

1. Suppose we have a NumPy array `y_true` of true targets and a NumPy array `y_pred` of predictions. Write a Python expression to compute:
    - a. `mean_squared_error` =
    - b. `mean_absolute_error` =
  2. Which of these models can make predictions like “these neighborhoods are more expensive than those neighborhoods” (without transforming the input features). **Circle all that can** and then **explain**.  
linear regression                  logistic regression  
decision tree                      random forest
  3. Suppose two classification models have roughly similar accuracy but model A has a much higher log loss (categorical cross-entropy) than model B. **What is going on?**
  4. Suppose you want to classify images as A, B, or C, but your model isn't giving you probabilities. Instead it's giving you a continuous score of how much it thinks the image is A, B, or C. e.g., maybe it says `score_A = 1`, `score_B = 2`, `score_C = -1`. Discuss with your neighbors how you might add a post-processing step that takes in `score_A`, `score_B`, and `score_C` and outputs probabilities `p_A`, `p_B`, and `p_C` that sum to 1.
- 

1. Suppose we have a NumPy array `y_true` of true targets and a NumPy array `y_pred` of predictions. Write a Python expression to compute:
  - a. `mean_squared_error` =
  - b. `mean_absolute_error` =
2. Which of these models can make predictions like “these neighborhoods are more expensive than those neighborhoods” (without transforming the input features). **Circle all that can** and then **explain**.  
linear regression                  logistic regression  
decision tree                      random forest
3. Suppose two classification models have roughly similar accuracy but model A has a much higher log loss (categorical cross-entropy) than model B. **What is going on?**
4. Suppose you want to classify images as A, B, or C, but your model isn't giving you probabilities. Instead it's giving you a continuous score of how much it thinks the image is A, B, or C. e.g., maybe it says `score_A = 1`, `score_B = 2`, `score_C = -1`. Discuss with your neighbors how you might add a post-processing step that takes in `score_A`, `score_B`, and `score_C` and outputs probabilities `p_A`, `p_B`, and `p_C` that sum to 1.

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

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

---

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

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