

- Suppose we're predicting the **sale price** of a house based on its **size** (number of square feet), **latitude**, and **longitude**. We have a training set of 1000 houses.
 - This is a (**circle one**) regression / classification problem.
 - X .shape = _____ y .shape = _____
 - Suppose we compute a predicted price by using $y_{pred} = X @ w + b$, **where X has the shape you wrote above**. What shapes must w and b have?

w .shape = _____ b .shape = _____
 - Cross-Entropy**: Suppose the ABC classifier predicts probabilities $p = [0.25, 0.25, 0.5]$ (corresponding to A, B, and C respectively) for a single example. *Use log base 2 so you can do the math in your head.*
 - What is the categorical cross-entropy loss if the image was actually an A? _____
 - What is the categorical cross-entropy loss if the image was actually a C? _____
 - Suppose two classification models have roughly similar accuracy but model Y has a *much higher log loss* (categorical cross-entropy) than model Z. **What is going on?**
 - (*preview*) The last layer in the ABC classifier model computes $scores = something @ w + b$, where something is something we'll discuss later. **What shape does scores have?** $scores.shape =$ _____. In light of that, $w.shape =$ _____ and $b.shape =$ _____.
-

- Suppose we're predicting the **sale price** of a house based on its **size** (number of square feet), **latitude**, and **longitude**. We have a training set of 1000 houses.
 - This is a (**circle one**) regression / classification problem.
 - X .shape = _____ y .shape = _____
- Suppose we compute a predicted price by using $y_{pred} = X @ w + b$, **where X has the shape you wrote above**. What shapes must w and b have?

w .shape = _____ b .shape = _____
- Cross-Entropy**: Suppose the ABC classifier predicts probabilities $p = [0.25, 0.25, 0.5]$ (corresponding to A, B, and C respectively) for a single example. *Use log base 2 so you can do the math in your head.*
 - What is the categorical cross-entropy loss if the image was actually an A? _____
 - What is the categorical cross-entropy loss if the image was actually a C? _____
- Suppose two classification models have roughly similar accuracy but model Y has a *much higher log loss* (categorical cross-entropy) than model Z. **What is going on?**
- (*preview*) The last layer in the ABC classifier model computes $scores = something @ w + b$, where something is something we'll discuss later. **What shape does scores have?** $scores.shape =$ _____. In light of that, $w.shape =$ _____ and $b.shape =$ _____.

Before you leave, pick a couple of these questions to react to:

- What was the most important concept from today for you?
- What was the muddiest concept today?
- How does what we did today connect with what you've learned before?
- What would you like to review or clarify next time we meet?
- What are you curious, hopeful, or excited about?

Before you leave, pick a couple of these questions to react to:

- What was the most important concept from today for you?
- What was the muddiest concept today?
- How does what we did today connect with what you've learned before?
- What would you like to review or clarify next time we meet?
- What are you curious, hopeful, or excited about?