

David Eriksson

Curriculum Vitae

✉ david.eriksson89@gmail.com
dme65.github.io/
in linkedin.com/in/davideriksson89/
github.com/dme65/
scholar.google.com/citations?user=SWQjkN4AAAAJ

Education

- 2014 – **Cornell University**
2018 *Ph.D. in Applied Mathematics*
• *Thesis*: Scalable kernel methods and their use in black-box optimization.
• TA award in Computer Science, Spring 2016.
• GPA: 4.21/4.0.
- 2012 – **Chalmers University of Technology**
2014 *M.Sc. in Engineering Mathematics and Computational Science*
• GPA: 5.0/5.0.
- 2008 – **Chalmers University of Technology**
2011 *B.Sc. in Mathematics*
• GPA: 4.92/5.0.

Research Interests

Bayesian Optimization, Gaussian Processes, AutoML, Machine Learning, Scientific Computing.

Selected Publications

- Discovering Many Diverse Solutions with Bayesian Optimization (*AISTATS 2023*, **Oral**).
- Sparse Bayesian optimization (*AISTATS 2023*).
- Bayesian Optimization over High-Dimensional Combinatorial Spaces via Dictionary-based Embeddings (*AISTATS 2023*).
- Bayesian optimization over discrete and mixed spaces via probabilistic reparameterization (*NeurIPS 2022*).
- Multi-Objective Bayesian Optimization over High-Dimensional Search Spaces (*UAI 2022*, **Oral**).
- High-Dimensional Bayesian Optimization with Sparse Axis-Aligned Subspaces (*UAI 2021*).
- Bayesian Optimization is Superior to Random Search for Machine Learning Hyperparameter Tuning: Analysis of the Black-Box Optimization Challenge 2020 (*NeurIPS 2021 Competition Proceedings*).
- Scalable Constrained Bayesian Optimization (*AISTATS 2021*).
- Fast Matrix Square Roots with Applications to Gaussian Processes and Bayesian Optimization (*NeurIPS 2020*).
- Scalable Global Optimization via Local Bayesian Optimization (*NeurIPS 2019*, **Spotlight**).
- Scaling Gaussian Process Regression with Derivatives (*NeurIPS 2018*).
- Scalable Log Determinants for Gaussian Process Kernel Learning (*NeurIPS 2017*).

Work Experience

Meta Aug. 2020 – Present
Staff Research Scientist Remote

- Tech lead focusing on using Bayesian optimization to tune the most important ML models at the company.
- Research on high-dimensional Bayesian optimization, neural architecture search, and probabilistic models.
- Main organizer of the NeurIPS 2020 black-box optimization competition (bbochallenge.com).
- Developer and open-source contributor to BoTorch and Ax.

Uber AI Labs Jan. 2019 – July 2020
Senior Research Scientist San Francisco, CA, USA

- Built and designed Uber's service for Bayesian optimization which tuned thousands of production models in 2019 and 2020.
- Lead author for our NeurIPS 2019 spotlight paper on TuRBO, a scalable high-dimensional Bayesian optimization method.
- Uber AI Quarterly Recognition Award (top 1%).
- Uber Engineering Quarterly Recognition Award (top 1%).
- Two patents on high-dimensional Bayesian optimization.

Google May. 2018 – Aug. 2018
Software Engineering Intern Mountain View, CA, USA

- Developed and implemented a high-dimensional Bayesian optimization method in Google Vizier.
- Developed an ML model for ad selection.

The MathWorks May. 2017 – Aug. 2017
Software Developer Intern Natick, MA, USA

- Designed and implemented `surrogateopt`, an asynchronous surrogate optimization framework in MATLAB.

NASA GSFC June 2011 – June. 2012, Summer 2013
Research Intern Greenbelt, MD, USA

- Research on surface loading and atmospheric ray tracing.
- Published two journal articles based on this work.

Open-Source Contributions

- BoTorch (github.com/pytorch/botorch).
- Ax (<https://github.com/facebook/Ax>).
- TuRBO (github.com/uber-research/TuRBO).
- GPyTorch (github.com/cornellius-gp/gpytorch).
- pySOT (github.com/dme65/pySOT).

Reviewing

- **Reviewer:** NeurIPS 19-22, ICML 21, AAAI 21, AISTATS 21, ICLR 21, UAI 20-21.
- **Area Chair:** AISTATS 23.