

Homework 5: Recipes for Recipe Classification

Due October 19th at 10pm

Part 1: Recipe Regression

In this assignment, you will build classifiers for predicting the course category of a recipe. I have given you a dataset of recipes scraped from Yummly by Min et al. (2016).

We are going to try to predict which course a recipe belongs in. In the original dataset, there are 12 course categories. To make this task easier, we will focus on predicting only the first 6 categories:

- 0: Main Dishes
- 1: Desserts
- 2: Beverages
- 3: Soups
- 4: Salads
- 5: Condiments and Sauces

I have filtered the dataset to include only recipes in these 6 categories, and split it into training and test files: **recipes_train.csv** and **recipes_test.csv**. Each recipe has the following features:

- file_id: the JSON file id, in case you want to get other features from the original recipe
- name: the recipe name
- rating: the average rating the recipe received
- servings: the number of servings the recipe makes
- sugar: how much sugar is in the recipe
- fat: how much fat is in the recipe
- protein: how much protein is in the recipe
- cuisine: a string consisting of all cuisine tags associated with the recipe
- cuisine_number_1: a numerical code for the recipe's first cuisine tag
- cuisine_number_2: a code for the recipe's second cuisine tag if there is one, otherwise -1
- cuisine_number_3: a code for the recipe's third cuisine tag if there is one, otherwise -1
- id: the recipe's original id
- time: how long the recipe takes to make
- category: the recipe's course category as a string
- category_number: a numerical code for the course category

Your goal is to write a multinomial regression classifier that uses a combination of a recipe's features to predict its category.

Task 1: Setting Up Your Regression Classifier (10 points)

To build your regression classifier, you should use the PyTorch library like we did in our class cat/dog classifier example. You may find it helpful to refer back to this file.

You might also find this PyTorch tutorial helpful.

I have given you a file called **yum_regression_starter.py**. It already contains code to read data from a recipe dataset file and process it into a Pandas dataframe.

You may need to install some packages to run this script. It has been tested with Python 3.12, PyTorch 2.5, and Pandas 2.2. **Let me know ASAP if you have issues at this step!**

Your task is to fill in the **RecipeClassifier** class to create a multinomial regression model. Your model should be very similar to the one we built in class for cat/dog classification.

Test out your model by training it with just one or two features as input.

You can specify the features to include by listing column names in `selected_features`:

```
selected_features = ["rating", "time", "servings"]
```

Important: You should **NOT** change how `train_labels` or `test_labels` are defined.

After each epoch, `evaluate` reports loss and accuracy on both datasets using the current model.

Record which features you tried, and how well your model did. You should record its loss and accuracy on both the train and validation sets.

Task 2: Understanding the Regression Equation (10 points)

Here is a set of model weights and biases:

Feature	0(Mains)	1(Desserts)	2(Beverages)	3(Soups)	4(Salads)	5(Sauces)
rating	0.035	-0.23	-0.077	-0.06	-0.10	0.04
time	0.3	0.3	0.3	0.3	0.3	0.3
servings	0.177	0.25	-0.023	0.13	0.18	-0.019
bias	0.45	0.094	-0.15	-0.13	0.08	-0.42

A. Given these model parameters, calculate the pre-softmax score for Recipe 24841, Roasted Eggplant Cannelloni, in Category 0 (Mains).

B. Given these model parameters, calculate the post-softmax probability for Recipe 24841, Roasted Eggplant Cannelloni, in Category 0 (Mains).

Task 3: Adding More Features (5 points)

When your model is working, try it with all the numerical features that I have given you: rating, time, servings, sugar, fat, protein, `cuisine_number_1`, `cuisine_number_2`, `cuisine_number_3`.

Your model should achieve around 60% accuracy on the training and test sets.

Record your model's loss and accuracy on both the train and validation sets.

Part 2: Feature Selection

Now that our basic model is working, it is time to think about our features. Which of these features are the most useful? In this part, you will improve your classifier by picking better features.

Save your work for this part of the homework by making a copy of your regression file. Call it `yum_regression_improved.py`.

Task 4: Interpreting Model Weights (10 points)

I have given you two functions to help you think about which features are most useful for the model. One is the **`stats_by_label` function**, which prints out descriptive statistics for each feature by category label. This can help you get an intuitive sense of which features you expect to be best at distinguishing among categories.

The second is the **`print_coefficients` function**. This function prints out a table of regression weights for a given model. The columns are categories; the rows are features.

A negative weight for a particular feature/category pair means that as the value of the feature increases, the likelihood of the category decreases. A positive weight indicates a correlation (positive linear relationship) between the feature and the category.

Identify 3 trends you see in the regression weights. For each observed trend, say whether you think the descriptive statistics support this weighting, go against it, or provide no evidence.

Task 5: Evaluation Metrics (10 points)

So far we have looked only at overall model performance. But our recipe collection suffers from **class imbalance**: there are way more recipes in the main course category than in the others.

Write a function called `print_performance_by_class` that prints a summary of accuracy by category. Your function should take two parameters: a set of labels (such as `test_labels`) and a set of predictions derived from calling `model.predict()` on the corresponding dataset.

Your function should print a summary like the one below:

Accuracy by Category:

Category 0 : 0.913

Category 1 : 0.000

Category 2 : 0.000

Category 3 : 0.161

Category 4 : 0.000

Category 5 : 0.000

Task 6: Feature Selection (10 points)

You may find that your model does well only on certain classes, as in the example output that I have shown above. Your next task is to do feature engineering to improve your model's performance.

Design some new features for your model. For full credit, your model must achieve 75% accuracy on the test set with non-zero accuracy on at least four classes.

There are several different approaches that you can take. Here are some suggestions:

- The recipe title is likely to be informative for this task. You could write functions to assign features based on strings that appear in recipe titles.
- The cuisine categories that I have given you are not very helpful, because the model can't assign different weights to different cuisines. You could add additional features to make better use of this information.

Your new features **should NOT make use of the category or category_number features**. This would let your model cheat by peeking at the answers!

Here is an example of adding a new feature (whether the recipe title contains "strawberry"):

```
def add_strawberry_column(df):
    names = df["name"]
    has_strawberry = [ ]
    for name in names:
        if "strawberry" in name.lower():
            has_strawberry.append(1)
        else:
            has_strawberry.append(0)
    df["strawberry"] = has_strawberry
    return df
train_data = add_strawberry_column(train_data)
test_data = add_strawberry_column(test_data)
```

Make sure to add the new column name ("strawberry") to `selected_features` as well.

As you add features, you can experiment with training the model for more epochs. If you train it too long, however, you may **overfit** to the training data. Monitor the test loss; if it increases for several epochs, your model may be overfitting.

Record which features you added, how many epochs you trained for, and your model's best performance. As before, record your loss and accuracy on both the train and validation sets. **You should also write a discussion of how these features impacted accuracy by recipe class.**

Part 3: *AI Needs You* reflections

Reading Reflection Question 1

Harding writes that Warnock's background as a philosopher rather than a scientist made her desirable as a committee chair, since a specialist on the topic of IVF might be considered too close to

the subject and unable to judge public opinion as clearly. Moreover, the committee of sixteen was carefully designed to have a “broad-based” membership (p. 93).

Imagine that you are in charge of assembling a committee to consider government regulation of future AI developments, for the good of society. **Who would you select?** You need not nominate actual individuals, but you should describe how you would select the sixteen members and the qualifications you would consider.

Reading Reflection Question 2

The Warnock Committee saw the value in having a “even playing field, a clear line, indisputable” that “everyone could not only understand, but could follow” (p. 97). Spend some time thinking about what this might look like when it comes to regulating future AI developments. **What factors would you propose to consider for distinguishing ethical and unethical AI development?**