Benchmarking BoTorch for Hyperparameter Optimization for Reinforcement Learning

> Project Presentation - Luou Wen lw658

BoTorch

- New framework for Bayesian Optimisation research based on PyTorch and GPyTorch
- Algorithms are designed to exploit auto-differentiation and hardware acceleration features of PyTorch
- Implementations of Monte-Carlo acquisition functions in addition to standard acquisition functions
- Sample average approximation for optimising MC acquisition functions
- Compared with existing frameworks on tasks including Hyperparameter Optimization

Hyperparameter Optimization Comparisons

- Tuning 5 parameters of a deep Q-network learning algorithm on the Cartpole task
- Tuning 6 parameter network network surrogate model for the UCI Adult data set
- Tuning 3 parameters of the Stochastic Weight Averaging procedure on the VGG-16 architecture
- All compared less than 10 parameters while Gaussian Processes can handle up to about 30 and more structured models can go up to a few hundred dimensions

Proposed project

- Aim: compare and evaluate performance for higher dimensional problems in reinforcement learning
- Benchmark BoTorch, Cornell-MOE, and GPyOpt acquisition functions on higher dimensional tasks
- Start with Inverted Pendulum Swing-up
- and possibly expand to locomotion tasks
- Compare the change in best observed function values over number of observations

(
	1

Project plan

- Initialising
 - Get all frameworks installed and running
- Build
 - Replicate implementations of acquisition functions described in the paper
- Benchmark
 - Find / implement RL benchmarking tasks starting with Inverted Pendulum Swingup
- Report
 - Report on findings

Questions?