Undergraduate Research/Master's project

Relative Performance Guidance

Synopsis

Much of state-of-the-art AI relies on large amounts of computation to build models, perform search, or do inference. In order to get the best performance, it is necessary to be able to automatically tune the parameters of such software, which requires building models of their performance. This research field is called meta-heuristics – a layer on top of the actual AI algorithms that enables us to use them better.

Meta-algorithmics involves running large-scale computational experiments to measure the behavior of algorithms in practice. However, in many cases, we cannot measure the time an algorithm takes to solve a problem exactly because it would take too long – we only know that it would take longer than a timeout.

You will investigate how to take such relative guidance ("I cannot tell you how long the run took, but it took longer than 300 seconds.") into account when building meta-algorithmic models. This will be based on recent research results in machine learning that develop a general framework to incorporate such relative guidance into the learning process. You will

- run experiments that measure the runtime of algorithms on problems;
- build machine learning models to predict the runtime of these algorithms based on existing relative guidance machine learning code;
- analyze whether and to what extent such models are better than standard meta-algorithmic approaches.

What you should bring to the project

You should be self-motivated and able to work independently, have strong programming and analytical skills, some experience running computational experiments (e.g. the final assignment in COSC 3020), and basic machine learning knowledge. Experience with a Linux environment and using large-scale computational resources such as Teton is not required, but a plus.

What you will get out of it

You will become familiar with using large-scale computational infrastructure and building machine learning models beyond using standard toolkits like scikit-learn. Depending on the obtained results, this project may lead to a scientific publication. This is a project with lots of details to be defined as part of it – you can bring in your own ideas and make it your own.

Interested? Talk to Lars Kotthoff <larsko@uwyo.edu>. Supported by NSF grant #1813537 and an REU supplement.