

CS273P Project 1  
Machine Learning & Data Mining: Spring 2018  
**Due: Tuesday May 29, 2018**

In this project you will work with a dataset of satellite-based meteorological measurements used to predict rainfall at a particular location. The dataset<sup>1</sup> has been pre-processed to extract features corresponding to a model actively used for predicting rainfall across the globe. Each data point corresponds to a particular lat-long location where the model thinks there might be rain. The extracted features include information such as IR (cloud) temperature at that location, and information about the corresponding cloud (area, average temperature, etc.). The target value is the amount of rainfall at the particular location..

### **Project Details.**

This course project will consist of groups of three students working together. We will use a class Kaggle competition <https://www.kaggle.com/c/uci-s2018-cs273p-1> to manage performance evaluation.

### **Step 1: Form teams.**

Please form teams of 3 students who share your interests and with whom you will directly collaborate. Piazza can help with locating other students in need of joining at team.

There is no barrier to collaboration on this project, so feel free to talk to other teams about what they are doing, how they are doing it, and how well it works (and you can see this last point on the leaderboard as well). A significant component of good performance can boil down to feature design and choices, so you may want to talk to other teams about their choices and how it affects performance. You are also free to use online code, or other sources. However, please make sure that your own entries are developed by your team members – the purpose of this project is to give you practical experience, so it will not be helpful if you just follow instructions from a friend.

### **Step 2: Download the data.**

Download the following files from the Kaggle website -

X\_train.txt - training data feature values  
Y\_train.txt - target values for training data  
X\_test.txt - test data feature values

You will train your models using X\_train and Y\_train, make predictions  $\hat{Y}$  using X\_test, and upload your predictions  $\hat{Y}$  to Kaggle. Kaggle will then score your predictions, and report your performance on a subset of X\_test, used for placing your team on the current leaderboard (the “leaderboard” data). After the competition, the score on the remainder of the data (X\_test) will be used to determine your final standing; this ensures that your scores are not affected by overfitting to the leaderboard data.

Your submission has to be a CSV file containing two columns separated by a comma. The first column should be the instance number, followed by the second column that is the predicted value

---

<sup>1</sup>The dataset is courtesy of Prof. Alex Ihler’s previous 273A course offerings and UC Irvine’s Center for Hydrometeorology and Remote Sensing.

for this instance. Further, the first line of the file should be “ID,Prediction”, i.e. the name of the two columns.

### Step 3: Choose techniques to explore.

You are required to employ at least two (2) different types of learners. Suggestions include

- **K-Nearest Neighbors.** KNN models for this data will need to overcome two issues: the large number of training & test examples, and the data dimension. Distance-based methods often do not work well in high dimensions, so you may need to perform some kind of feature selection process to decide which features are most important. Also, computing distances between all pairs of training and test instances may be too slow; you may need to reduce the number of training examples somehow (for example by clustering), or use more efficient algorithms to find nearest neighbors. Finally, the right “distance” for prediction may not be Euclidean in the original feature scaling (these are raw numbers); you may want to experiment with scaling features differently.
- **Neural networks.** The key to learning a good NN model on these data will be to ensure that your training algorithm does not become trapped in poor local optima. You should monitor its performance across backpropagation iterations on training/validation data, and verify that predictive performance improves to reasonable values. Start with few layers (2-3) and moderate numbers of hidden nodes (100-1000) per layer, and verify improvements over baseline linear models.
- **Random forests.** Learn a collection of models and combine them to get a more accurate and stable prediction.
- **Boosted learners.** Use AdaBoost, gradient boosting, or another boosting algorithm to train a boosted ensemble of some base learner.
- **Other.** You may pick any other algorithm/method of your choosing.

You may use any platform, environment, language, library or code you choose, as long as you list any any third party software/library/tools you used in your report. Or you may implement the algorithms yourself (although I caution against it).

However, a key point is that you must *explore* your approach(es); so you must do more than simply download a publicly available package and run it with default settings, or trying a few values for regularization or other basic parameter settings. You must at least explore the method fully enough to understand how changes might affect its performance, verify that your findings make sense, and then use your findings to optimize your performance; ideally, you should repeat this process several times, exploring several different ideas. In your report, you should describe why you decided to explore this aspect, what you expected to find, and how your findings matched (or didn't match) your expectations.

### Step 3: Build your learners, Evaluate, Repeat.

Using the techniques you chose to focus on, construct predictive models for your target(s). For each learner, you should do enough work to make sure that it achieves “reasonable” performance. Then, take your best learned models, and combine them using a blending or stacking technique. This could be done via a simple average/vote, or a weighted vote based on another learning algorithm. Feel free to experiment and see what performance gains are possible.

Be aware of the positive and negative aspects of the learners we have discussed. For example, nearest neighbor methods can be very powerful, but can also be very slow for large data sets; a similar statement applies to dual-form SVMs. For such learners, dealing with the large data set may be a significant issue – perhaps you could reduce the data in some way without sacrificing performance? On the other hand, linear methods are very fast, but may not have enough model complexity to provide a good fit. For such learners, you may need to try to generate better features, etc.

Output your predictions and evaluate them on Kaggle:

```
fh = open('predictions.csv', 'w')    # open file for upload
fh.write('ID,Prediction\n')         # output header line
for i,yi in enumerate(Yte):
    fh.write('{}{}\n'.format(i,yi))  # output each prediction
fh.close()                           # close the file
```

You can check the leaderboard to see your “test” performance.

**NOTE:** You’re limited to a few ( $\approx 2$ ) submissions per day to avoid over-loading their servers and to enforce good practices. Thus, you should not try to upload every possible model with every possible parameter setting – use validation data, or cross-validation, to assess *which* models are worth uploading, and just use the uploads to verify your performance on those.

## Step 5: Write it up

Your team will produce a single write-up document (one report per team), approximately **6 pages** long, describing the problem you chose to tackle and the methods you used to address it, including which model(s) you tried, how you trained them, how you selected any parameters they might require, and how they performed in on the test data. Consider including tables of performance of different approaches, or plots of performance used to perform model selection (i.e., parameters that control complexity).

Within your document, please try to describe to the best of your ability who was responsible for which aspects (which learners, etc.), and how the team as a whole put the ideas together.

You are free to collaborate with other teams, *including* sharing ideas and even code, but please document where your predictions came from. (This also relaxes the proscription from posting code or asking for code help on Piazza, at least for project purposes.) For example, for any code you use, please say in your report who wrote the code and how it was applied (who determined the parameter settings and how, etc.) Collaboration is particularly true for learning ensembles of predictors: your teams may each supply a set of predictors, and then collaborate to learn an ensemble from the set.

## Requirements / Grading

I am looking for several elements to be present in any good project. These are:

- (a) Exploration of techniques/elements on which we did not spend significant time in class. For example, using neural networks, or random forests are great ideas; but if you do this, you should explore in some depth the various options available to you for parameterizing the model, controlling complexity, etc. (This should involve more than simply varying a parameter and showing a plot of results.) Other options might include feature design, or optimizing your models to deal with special aspects of the data (e.g. possible outlier data; etc.). Your report should describe what aspects you chose to focus on.

- (b) Performance validation. You should practice good form and use validation or cross-validation to assess your models' performance, do model selection, combine models, etc. You should *not* simply upload hundreds of different predictors to the website to see how they do. Think of the website as “test” performance – in practice, you would only be able to measure this once you go live.
- (c) Adaptation to under- and over-fitting. Machine learning is not very “one size fits all” – it is impossible to know for sure what model to choose, what features to give it, or how to set the parameters until you see how it does on the data. Therefore, much of machine learning revolves around assessing performance (e.g., is my poor performance due to underfitting, or overfitting?) and deciding how to modify your techniques in response. Your report should describe how, during your process, you decided how to adapt your models and why.

Your project grade will be based on the quality of your written report, and groups whose final prediction accuracy is mediocre may still receive a high grade, if their work and results are described and analyzed carefully. But, some additional bonus points will also be given to the teams at the top of the leaderboard.