

In this worksheet you'll practice fitting and evaluating different predictive models using root mean square error (RMSE). You might potentially see overfitting occurring before your very eyes! Today's data is available in the Files tab on bCourses. You may load it in as `original`.

1. Visualize the relationship between x and y with `ggplot2` code. Sketch the plot below.
2. Comment on what you see. Specifically, describe: the strength, direction, and shape/form of the association.
3. Split the original data into two data sets: `train` and `test`. *Seventy percent* of the data should be allocated to `train`.
4. Fit a linear model between x and y and save it into `m1`. Write the code you used below.
5. Write a `dplyr` pipeline starting with `train` to find the *training* RMSE for `m1`. The `fitted()` function will be helpful for you here!

6. Now, write a `dplyr` pipeline starting with `test` to find the testing *RMSE* for `m1`. *Hint: now we are truly making predictions rather than finding fitted values! What function should we replace `fitted()` with in our pipeline?*
7. Fit a second-degree polynomial predicting `y` with `x` and save it into `m2`. Copy the code you used below.
8. Using similar code from before, report the training and testing *RMSE*'s for this second-degree model in the space below. You do not need to copy your code down.
9. Fill out the table shown below. You will fit a model with a polynomial having the degree specified, and report the testing and training *RMSE* in each case. You should already be able to complete the first two rows!

Degree	Training RMSE	Testing RMSE
1		
2		
3		
4		
5		
10		
20		
25		

10. In at least three sentences, describe the pattern in the results you see and describe where overfitting is taking place.