

Final Project

TA email: ml_ta@csie.ntu.edu.tw

RELEASE DATE: 04/16/2019

COMPETITION END DATE: **06/17/2019 NOON ONLINE**

REPORT DUE DATE: **06/25/2019 NOON ONLINE**

Unless granted by the instructor in advance, no late submissions will be allowed.

Any form of cheating, lying, or plagiarism will not be tolerated. Students can get zero scores and/or fail the class and/or be kicked out of school and/or receive other punishments for those kinds of misconducts.

You should write your solutions in English or Traditional Chinese with the common math notations introduced in class or in the problems. We do not accept solutions written in any other languages.

Introduction

In this final project, you are going to be part of an exciting machine learning competition. A team of scientists tries to understand the behaviors of diffusion trajectories in a periodic two-dimensional mesh potential. The trajectories are typically characterized by a model called fractional Brownian Motion (fBM), which is a type of random walk with memory. That is, the step size of the next motion depends on the history of previous steps. Two attributes, namely “mesh size” and “penetration rate”, describe the property of the mesh potential, and another attribute, namely “alpha”, represents the amplitude of the memory within fBM.

The goal of the scientific discovery is to infer the attributes from the observed trajectory (of 100,000 steps). To facilitate the inference task, the scientists have built a simulator that can generate possible trajectories for any given attributes (“alpha”, “mesh size”, “penetration rate”). The simulator is then used to generate trajectories from many different attribute combinations. After generating the trajectories, 10,000 features are extracted from each trajectory. So simply speaking, the task is to infer the three attributes from the 10,000 features.

Now, having collected some data from the simulator, the team of scientists decides to ask you, a professional machine learning researcher, to help with the scientific discovery. You need to fight for the most accurate inference on the score board. Then, you need to submit a comprehensive report that describes not only the recommended approaches, but also the reasoning behind your recommendations. Promising ideas can help you not only win the competition, but also win opportunity to work further with those scientists. Well, let’s get started!

Data Set

The problem is formalized as a three-output regression problem, where the goal is to predict the attribute “truth” of each trajectory accurately. We will have two tracks of competition. The details of the tracks, which would differ by evaluation criteria (i.e. error functions), will be announced later. The data will be divided to the training set and the test set. For the test set, the attribute “truth” will be hidden.

The data set is provided by Dr. Chia-Lung Hsieh’s team at Academia Sinica. We thank the team for helping with this competition.

Survey Report

You are asked by the board to study at least THREE machine learning approaches using the training set above. Then, you should make a comparison of those approaches according to some different perspectives, such as efficiency, scalability, popularity, and interpretability. In addition, you need to recommend THE BEST ONE of those approaches as your final recommendation **for each track** and provide the “cons and pros” of the choice.

The survey report should be no more than SIX A4 pages with readable font sizes. The most important criterion for evaluating your report is replicability. Thus, in addition to the outlines above, you should also describe how you pre-process your data, such as the features you build; introduce the approaches

you tried and provide specific references, especially for those approaches that we didn't cover in class; list your experimental settings and the parameters you used (or chose) clearly. Other criteria for evaluating your survey report would include, but are not limited to, clarity, strength of your reasoning, "correctness" in using machine learning techniques, the work loads of team members, and properness of citations.

Our sincere suggestion: *Think of your TAs as your boss who wants to be convinced by your report.*

For grading purposes, a minor but required part in your survey report for a two- or three-people team (see the rules below) is how you balance your work loads.

Competition

The submission site would be announced later. Use your submissions wisely—you *do not want to leave the TAs with a bad impression that you just want to "query" or "overfit" the test examples*. After submitting, there will be a score board showing the test error on a random half of the data set. The "hidden" test error on the other half will eventually be used to evaluate your performance.

The competition ends at noon on 06/17/2019. We'll have a mini-ceremony to honor the best team(s) on 06/18/2019. The competition site will continue to be open until the due day of the report.

Misc Rules

Report: Please upload one report per team electronically on Gradescope. You do not need to submit a hard-copy. The report is due at noon on 06/25/2019.

Teams: By default, you are asked to work as a team of size THREE. A one-person or two-people team is allowed only if you are willing to be as good as a three-people team. It is expected that all team members share balanced work loads. Any form of unfairness, such as the intention to cover other members' work, is considered a violation of the honesty policy and will cause some or all members to receive zero or negative score.

Algorithms: You can use any algorithms, regardless of whether they were taught in class.

Packages: You can use any software packages for the purpose of experiments, but please provide proper references in your report for replicability.

Source Code: You do not need to upload your source code for the final project. Nevertheless, please keep your source code until 08/01/2019 for the graders' possible inspections.

Grade: The final project is worth 400 points. That is, it is equivalent to two usual homework sets. At least 360 of them would be reserved for the report. The other 40 may depend on some minor criteria such as your competition results, your discussions on the boards, your work loads, etc..

Collaboration: The general collaboration policy applies. In addition to the competitions, we still encourage collaborations and discussions between different teams.

Data Usage: You can use only the data sets provided in class for your experiments, and you should use the data sets properly. Getting other forms of the data sets is strictly prohibited and is considered a serious violation of the honesty policy. Using any tricks to query the labels of the test set is also strictly prohibited.