

ERIC HANS LEE

(425) 241-2969, erichanslee@gmail.com, <https://erichanslee.github.io/>,
<https://github.com/erichanslee>, <https://www.linkedin.com/in/erichanslee>

Machine learning research scientist specializing in numerical optimization and decision-making under uncertainty. Three-plus years of experience engineering and implementing production machine learning and deep learning systems. Experience building graph neural network and computer vision models. Passionate about both research and engineering.

WORK EXPERIENCE

AI Research Scientist, *Intel*, 2020-present

Note: Started job at Sigopt, which was then acquired by Intel.

- Research, productionize, and/or publish methods for black-box search and optimization.
- Improve efficiency and capabilities of numerical optimization code.
- Write and release open source code and tutorials for machine learning community.
- Collaborate with other teams to incorporate better numerical optimization into their pipelines. Worked promising discovery problems to incorporate machine learning into materials and battery design.

Applied Science Research Intern, *Amazon*, Q3-Q4 2019

- Researched neural network hyperparameter tuning with constraints on tuning budget.
- Developed and implemented a successful cost-aware algorithm that tunes computer vision models faster than existing state-of-the-art and published results.

Research Intern, *SigOpt*, Q1-Q2 2019

- Researched applications of reinforcement learning to Bayesian optimization.
- Developed a method to accelerate Bayesian optimization using rollout and published results.

Research Intern, *Cray Incorporated*, Summer 2017

- Investigated algorithms for high performance matrix factorizations.
- Developed QR factorization in C and MPI 2x faster than the state-of-the-art on Cray supercomputer.

OPEN-SOURCE CONTRIBUTIONS

- **Sigopt Open Source** (<https://github.com/sigopt/sigopt-server>): a black-box optimization framework designed for ease-of-use and robustness.
- **BoTorch Tutorials** (<https://botorch.org/tutorials/>): a Bayesian optimization library built on top of PyTorch, designed for flexibility and extensibility.

EDUCATION

Cornell University, 2015-2020

M.S. and Ph.D. in Computer Science

Area: Machine Learning

Thesis: Budget-constrained Bayesian Optimization

UC Berkeley, 2011-2015

B.A. in Applied Mathematics, Magna Cum Laude

SKILLS

Machine Learning:

Machine learning, active learning, deep learning, Bayesian optimization, graph neural networks, computer vision, gradient-boosted trees.

Statistics:

Monte Carlo, statistical testing / analysis, Bayesian inference, density estimation, sampling and random number generation, uncertainty quantification.

Numerics:

Kernel methods, numerical optimization, Markov decision processes, quadrature, applied statistics, matrix computations, dynamical systems.

Programming:

Python, Numpy, Pytorch, Scipy, Scikit-learn, Matplotlib, Pandas, Jupyter, Colab, C++, C.

Distributed Computing:

AWS, GCP, Hadoop, Spark.

High Performance Computing:

Vectorization and parallelization of numeric code, OpenMP, MPI, CUDA, Slurm, Torque.

Software Engineering Tools:

Unix, Git, Pytest, Circle, Jira, CMake.

PUBLICATIONS

- A Sorokin, X Zhu, **EH Lee**, B Cheng, *SigOpt Mulch: An intelligent system for AutoML of gradient boosted trees*, Knowledge-based systems, 2023.
- X Zhu, L Huang, **EH Lee**, CA Ibrahim, D Bindel, *Bayesian Transformed Gaussian Processes*, Transactions on machine learning research, 2023.
- **EH Lee**, B Cheng, M McCourt, *Achieving Diversity in Objective Space for Sample-Efficient Search of Multiobjective Optimization Problems*, Winter Simulation Conference, 2022.
- C Gonzales, **EH Lee**, KLK Lee, J Tang, S Miret, *Hyperparameter Optimization of Graph Neural Networks for the OpenCatalyst Dataset: A Case Study*, AI for Accelerated Materials Design NeurIPS 2022 Workshop.
- **EH Lee**, D Eriksson, V Perrone, M Seeger, *A nonmyopic approach to cost-constrained Bayesian optimization*, Uncertainty in artificial intelligence, 2021.
- G Malkomes, B Cheng, **EH Lee**, M Mccourt, *Beyond the pareto efficient frontier: Constraint active search for multiobjective experimental design*, International conference on machine learning, 2021.
- **EH Lee**, V Perrone, C Archambeau, M Seeger, *Cost-aware Bayesian optimization*, International conference on machine learning, 2020.
- **EH Lee**, D Eriksson, B Cheng, M McCourt, D Bindel, *Efficient Rollout Strategies for Bayesian Optimization*, Uncertainty in artificial intelligence, 2020.
- D Eriksson, K Dong, **EH Lee**, D Bindel, AG Wilson, *Scaling Gaussian process regression with derivatives*, Neural information processing systems, 2019.
- B Kalantari, **EH Lee**, *Newton–Ellipsoid polynomiography*, Journal of mathematics and the arts, 2019.