

# Undergraduate Research/Master's project

## Quicksort Tuning

### Synopsis

Most standard library sorting implementations use some version of quicksort, which is usually combined with insertion sort for small lists to avoid some of the overhead of repeated partitioning and recursive calls. This improves performance in practice, but the threshold size of when to switch from quicksort to insertion sort is hard-coded.

You will investigate whether this static choice is best across different workloads and machine types. You will then apply state-of-the-art AI and machine learning techniques to tune this threshold for a given workload on a given machine. This involves

- running experiments with different threshold values for different workloads on different machines;
- using techniques from statistics to analyze the results of these experiments;
- applying parameter-tuning techniques to dynamically find the best threshold in a given scenario.

### 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). You should have basic familiarity with statistical techniques. Experience with a Linux environment and using large-scale computational resources such as Mount Moran is not required, but a plus.

### What you will get out of it

You will become familiar with state-of-the-art methods for parameter tuning in artificial intelligence and machine learning, and learn how to use large-scale computational infrastructure. You will also develop familiarity with statistical and data science techniques, which are highly sought-after in industry. 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>.