

Crowdsourcing Incentives with Submodular Reward Functions

Supervisor: Aleksei Triasteyn

Keywords: machine learning, crowdsourcing, submodular functions, mobile

Description

In many crowdsourcing scenarios, participants are asked to perform certain tasks for a pre-defined reward (a good example is Amazon Mechanical Turk [1]). However, when the quality of work is important (e.g. data labeling tasks), the uniform reward scheme may not be the best solution as it doesn't elicit higher quality but rather higher quantity. Therefore, we are going to investigate reward schemes based on submodular functions [2] that approximate the quality of user inputs with regards to future learning performance.

In this project, a student will create a mobile app, potentially in a form of a game, where users would provide (or select) training examples for a certain learning problem (e.g. draw characters for handwritten character recognition), and will be rewarded according to the "teaching value" of their examples calculated by a submodular reward function. The app will have to provide instructions and indicate rewards to incentivise users. The final goal is to determine whether or not a non-uniform rewarding scheme yields significant improvements in future learning performance.

Project Objectives

- Design and implement a crowdsourcing mobile app.
- Get familiar with the concept of submodular set functions.
- Choose (or design) an appropriate submodular function for the problem.
- Perform the final experiment with two user control groups.

References

- [1] <https://www.mturk.com/mturk/welcome>
- [2] A. Krause, D. Golovin, Submodular function maximization. In *Tractability: Practical Approaches to Hard Problems*, 2012.