



LA-UR-24-28098

Author(s): Christin Whitton

Mentors: Nathan DeBardleben, Vanessa Job

Title: Are we there yet? Predicting the Queue Wait Times and Job Runtimes for HPC Jobs

Abstract: This project utilizes historical data from the Grizzly cluster, a medium-sized supercomputer here at LANL, to predict both the job runtime (what the user roughly predicts when submitting their job with their “wallclock limit”) and to predict the minutes a job will wait in the queue. SLURM data from 2018 and 2022 were used to build models to predict these variables. Multiple models built with different machine-learning algorithms, including Random Forest and Adaboost, were compared. Using the built models, we show that machine learning algorithms offer an improved estimate over the user-predicted runtime and over a simulated queue wait-time.