# Module 4

Lecture 7

# How do I make my ML model better?

Sarah Groves, Ph.D.

Shayn Peirce-Cottler, Ph.D.

**1. DOWLOAD TEST DATA FROM CANVAS**

**2. FILL OUT SET COURSE EVALUATION (email)**

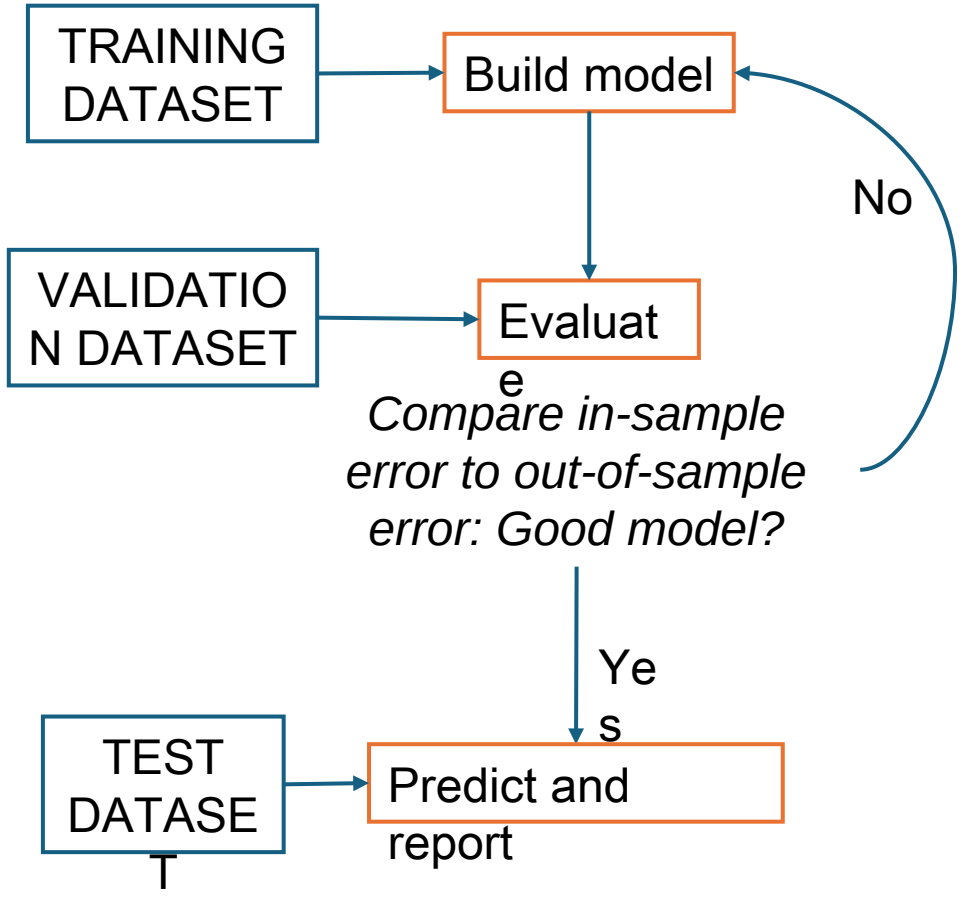
**3. I**

**MO**

**URVEY**



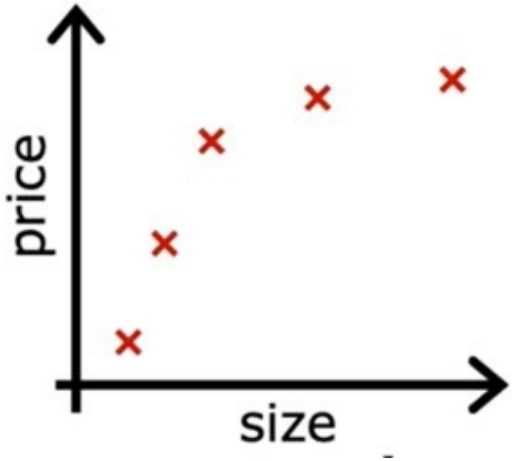
# How do I fix my model? (Validation step)




- If your model is underfitting…
  - ➢ *If you aren't seeing clear differences between classes in training or validation data*
  - ➢ *If your regression model doesn't predict well on your training or validation data*
  - ✓ Increase model complexity (linear → nonlinear regression, change decision boundary or model type)
  - ✓ Add features (new genes)
  - ✓ Change/Increase the number of the features
  - ✓ Change/decrease the number of classes
  - ✓ May signify your data is a continuum, not distinct clusters (can you try regression?)
  - ✓ Change classification threshold
  - ✓ Some classification methods support regularization (like Logistic Regression)
- If your model is overfitting…
  - ➢ *Low error for training data but high for validation data*
  - ➢ *You can visually see distinct clusters but out of sample error is high*
  - ✓ Try **regularization**
  - ✓ Reduce model complexity: cutoff decision trees earlier, change classification threshold

Once your model performs well, apply to unseen **test data** to get final out-of-sample error!

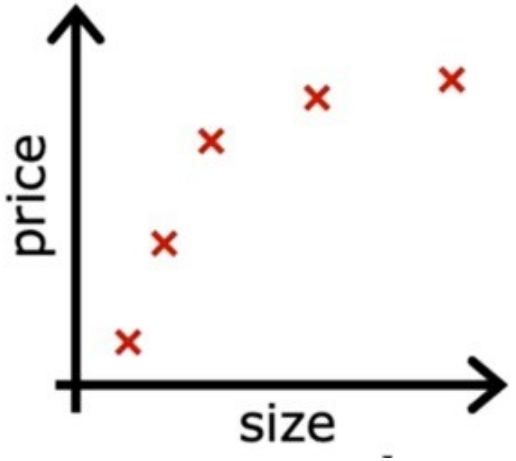
# Picking the right model: in sample vs out of sample error




Low complexity model

**Underfit model**

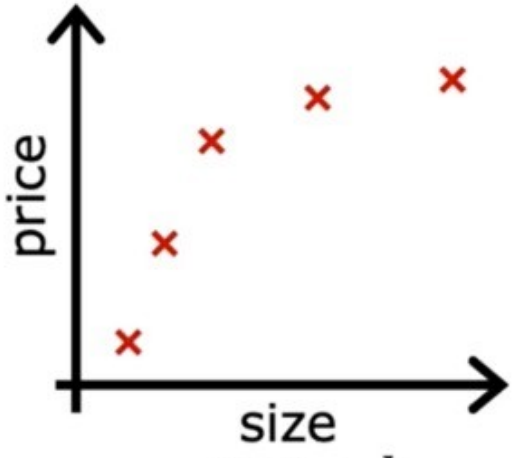
- Does not fit the training set well
- Not generalizable




Medium complexity model

**Well-fit model**

- Fits the training set well
- Generalizable to new data

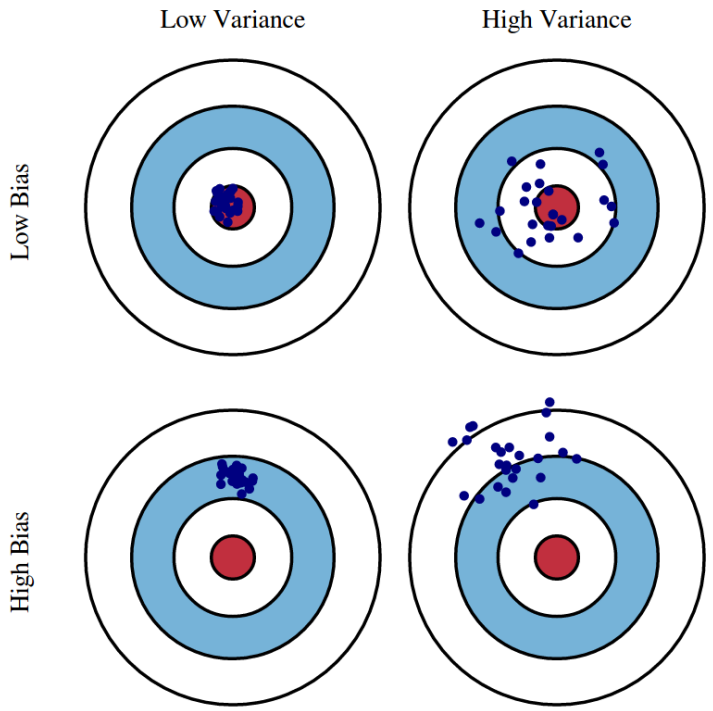


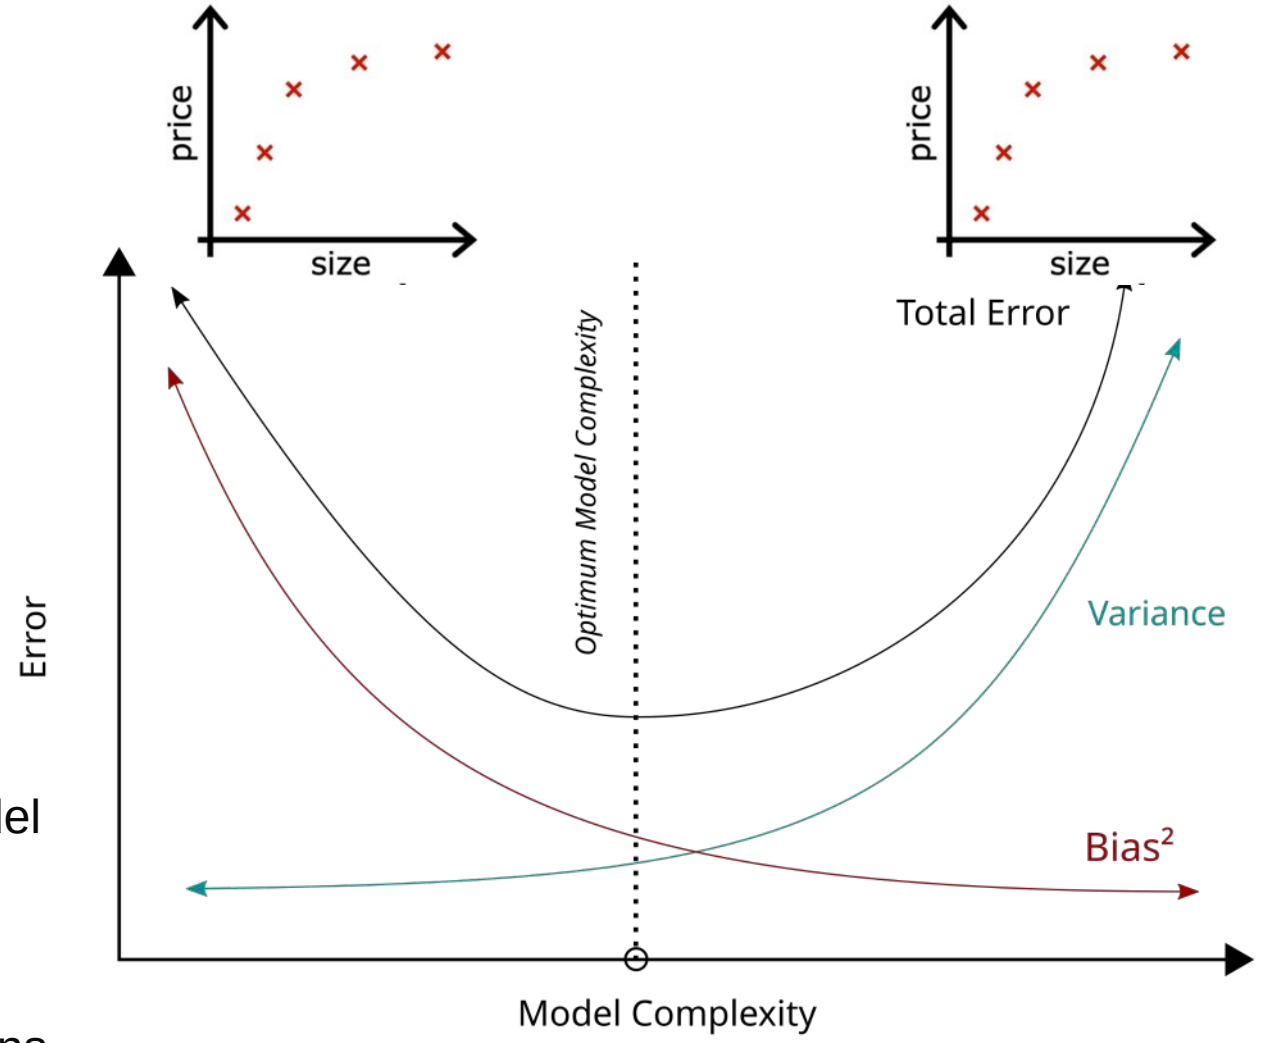

High complexity model

**Overfit model**

- Fits the training set extremely well
- Does not generalize well

# What do I do if my model is **overfitting**?

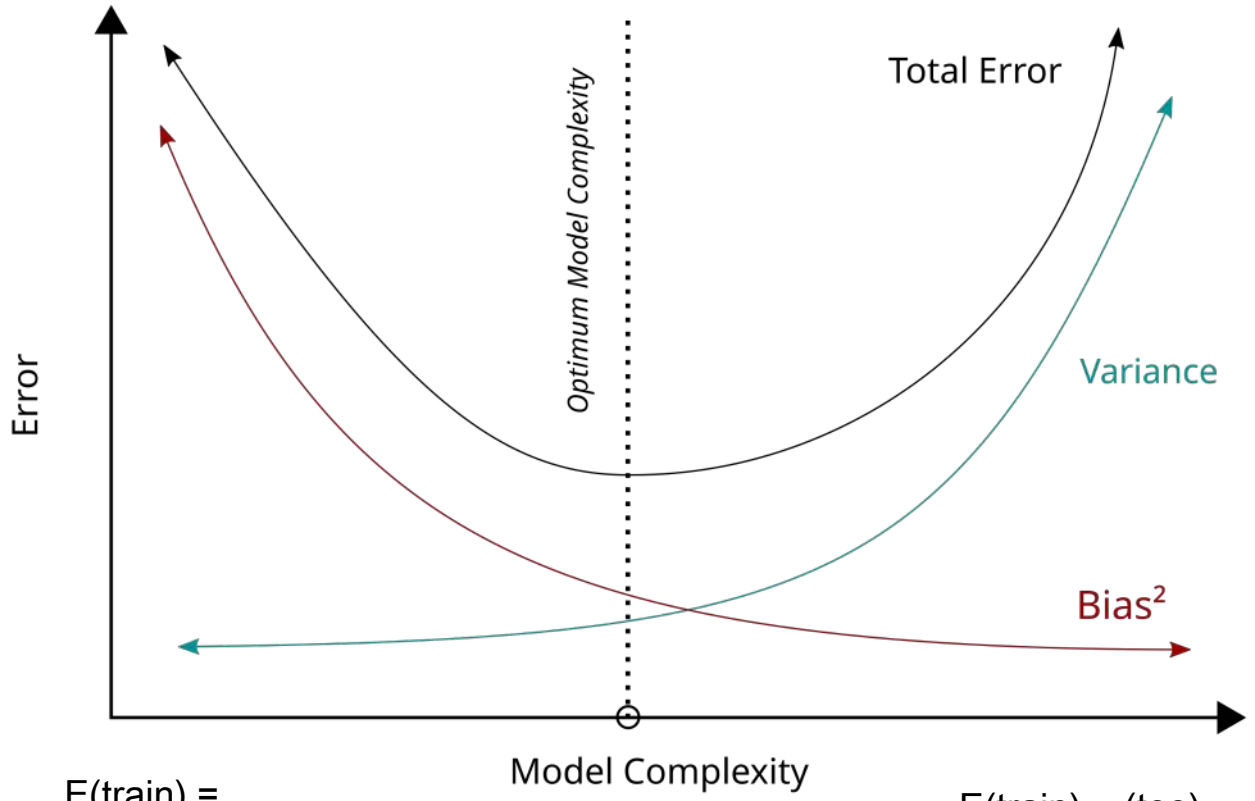



- **Bias:**
  - error from overly simplistic assumptions in the model
  - how far off is the average prediction from the true value?
- **Variance:**
  - error from the model's sensitivity to small fluctuations in the training data
  - how much would the model's predictions change if it were trained on a different dataset?

# Regularization on regression models

- Overfitting occurs when a model that is *too complex* is used, resulting in high variance
- Regularization is used to control the size of features
  - Forcing some feature coefficients to be close to zero reduces the complexity of the model, reducing possible overfitting



# How do we implement regularization?

- Add it to the cost function!
- Now our cost function balances between minimizing the SSE (or other cost, depending on application) and choosing small coefficients so we don't over-fit

Linear Regression

$$min[J(a_0, a_1)] = \frac{1}{2m} SSE + Regularization$$

$$= \frac{1}{2m} \sum_{i=1}^{m} (y_i - f_i)^2 + \frac{\lambda}{2m} \sum_{j=1}^{n} a_j^2$$

"Ridge regression"
"L2 norm regularization"

# Lots of options for regularization that you might hear about

- Ridge regression: decrease coefficients (i.e. w → small)

$$\min_w ||Xw - y||_2^2 + \alpha||w||_2^2$$

||x||2 = Euclidean distance

- From sklearn API
- Xw is the prediction from our function f (i.e. f(x) that we are comparing to true labels y)

- Lasso regression: fewer coefficients (i.e. w → 0)

$$\min_w \frac{1}{2n_{\text{samples}}}||Xw - y||_2^2 + \alpha||w||_1$$

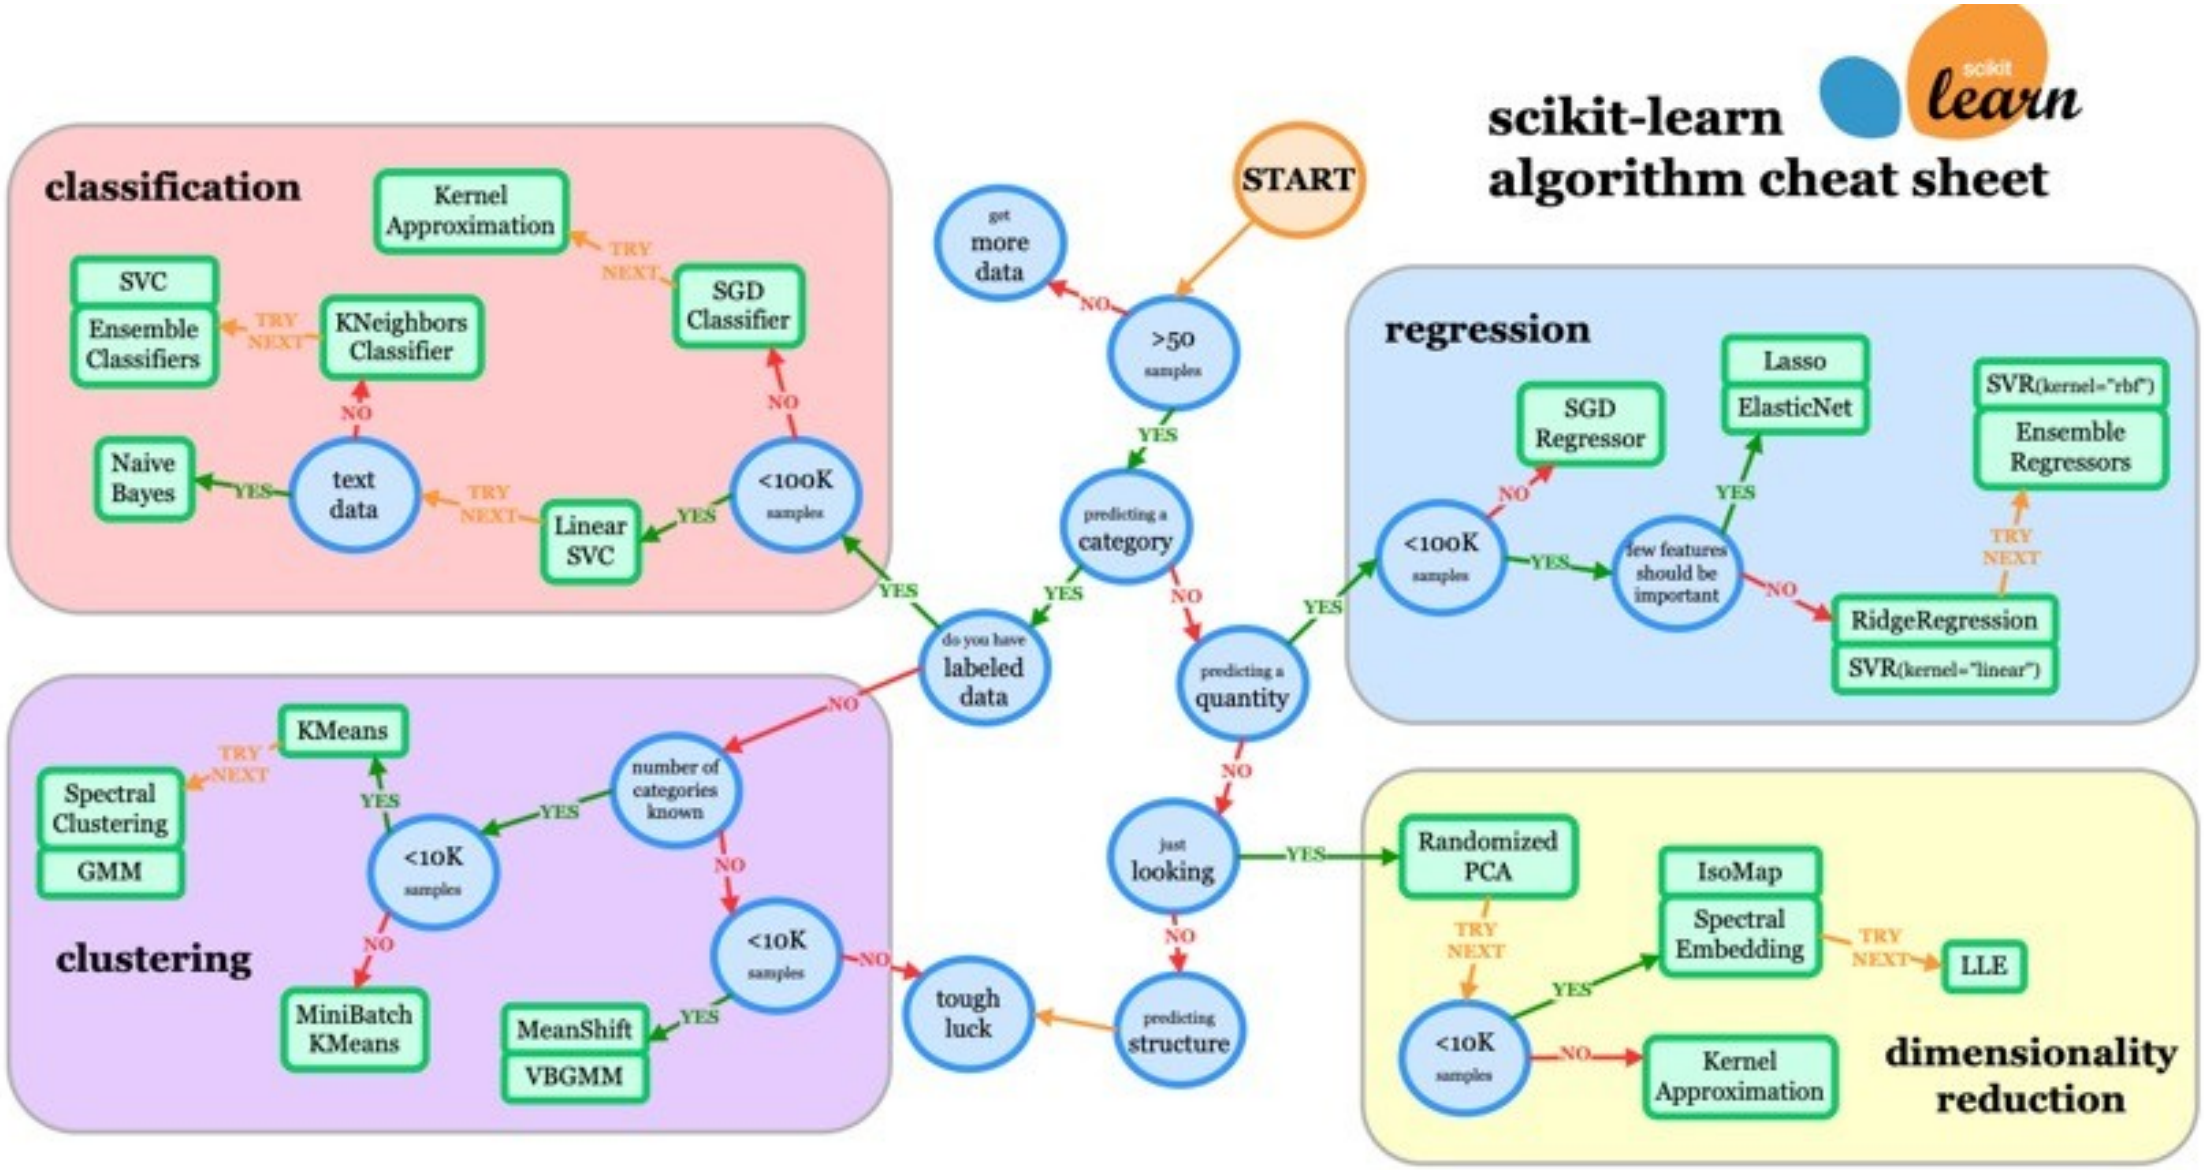
$||x||_1$ Manhattan distance

- Elastic net regression

$$\min_w \frac{1}{2n_{\text{samples}}}||Xw - y||_2^2 + \alpha\rho||w||_1 + \frac{\alpha(1-\rho)}{2}||w||_2^2$$

Tradeoff between Lasso and Ridge

# Use the cheat sheet if you want to try a more complex model!



# Is your model overfitting, underfitting, or performing well?

- You and your partner should discuss with one other pair in your company
- If your model is overfitting or underfitting, decide on a technique to try to improve your model
- If your model is performing well, try to make it even better with a small change in either direction, and be ready to defend its performance!
- Help each other!
- Ask a TA if you are unsure how your model is doing

# Final push on final project!

- Final notebooks due on **Tuesday MIDNIGHT May 5**
  - You should try to make AT LEAST ONE improvement to your model to lower errors before predicting on test data
    - If your model performs poorly no matter what you try, you should explain in conclusions/future directions why you think that is.
  - Final predictions on test data reported with error estimates/metrics
  - Explore the literature to see if you can externally validate your model results (i.e. does the literature draw the same conclusions as you?)
  - What do your model predictions mean for answering your project question?
  - Ethics and conclusions section of the notebook
- KAMEN: Presentations on **Tuesday 2PM**