DESCRIPTION FOR THE GENERAL PUBLIC

The goal of the project is to develop and analyze new machine learning algorithms for complex prediction problems within a theoretical framework of *online learning*. Online learning concerns *sequential prediction problems*. Online movie recommendation, predicting stock market prices, car navigation and routing in the traffic, weather prediction, playing games, deciding on the action of a robot (e.g., self-driving car) in changing environment, deciding which ads to present on a web page or which results to show in the search engine, are all examples of sequential prediction problems.

As an illustrative example, consider an online learning problem, in which at each day, the algorithm should predict whether it will rain tomorrow or not. The algorithm has access to several weather forecasting experts. Every day, each expert predicts whether it will rain or not, and the algorithm's prediction is a combination of experts' predictions. In the next day, the actual weather is observed, and the algorithm evaluates whether its prediction was correct or not. As more and more data is observed, the algorithm will adapt its decisions by putting more confidence on weather experts which are often right, and less on those which are often wrong. The goal of the algorithm is to have its predictive accuracy close to that of the expert which will turn out to be the best after all the data (states of the weather) have been revealed; and we want this to hold no matter what the experts' predictions and actual weather are! While this goal seems hard to achieve, there are algorithms, which exactly do that, by progressively concentrating their confidence on the best so far expert, while hedging their bets with some small weight put on the remaining experts.

While the problem described above might look a bit artificial, online learning methods are nowadays routinely used in machine learning, and have already been successfully applied to a diverse range of practical problems. Impressive results include: scalable (billions of training examples described by trillions of features) learning methods, state-of-the-art data compression software, recommender systems, online advertisements, ranking of search query results, gambling, competitive stock market strategies, and many others. Thus, not surprisingly, these algorithms are present in most major machine learning software packages.

There are three main concepts which constitute an online learning problem: i) the loss function, defined for all possible predictions and outcomes, ii) the structure of the data and the parameters of the algorithm, and iii) the set of competing strategies. The first concept determines the set of actions (predictions) the algorithm can take and the function which evaluates the quality of prediction after the true outcome is revealed. For instance, in the problem of weather forecasting, the algorithm predicts weather it rains or not, and after the true state of the nature (weather) is revealed, the algorithm suffers a unit of loss if its prediction was incorrect, and zero loss otherwise. The second concept specifies the structure of the data and the corresponding set of parameters. In our example, the data are experts' predictions, and the parameter of the algorithm is the amount of confidence put on each expert (which is formally a probability distribution). Finally, the third concept specifies the class of strategies, to which the performance of our algorithm will be compared. In weather forecasting example, this would be just a set of weather experts which predictions the algorithm uses.

Within the scope of this project, we plan to pursue research on learning scenarios in which at least one of the above constituents – the loss function, the data and parameters, or the reference set of strategies – leads to challenging research problems. Results of this project will lead to improvement in efficiency of existing algorithms, as well as developing novel and original online learning methods for problems not sufficiently explored in the online learning theory.