

Lab 2: Empirical Risk and Simple Linear Regression Solutions

EECS 245, Spring 2026 at the University of Michigan

due for completion at 11:59PM Ann Arbor Time on Monday, May 11th, 2026

Name: _____

username: _____

Each lab worksheet will contain several activities, some of which will involve writing code and others that will involve writing math on paper. To receive credit for a lab, you must complete as many of the activities as you can in 2 hours and submit a PDF of your work to Gradescope. We will provide specific instructions on how to submit programming activities (e.g. submitting the notebook or including a screenshot of some output).

Feel free to work with others in the course, but you must submit individually.

Recap: The Modeling Recipe

In [Chapter 1.3](#), we introduced the three-step modeling recipe for finding optimal model parameters, which ultimately helps us make the best possible predictions.

1. Choose a model.

$$\underbrace{h(x_i) = w}_{\text{constant model}} \quad \underbrace{h(x_i) = w_0 + w_1 x_i}_{\text{simple linear regression model}}$$

2. Choose a loss function.

$$\underbrace{L_{\text{sq}}(y_i, h(x_i)) = (y_i - h(x_i))^2}_{\text{squared loss}} \quad \underbrace{L_{\text{abs}}(y_i, h(x_i)) = |y_i - h(x_i)|}_{\text{absolute loss}}$$

3. Minimize *average loss* (also called *empirical risk*) to find optimal model parameters.

- Constant model, squared loss: $R_{\text{sq}}(w) = \frac{1}{n} \sum_{i=1}^n (y_i - w)^2 \implies w^* = \bar{y}$
- Constant model, absolute loss: $R_{\text{abs}}(w) = \frac{1}{n} \sum_{i=1}^n |y_i - w| \implies w^* = \text{Median}(y_1, y_2, \dots, y_n)$
- Simple linear regression model, squared loss:

$$R_{\text{sq}}(w_0, w_1) = \frac{1}{n} \sum_{i=1}^n (y_i - (w_0 + w_1 x_i))^2 \implies w_1^* = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sum_{i=1}^n (x_i - \bar{x})^2}, \quad w_0^* = \bar{y} - w_1^* \bar{x}$$

Activity 1: Relative Squared Loss

Suppose we'd like to find the optimal parameter, w^* , for the constant model $h(x_i) = w$. To do so, we use the following loss function, called the **relative squared loss**:

$$L_{\text{rsq}}(y_i, h(x_i)) = \frac{(y_i - h(x_i))^2}{y_i}$$

What value of w minimizes the average loss (i.e. empirical risk) when using the relative squared loss function – that is, what is w^* ? Your answer should only be in terms of the variables $n, y_1, y_2, \dots, y_n$, and any constants.

Solution:

Since $h(x_i) = w$ for the constant model, relative squared loss for the constant model is:

$$L_{\text{rsq}}(y_i, w) = \frac{(y_i - w)^2}{y_i}$$

and so average relative squared loss for the constant model is:

$$R_{\text{rsq}}(w) = \frac{1}{n} \sum_{i=1}^n \frac{(y_i - w)^2}{y_i}$$

To find the value of w that minimizes $R_{\text{rsq}}(w)$, we'll first find its first derivative and set it to zero. The first derivative of $R_{\text{rsq}}(w)$ is:

$$\begin{aligned} \frac{d}{dw} R_{\text{rsq}}(w) &= \frac{d}{dw} \left(\frac{1}{n} \sum_{i=1}^n \frac{(y_i - w)^2}{y_i} \right) \\ &= \frac{1}{n} \sum_{i=1}^n \frac{d}{dw} \left(\frac{(y_i - w)^2}{y_i} \right) \end{aligned}$$

At this point, it'll be useful to step aside and find the derivative of $L_{\text{rsq}}(y_i, w)$ with respect to w , as this is the expression being summed. The derivative of $L_{\text{rsq}}(y_i, w)$ with respect to w is:

$$\begin{aligned} \frac{d}{dw} L_{\text{rsq}}(y_i, w) &= \frac{d}{dw} \frac{(y_i - w)^2}{y_i} \\ &= \frac{1}{y_i} \cdot \frac{d}{dw} (y_i - w)^2 \\ &= \frac{1}{y_i} \cdot 2(y_i - w) \cdot (-1) \\ &= -2 \cdot \frac{y_i - w}{y_i} \\ &= \boxed{2 \cdot \frac{w}{y_i} - 2} \end{aligned}$$

Back to $\frac{d}{dw} R_{\text{rsq}}(w)$, we have:

$$\begin{aligned} \frac{d}{dw} R_{\text{rsq}}(w) &= \frac{1}{n} \sum_{i=1}^n \frac{d}{dw} \left(\frac{(y_i - w)^2}{y_i} \right) \\ &= \frac{1}{n} \sum_{i=1}^n \left(2 \cdot \frac{w}{y_i} - 2 \right) \\ &= \frac{2w}{n} \sum_{i=1}^n \left(\frac{1}{y_i} \right) - \frac{1}{n} \sum_{i=1}^n 2 \\ &= \frac{2w}{n} \sum_{i=1}^n \left(\frac{1}{y_i} \right) - 2 \end{aligned}$$

Solution: (continued) Setting this equal to 0 yields:

$$\frac{2w}{n} \sum_{i=1}^n \left(\frac{1}{y_i}\right) - 2 = 0$$

$$\frac{w}{n} \sum_{i=1}^n \left(\frac{1}{y_i}\right) = 1$$


$$w^* = \frac{1}{\frac{1}{n} \sum_{i=1}^n \left(\frac{1}{y_i}\right)}$$

$$w^* = \boxed{\frac{n}{\sum_{i=1}^n \frac{1}{y_i}}}$$

This is known as the **harmonic mean** of $y_1, y_2, \dots, y_n$.

Activity 2: Rapid Fire

Consider a dataset of n **integers**, $y_1, y_2, \dots, y_n$, whose histogram is given below:



a) Which of the following is closest to the constant prediction w^* that minimizes:

$$\frac{1}{n} \sum_{i=1}^n \begin{cases} 0 & y_i = w \\ 1 & y_i \neq w \end{cases}$$

- ☐ 1
 ☐ 5
 ☐ 6
 ☐ 7
 ☐ 11
 ☐ 15
 ☐ 30

Solution: 30.

The minimizer of average 0-1 loss is the **mode**.

See: [Chapter 1.4: Beyond Absolute and Squared Loss](#)

b) Which of the following is closest to the constant prediction w^* that minimizes:

$$\frac{1}{n} \sum_{i=1}^n |y_i - w|$$

- ☐ 1
 ☐ 5
 ☐ 6
 ☐ 7
 ☐ 11
 ☐ 15
 ☐ 30

Solution: 7.

The minimizer of average absolute loss is the **median**. The outliers near 30 shift it from 6 to 7.

c) Which of the following is closest to the constant prediction w^* that minimizes:

$$\frac{1}{n} \sum_{i=1}^n (y_i - w)^2$$

☐ 1 ☐ 5 ☐ 6 ☐ 7 ☐ 11 ☐ 15 ☐ 30

Solution: 11.

The minimizer of average squared loss is the **mean**, pulled upward by the heavy right tail, so it's above the median (7) and closest to 11.

d) Which of the following is closest to the constant prediction w^* that minimizes:

$$\lim_{p \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n |y_i - w|^p$$

Hint: Think about the effect of outliers.

☐ 1 ☐ 5 ☐ 6 ☐ 7 ☐ 11 ☐ 15 ☐ 30

Solution: 15.

As $p \rightarrow \infty$, the minimizer is the **midrange**, halfway between min and max.

Activity 3: Slope of Mean Absolute Error

Consider a dataset of 8 points, $y_1, y_2, \dots, y_8$ that are in sorted order, i.e. $y_1 < y_2 < \dots < y_8$.

Recall that mean absolute error, $R_{\text{abs}}(w)$, for the constant model $h(x_i) = w$ is defined as:

$$R_{\text{abs}}(w) = \frac{1}{n} \sum_{i=1}^n |y_i - w|$$

This is a piecewise linear function that changes slope at each data point. The slope of $R_{\text{abs}}(w)$ at any w that is not a data point is:

$$\frac{d}{dw} R_{\text{abs}}(w) = \frac{\# \text{ left of } w - \# \text{ right of } w}{n}$$

Suppose that $y_4 = 10$, $y_5 = 14$, $y_6 = 22$, and $R_{\text{abs}}(11) = 9$. What is $R_{\text{abs}}(22)$?

Hint: You don't have all 8 of the y -values, so you can't find $R_{\text{abs}}(22)$ just by plugging in numbers into the formula for $R_{\text{abs}}(w)$. Instead, think about how to use the slope formula.

Solution:

$$R_{\text{abs}}(22) = 11.$$

We can write the points given to us as:

$$y_1, y_2, y_3, 10, 14, 22, y_7, y_8$$

Since there are an even number of data points ($n = 8$), the minimizer of absolute error is not a single point but the entire interval between the two middle points. Here, the middle two are 10 and 14, so every $w \in [10, 14]$ minimizes $R_{\text{abs}}(w)$. This explains why the error is *flat* inside that interval: the number of points on the left equals the number on the right, so shifting w around does not change the error. As a result, $R_{\text{abs}}(11) = 9$ and $R_{\text{abs}}(14) = 9$.

Once we move beyond 14, the balance breaks. There are now five points to the left and only three to the right, so the slope of $R_{\text{abs}}(w)$ becomes positive. The slope formula tells us:

$$\frac{d}{dw} R_{\text{abs}}(w) = \frac{\# \text{ left of } w - \# \text{ right of } w}{n}$$

so for any $w \in (14, 22)$ we have

$$\frac{d}{dw} R_{\text{abs}}(w) = \frac{5 - 3}{8} = \frac{1}{4}.$$

This means that for every one unit we move to the right of $w = 14$, the error increases by $\frac{1}{4}$. Moving from $w = 14$ to $w = 22$ is a distance of $22 - 14 = 8$ units, so the error increases by

$$8 \cdot \frac{1}{4} = 2.$$

Adding this to the baseline error of $R_{\text{abs}}(14) = 9$, we get:

$$\begin{aligned} R_{\text{abs}}(22) &= R_{\text{abs}}(14) + (22 - 14) \cdot \frac{1}{4} \\ &= 9 + 2 = 11. \end{aligned}$$

Activity 4: Programming

Complete the tasks in the lab02.ipynb notebook.

There are two ways to access the supplemental Jupyter Notebook:

- **Option 1 (preferred):** Set up a Jupyter Notebook environment locally, use git to clone our [course repository](#), and open labs/lab02/lab02.ipynb. For instructions on how to do this, see the [Environment Setup](#) page of the course website.
- **Option 2:** Click [here](#) to open lab02.ipynb on DataHub. Before doing so, read the instructions on the [Environment Setup](#) page on how to use the DataHub.

Once you're done, include a screenshot of your completed Activity 4 implementation in your PDF submission of Lab 2 to Gradescope, making sure to include proof that the (local) autograder passed.

Activity 5: Visualizing Changes in the Data


The problems in this final activity will help you visualize how changes in the data affect the optimal simple linear regression line. To recap, this is the line $h(x_i) = w_0 + w_1 x_i$ defined by:

$$w_1^* = r \frac{\sigma_y}{\sigma_x} = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sum_{i=1}^n (x_i - \bar{x})^2} \quad w_0^* = \bar{y} - w_1^* \bar{x}$$

r is the correlation coefficient between x and y , σ_x is the standard deviation of x , and σ_y is the standard deviation of y .

Assume all data is in the first quadrant, i.e. all x_i and y_i are positive.

- a) In each dataset shown below, how will the slope and intercept of the regression line change if we move the red point in the direction of the arrow?



Solution:

Dataset 1 (top): If this point is moved upward, the slope of the regression line will increase and the intercept will decrease.

Dataset 2 (bottom): If this point is moved downward, the slope of the regression line will increase and the intercept will decrease, just like in Dataset 1.


In Dataset 1, the point on the right end shifts upward, and pulls the right side of the line up. Similarly, in Dataset 2 the left point being shifted down causes the line to be pulled down on the left side. In both of these cases, we can think of the line as becoming **more flat** as a result of the shifting point in both datasets.

Because both lines have a negative slope, if they become less steep the intercept will shift downwards.

b) Compare two different possible changes to the dataset shown below.

- Move the dashed point down c units.
- Move the solid point down c units.

Which move will change the slope of the regression line more? Why? *Hint: We're not looking for a formal proof. But, if you want to read more, look at [Chapter 2.3](#).*



Solution:

The slope of the regression line is given by

$$b = \frac{\sum (x_i - \bar{x})(y_i - \bar{y})}{\sum (x_i - \bar{x})^2}.$$

When we move a single point down by c units, the only part of the slope formula that changes is the numerator

$$\sum (x_i - \bar{x})(y_i - \bar{y}),$$

which measures the covariance between x and y . The change in covariance due to moving one point is proportional to $(x_i - \bar{x})(-c)$.

- The dashed point is closer to the mean $\bar{x}$, so $|x_i - \bar{x}|$ is relatively small. Moving this point has only a modest effect on the slope.
- The solid point is farther from the mean $\bar{x}$, so $|x_i - \bar{x}|$ is larger. Moving this point has a much larger effect on the slope.

$$\Delta \text{slope} \propto (x_i - \bar{x})(-c).$$

Therefore, **moving the solid point down c units will change the slope of the regression line more than moving the dashed point.** This is because points farther from the mean of x have greater *leverage* on the slope.

The rest of this worksheet is extra practice. Don't feel pressured to answer all of these problems in lab, but make sure to attempt them at some point.

Activity 6: Relative Squared Loss, Continued

Recall the formula for **relative squared loss** from Activity 1:

$$L_{\text{rsq}}(y_i, h(x_i)) = \frac{(y_i - h(x_i))^2}{y_i}$$

- a) Let $C(y_1, y_2, \dots, y_n)$ be your minimizer w^* from Activity 1. That is, for a particular dataset $y_1, y_2, \dots, y_n$, $C(y_1, y_2, \dots, y_n)$ is the value of w that minimizes empirical risk for relative squared loss on that dataset.

What is the value of $\lim_{y_4 \rightarrow \infty} C(1, 3, 5, y_4)$ in terms of $C(1, 3, 5)$? Your answer should involve the function C and/or one or more constants.

Hint: To notice the pattern, evaluate $C(1, 3, 5, 100)$, $C(1, 3, 5, 10000)$, and $C(1, 3, 5, 1000000)$.

Solution:

$$\begin{aligned} \lim_{y_4 \rightarrow \infty} C(1, 3, 5, y_4) &= \lim_{y_4 \rightarrow \infty} \frac{4}{\frac{1}{1} + \frac{1}{3} + \frac{1}{5} + \frac{1}{y_4}} \\ &= \frac{4}{\frac{1}{1} + \frac{1}{3} + \frac{1}{5} + 0} \\ &= \frac{4}{\frac{3}{3} \cdot \frac{1}{\frac{1}{1} + \frac{1}{3} + \frac{1}{5}}} \\ &= \frac{4}{3} \cdot C(1, 3, 5) \end{aligned}$$

- b) What is the value of $\lim_{y_4 \rightarrow 0} C(1, 3, 5, y_4)$? Again, your answer should involve the function C and/or one or more constants.

Solution:

$$\begin{aligned} \lim_{y_4 \rightarrow 0} C(1, 3, 5, y_4) &= \lim_{y_4 \rightarrow 0} \frac{4}{\frac{1}{1} + \frac{1}{3} + \frac{1}{5} + \frac{1}{y_4}} \\ &= \frac{4}{\frac{1}{1} + \frac{1}{3} + \frac{1}{5} + \infty} \\ &= \frac{4}{\infty} = 0 \end{aligned}$$

- c) Based on the results of the previous two parts, when is the prediction $C(y_1, y_2, \dots, y_n)$ robust to outliers? When is it not robust to outliers?

Solution:

$C(y_1, y_2, \dots, y_n)$ is great at ignoring large outliers. No matter how large you make any particular value, $C(y_1, y_2, \dots, y_n)$ is upper-bounded by $\frac{n}{n-1}$ multiplied by the value of C applied to all data points excluding the large outlier. This is as opposed to the regular “arithmetic mean”, where if you make a single data point arbitrarily large, the mean also becomes arbitrarily large (i.e. if $y_n \rightarrow \infty$, then $\text{Mean}(y_1, y_2, \dots, y_n) \rightarrow \infty$ too).

However, $C(y_1, y_2, \dots, y_n)$ is not robust to small outliers. As a particular data point approaches 0, the value of $C(y_1, y_2, \dots, y_n)$ also approaches 0 no matter how large the other data points are.