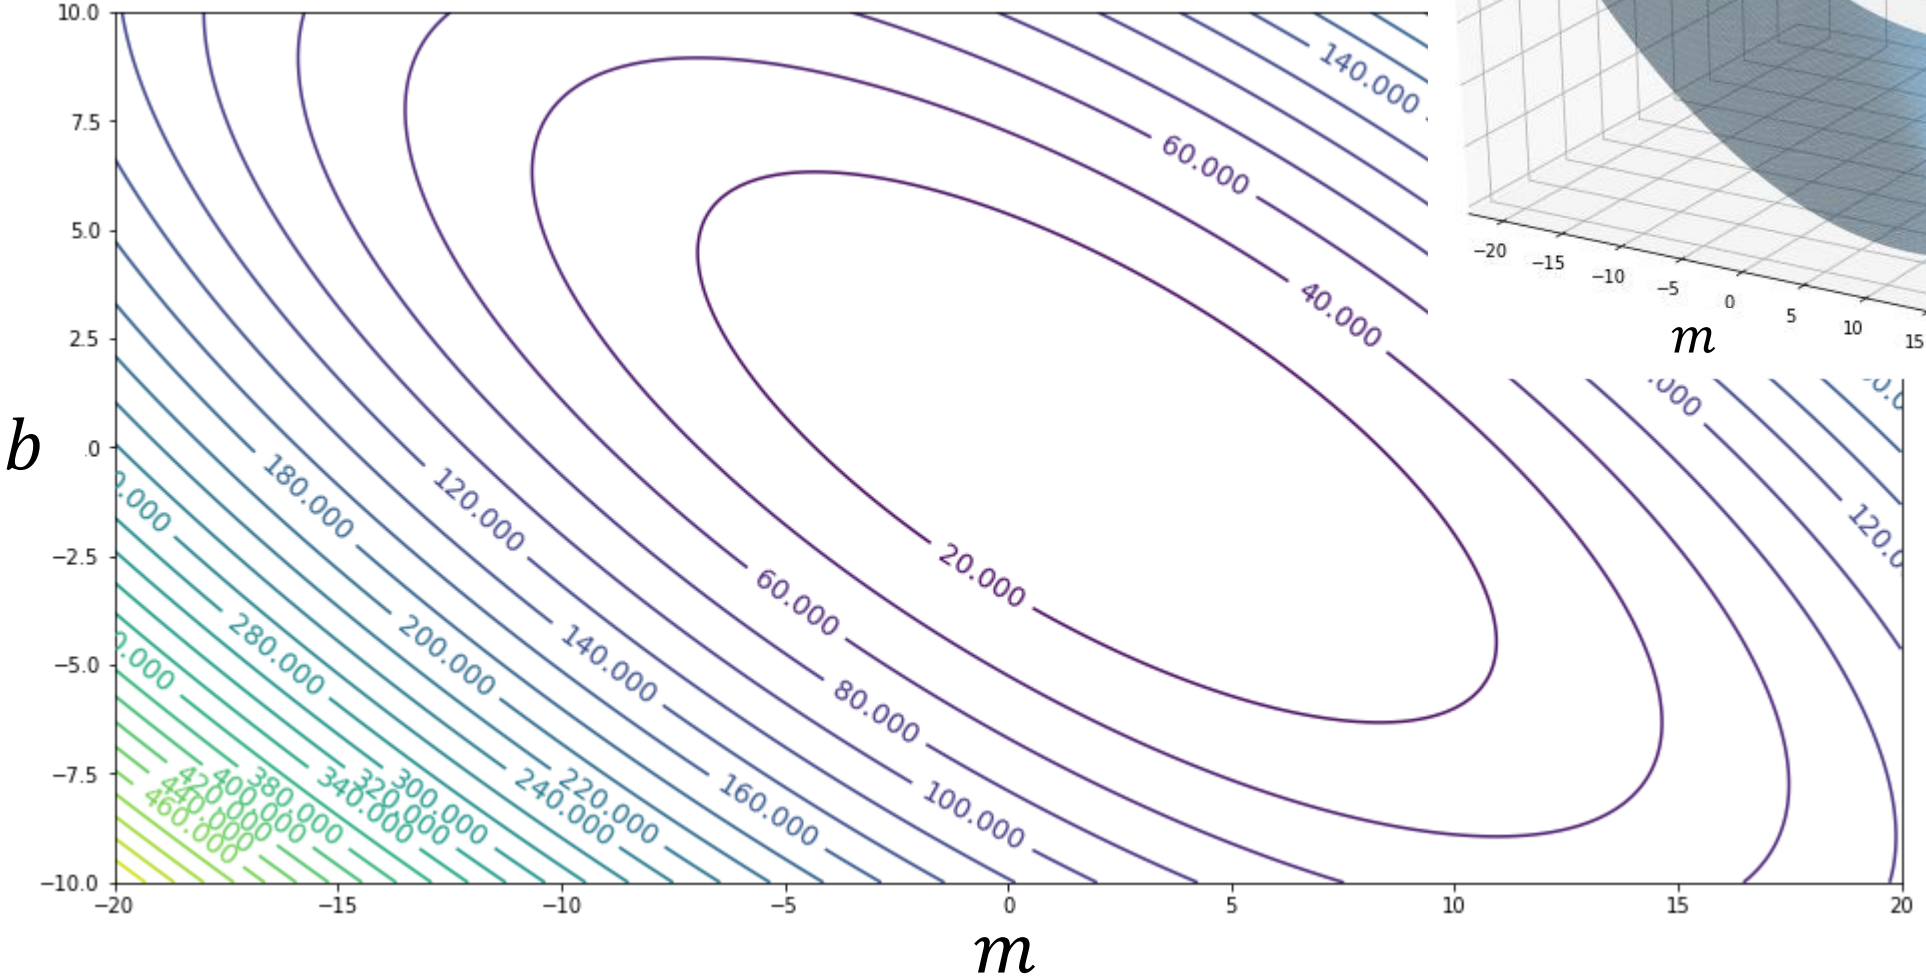

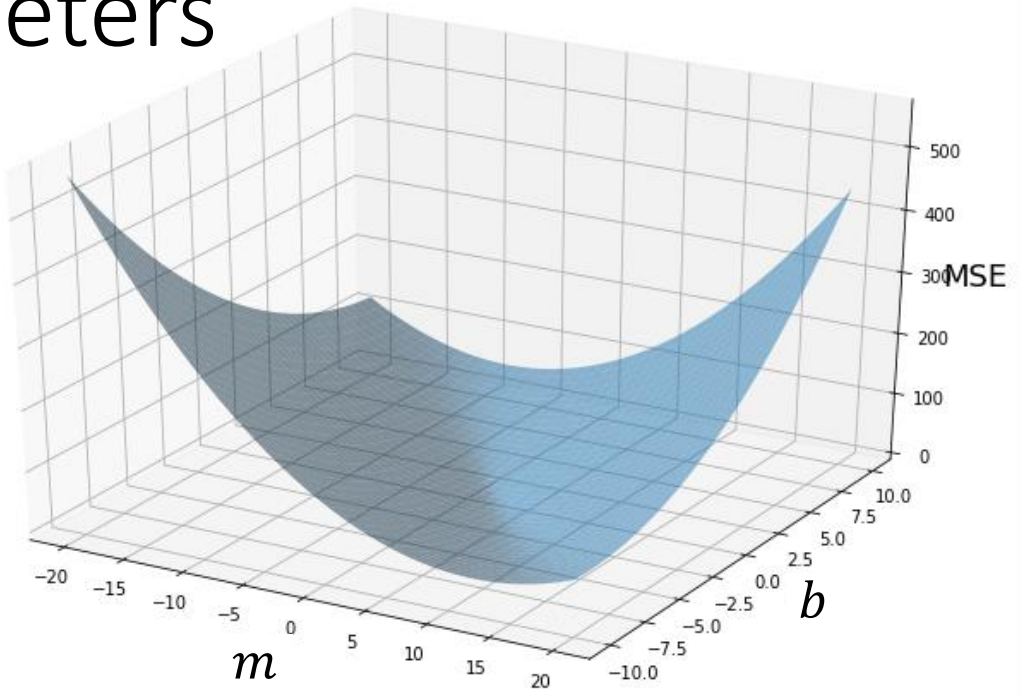

Linear Regression

Methods for optimizing the objective

- Grid search
- Random search
- (Solve for the minimum of this paraboloid)
- Gradient descent

Performance measure:
Mean Squared Error

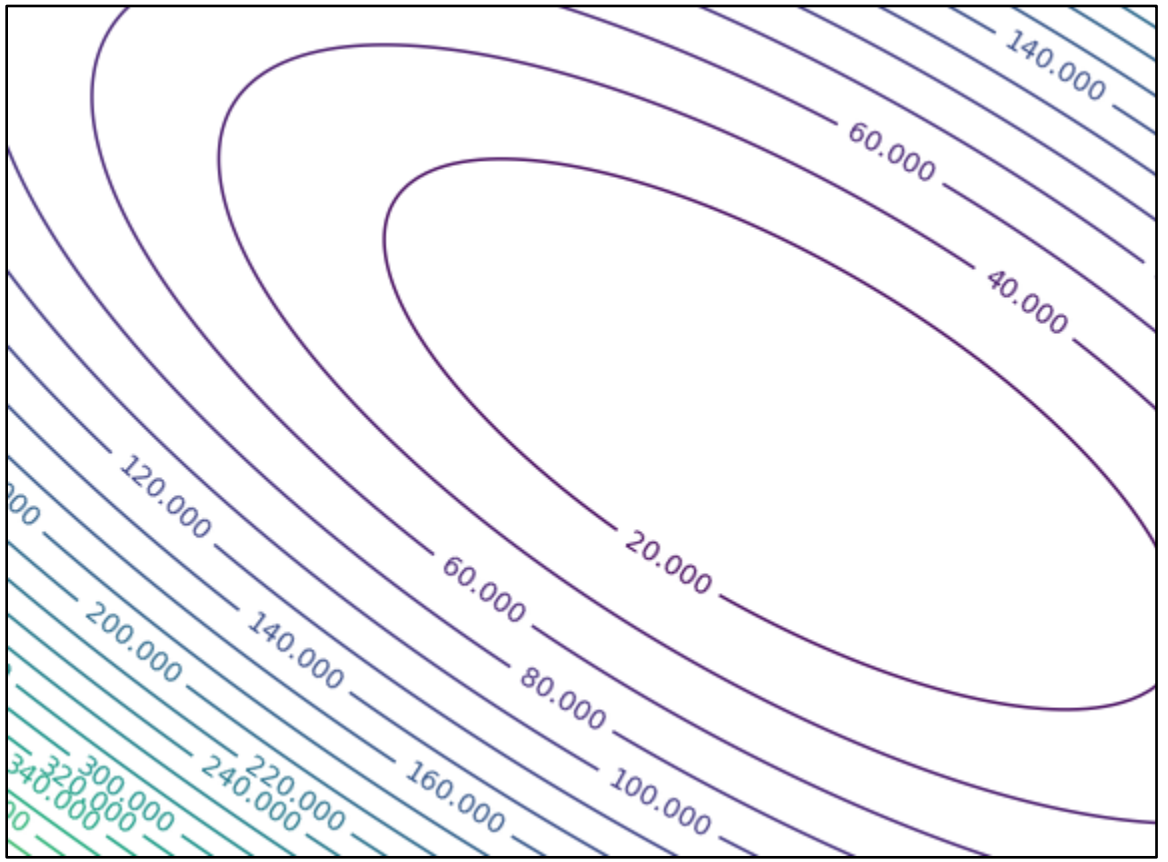

b

m

Gradient

Shared Drive Folder / Lec_9_10_Regression / Group N

regression_3_search_visualization.ipynb

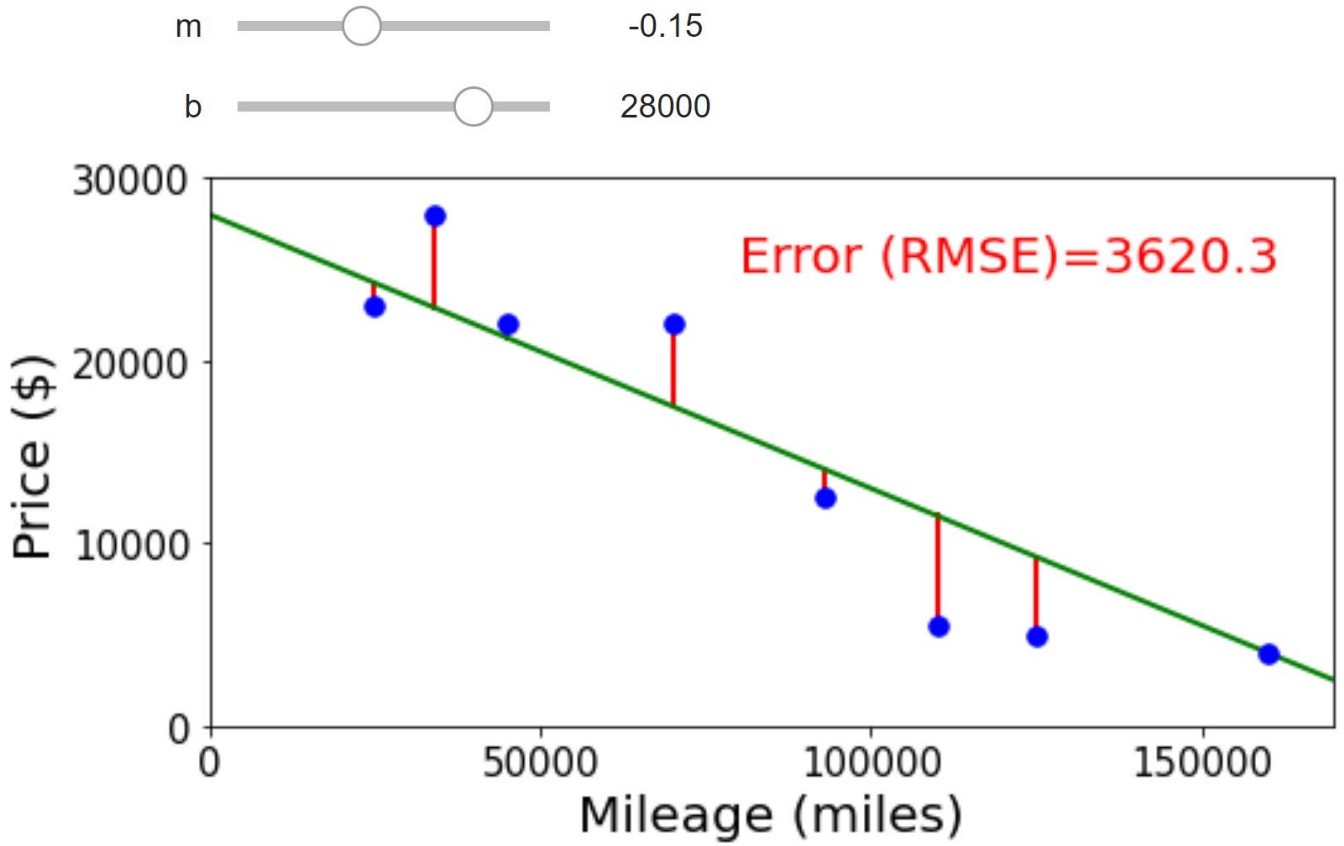

Optimization

Gradients

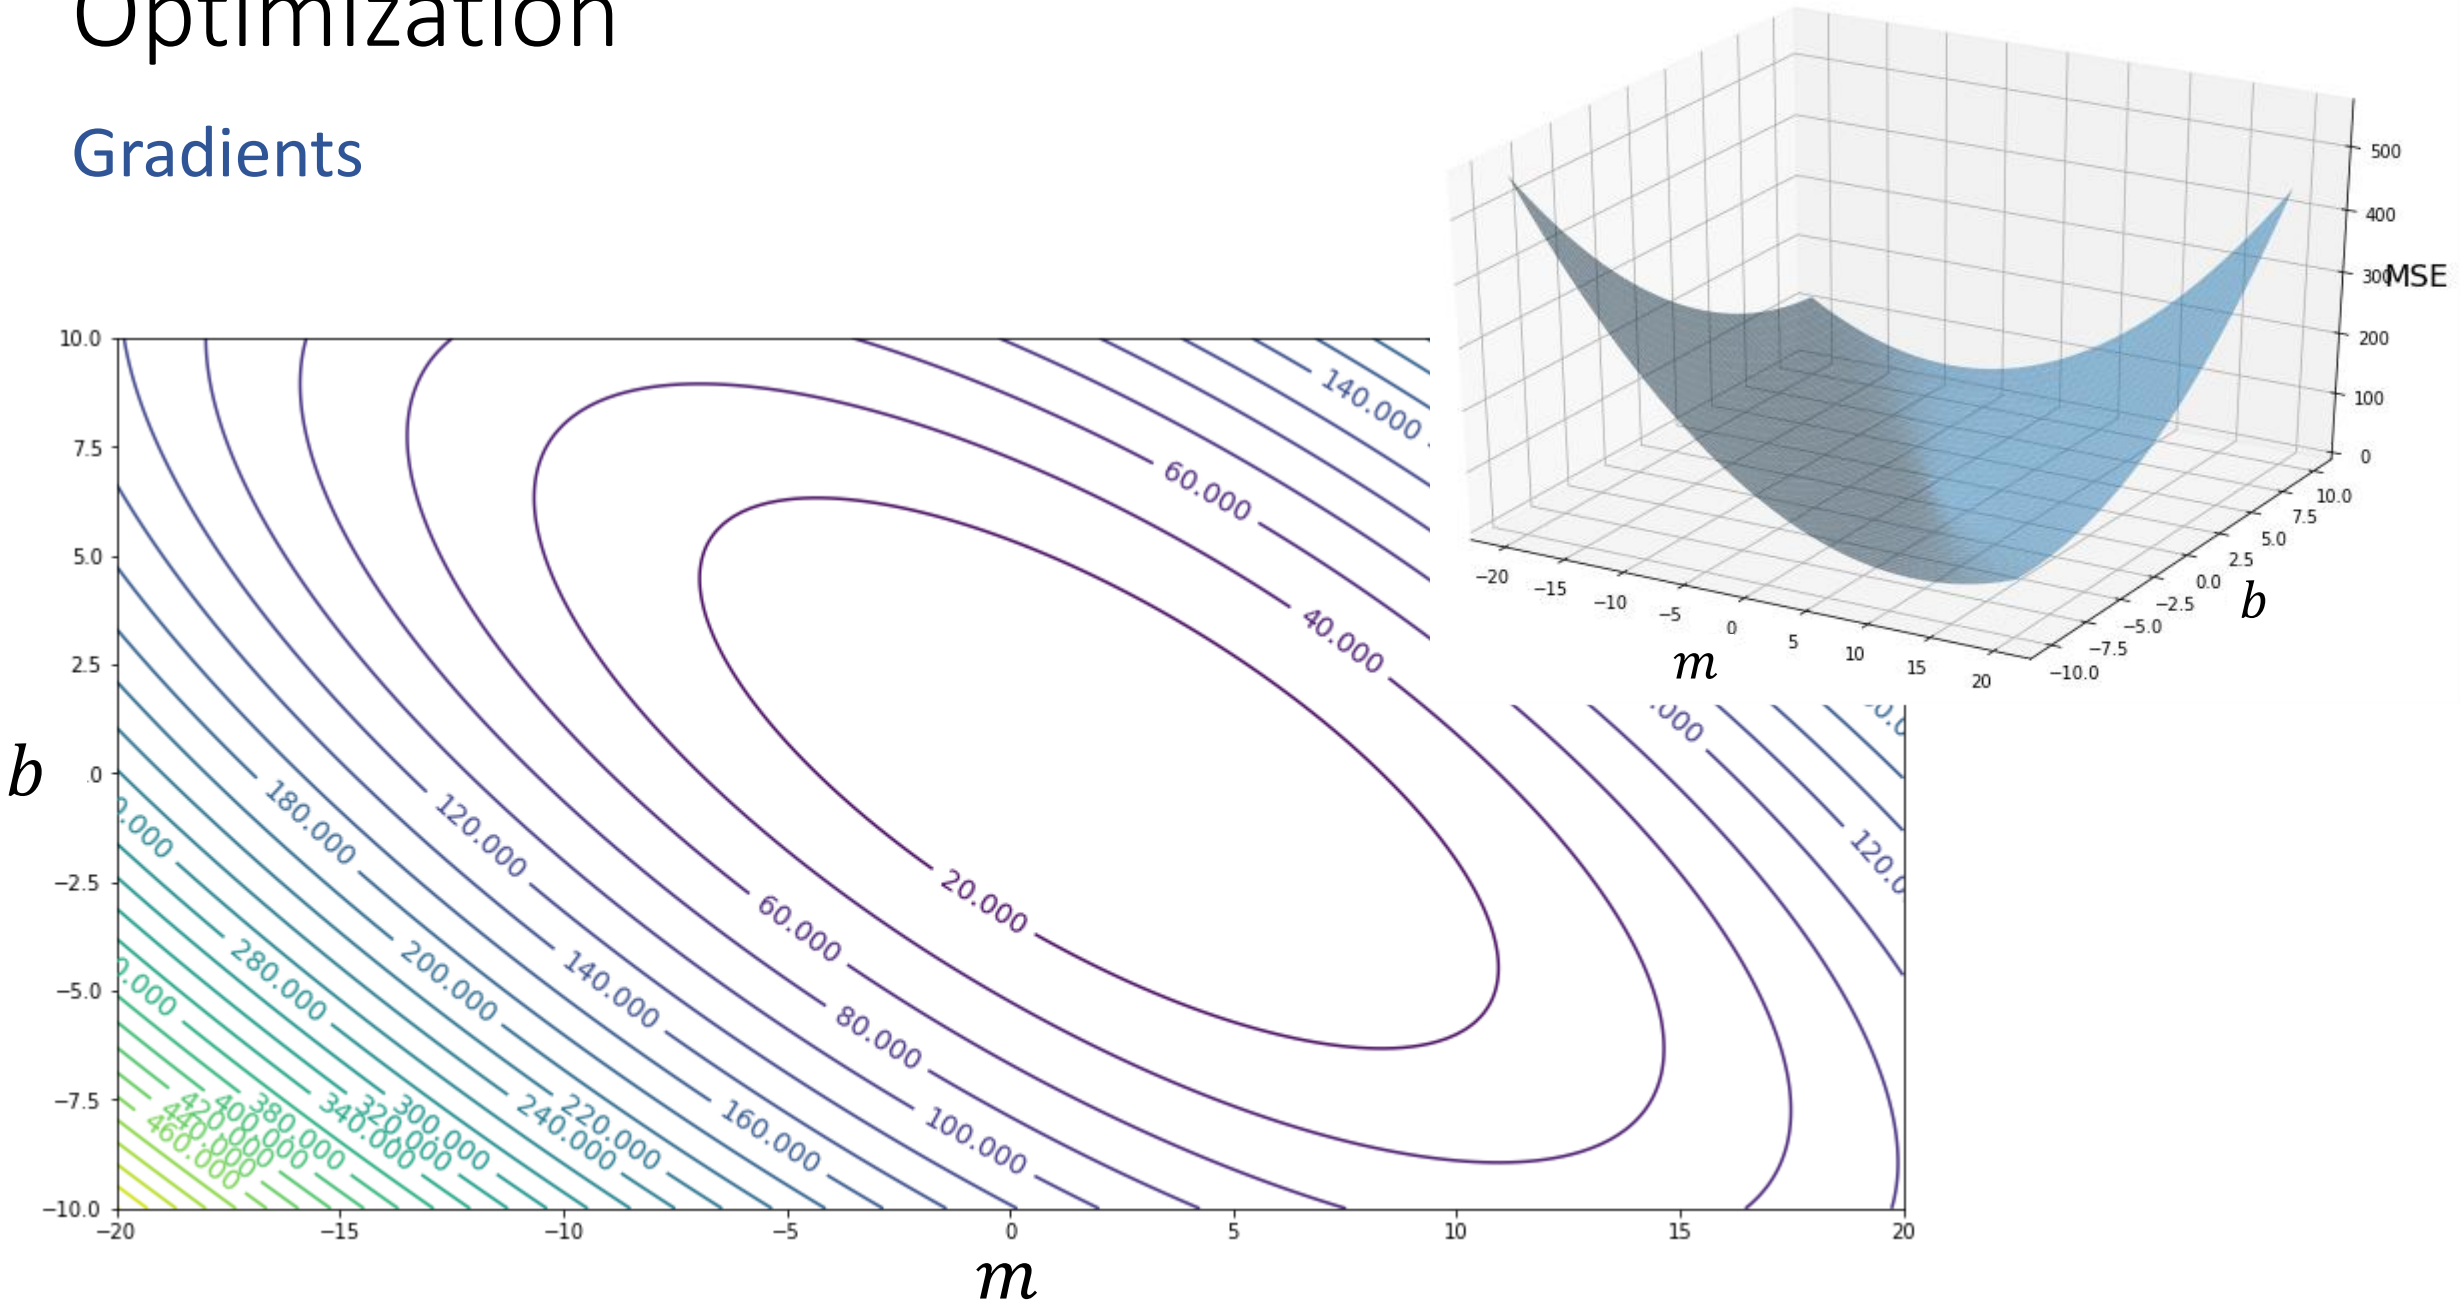

Optimization

Gradients

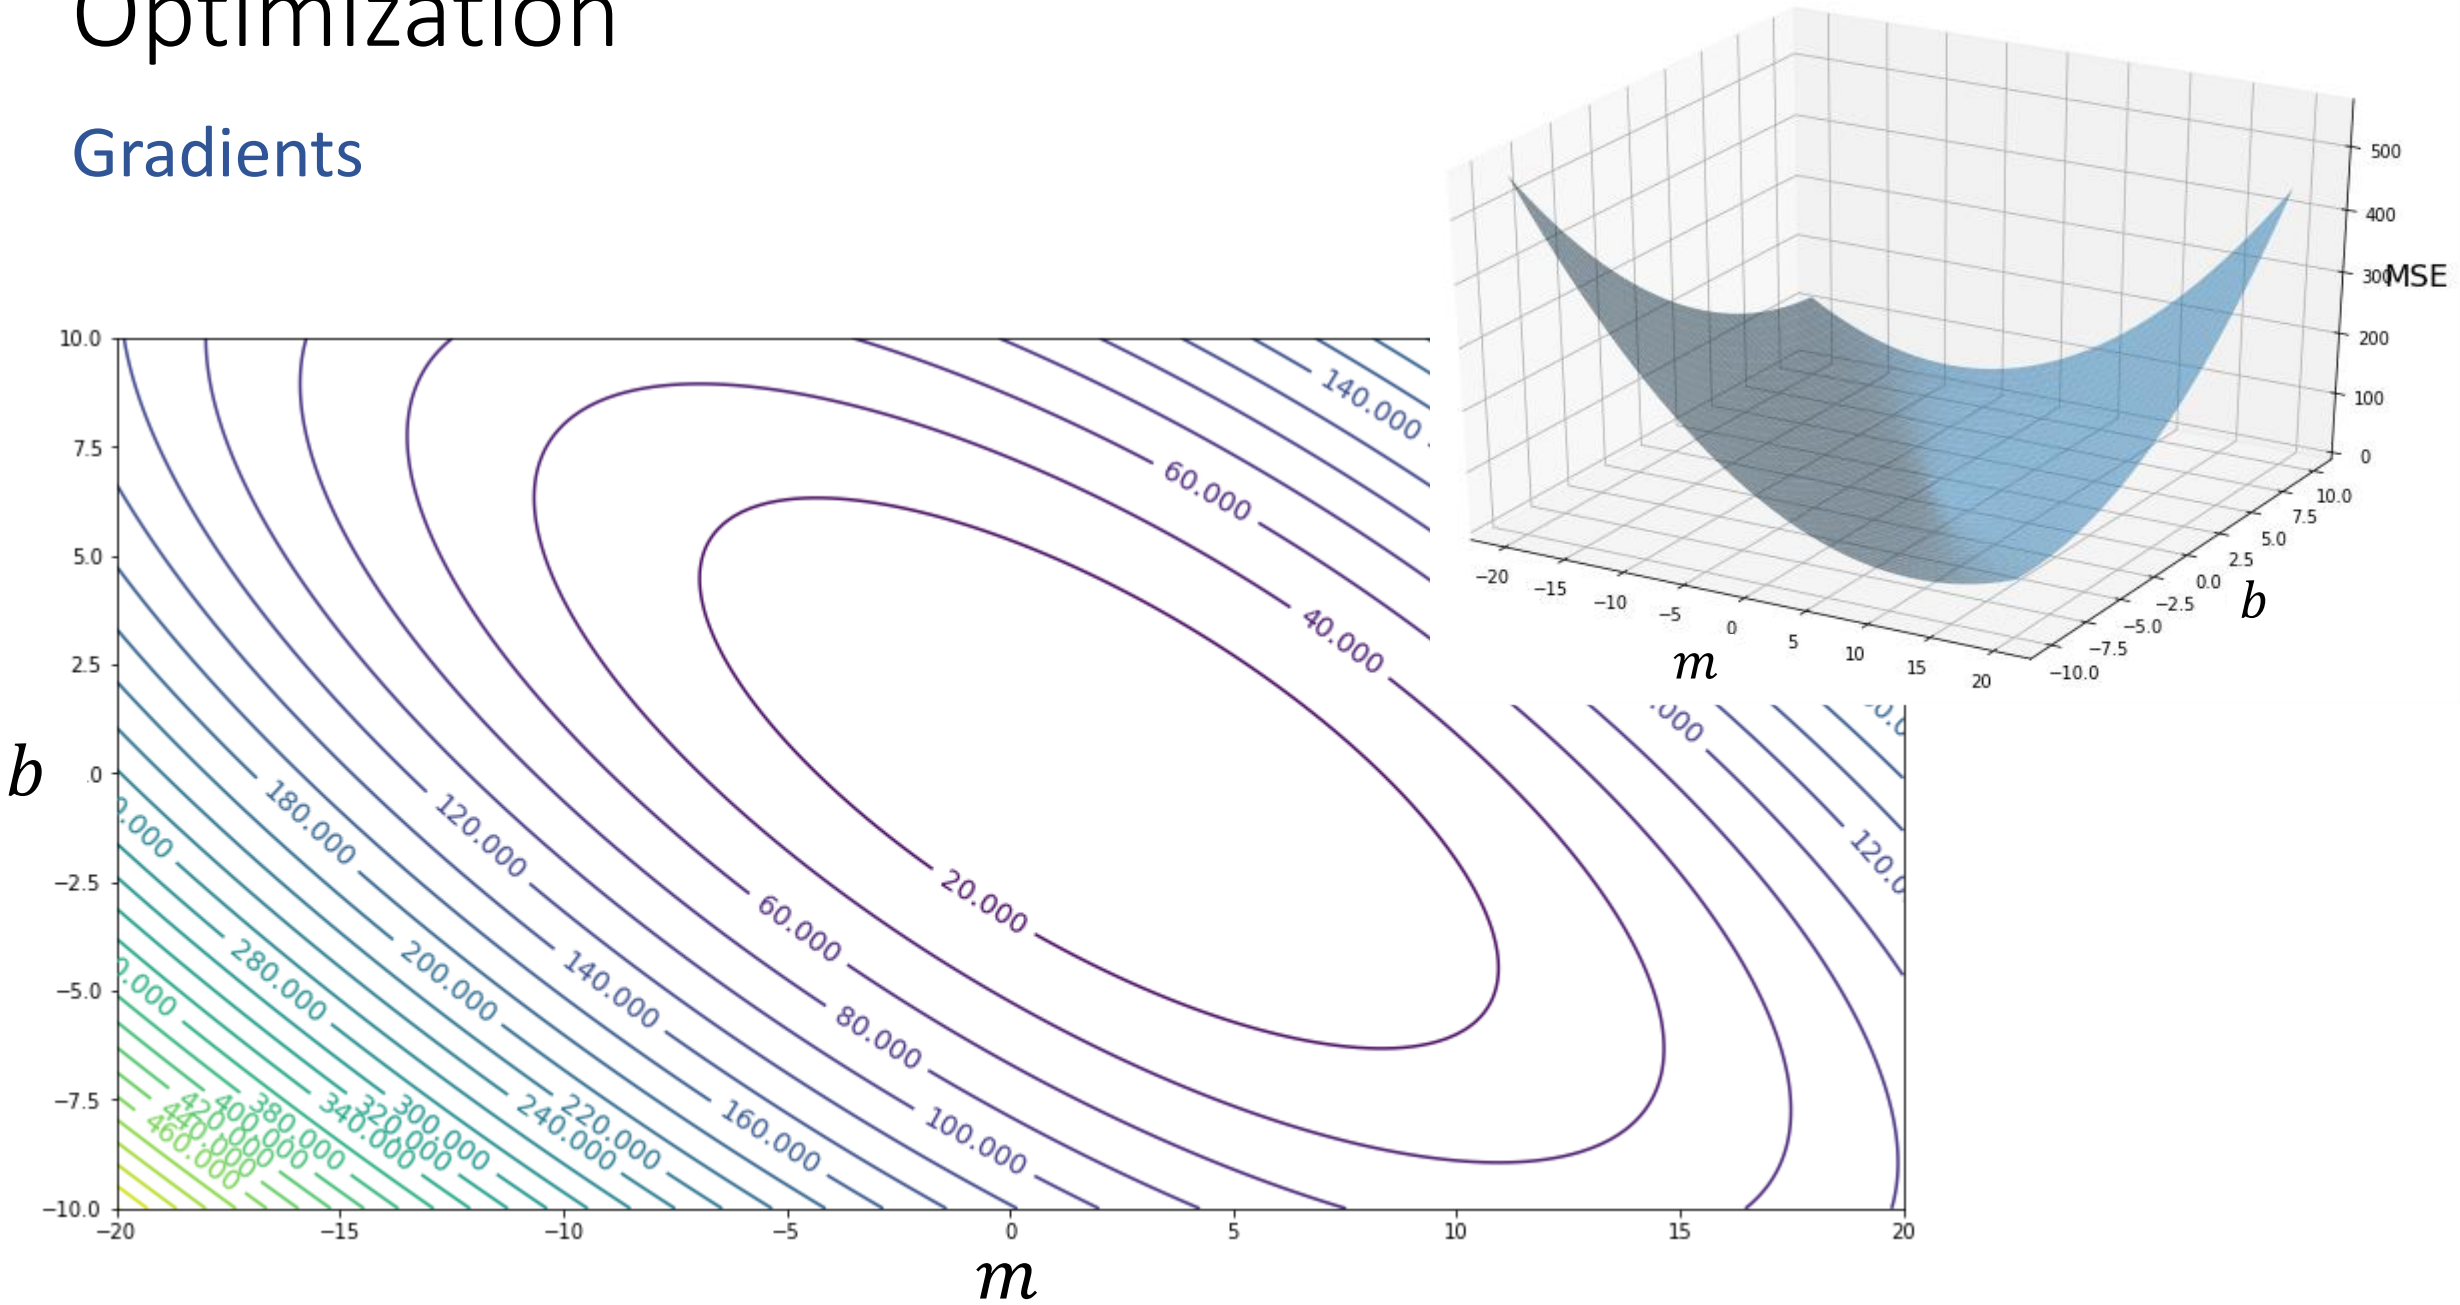

Optimization

Gradients

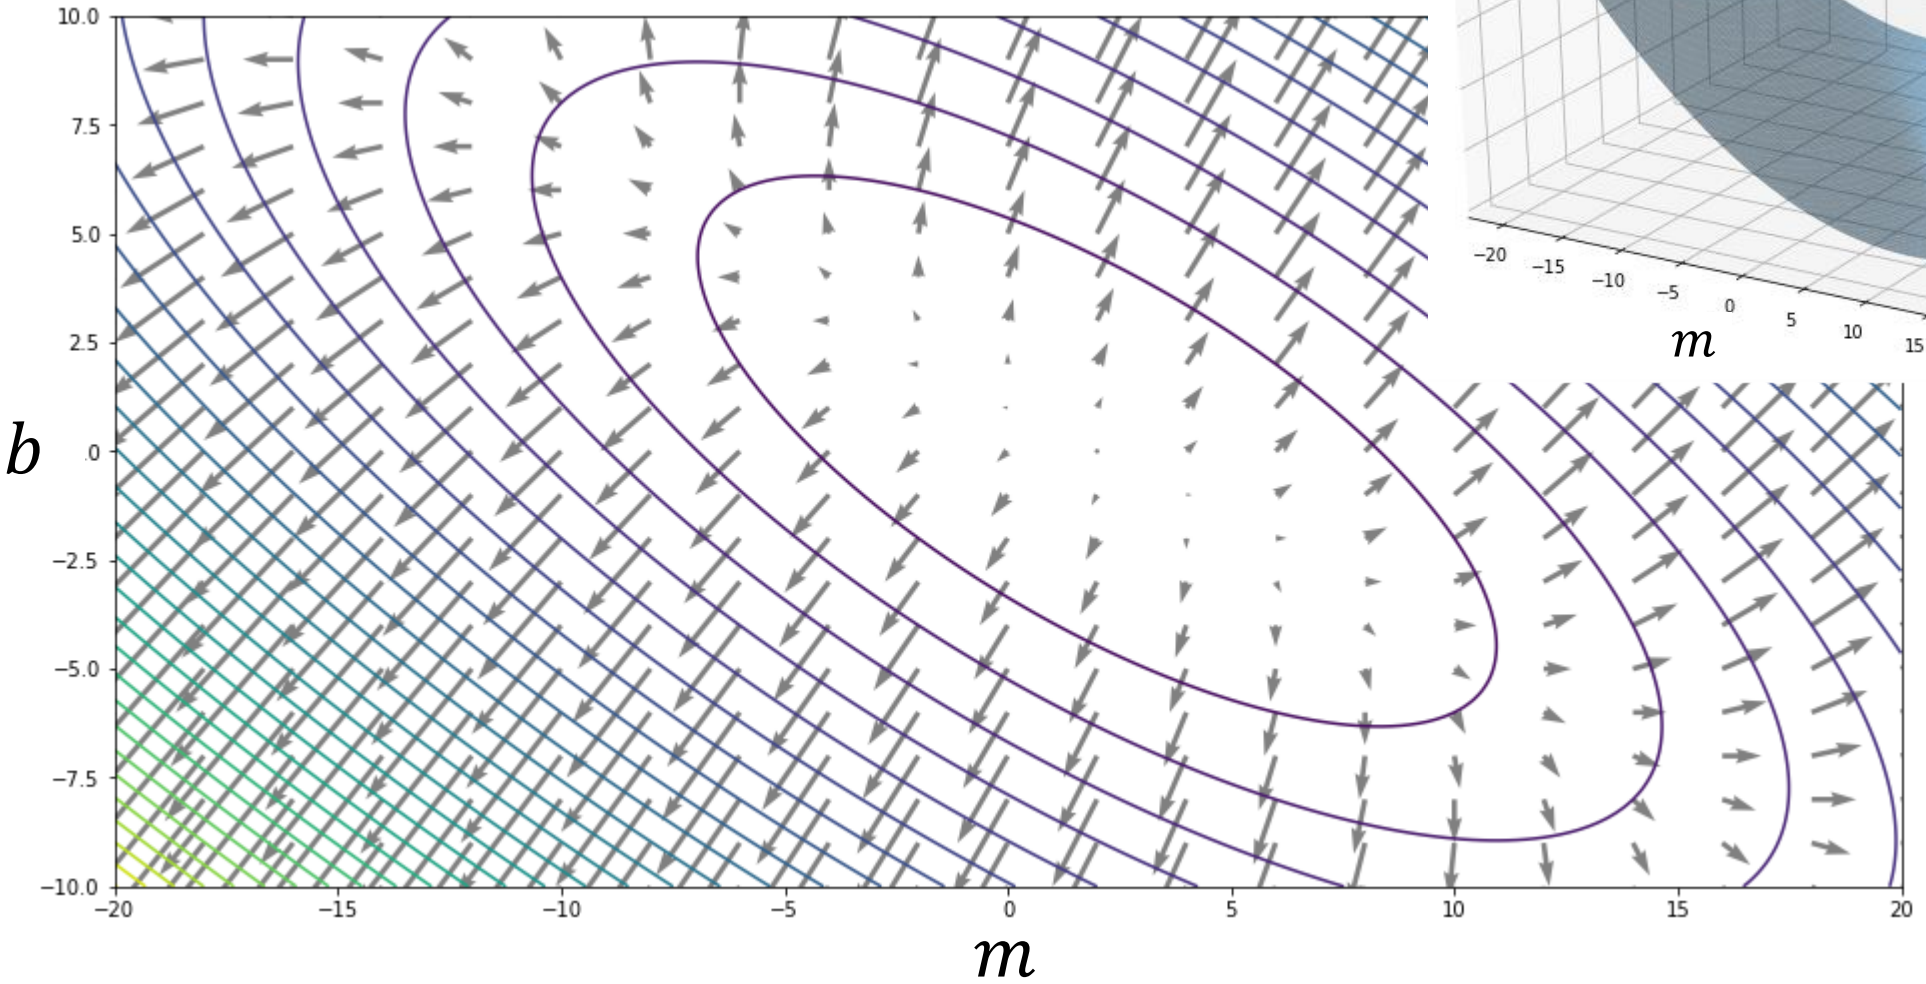

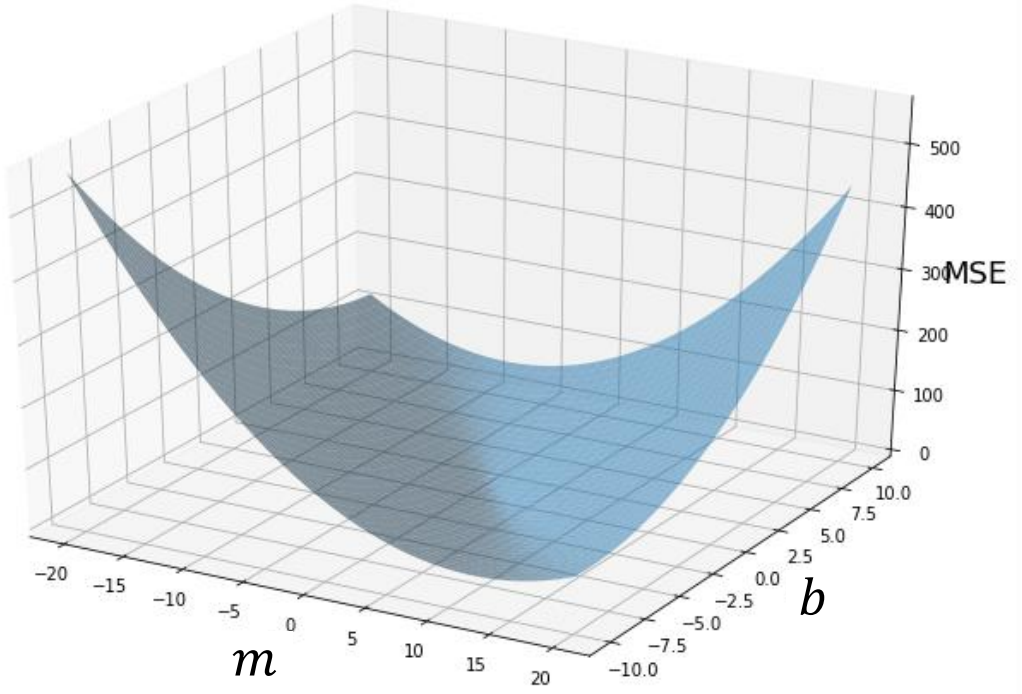

Adding More Input Features

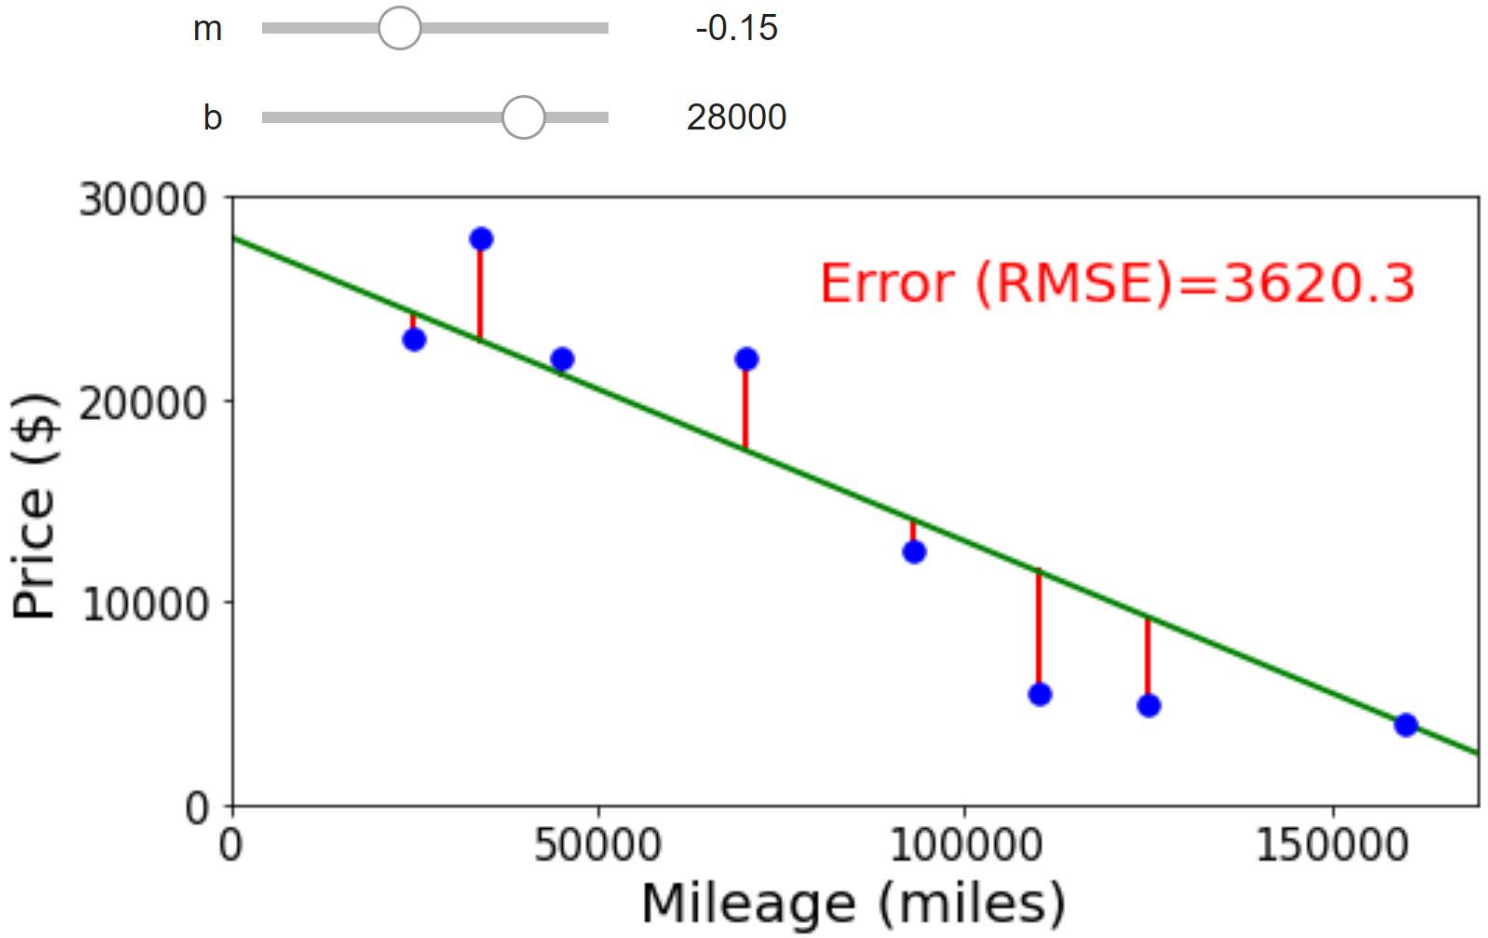

Adding More Input Features

Coin values

Adding More Input Features

Alien coins