

Carson Eisenach

Homepage: <https://carson.eisenach.org>
GitHub: <https://github.com/ceisenach>
Email: carson@eisenach.org

New York, NY
(609) 651-1285

PROFILE AI researcher with a PhD in Statistics and Machine Learning. Working on theory and methodology for deep reinforcement learning as applied to supply chain optimization.

EDUCATION **Princeton University** Princeton, NJ
Ph.D., ORFE, June 2019.
M.A., ORFE, September 2016.
Dissertation: *Modern Optimization for Statistics and Learning*
Adviser: Han Liu
Coursework: *Advanced Algorithm Design, Theoretical Machine Learning, Linear and Convex Optimization, Advanced Optimization, Statistical Theory and Methods, Statistical Learning and Nonparametric Estimation, Probability Theory, Stochastic Calculus*

Williams College Williamstown, MA
B.A., Mathematics & Computer Science, Cum Laude with Honors in Mathematics, May 2014

RESEARCH INTERESTS Deep reinforcement learning, Optimization for Statistical Learning, Machine Learning.

PROFESSIONAL EXPERIENCE **Amazon** New York, NY
Senior Applied Scientist January 2020 - Present
Developed SOTA architectures for neural forecasting, used at Amazon to forecast demand for millions of products worldwide. More recently, have been working on applying reinforcement learning to various problems in supply chain including periodic review inventory control. Constrained RL and simulator fidelity are ongoing research directions.

Tencent AI Lab Bellevue, WA
Research Intern September 2017 - August 2018
Reinforcement learning theory and applications research. Led a team implementing RL algorithms, and underlying infrastructure, to solve real-time-strategy games.

Princeton University SMILE Lab Princeton, NJ
Researcher September 2015 - August 2017
High-dimensional statistics, optimization and deep learning research. Wrote papers, derived results and developed software packages implementing novel methodology.

uFora (formerly Broad Street Analytics) New York, NY
Software Engineer June - August 2012
Designed software tools for numerical analysis on large datasets.

RESEARCH EXPERIENCE **Selected Publications and Preprints**
[1] Dhruv Madeka, Kari Torkkola, Carson Eisenach, Anna Luo, Dean P. Foster, Sham M. Kakade. "Deep Inventory Management". arXiv:2210.03137, 2022.
[2] Carson Eisenach, Yagna Patel and Dhruv Madeka. "MQTransformer: Multi-Horizon Forecasts with Context Dependent and Feedback-Aware Attention". arXiv:2009.14799, 2020

- [3] Carson Eisenach, and Han Liu. “Efficient, Certifiably Optimal Clustering with Applications to Latent Variable Graphical Models”. *Mathematical Programming Series B*, 2020.
- [4] Carson Eisenach, Haichuan Yang, Ji Liu, and Han Liu. “Marginal Policy Gradients: A Unified Family of Estimators for Bounded Action Spaces with Applications”. *International Conference on Learning Representations (ICLR’19)*, New Orleans, USA, 2019.
- [5] Carson Eisenach, Florentina Bunea, Yang Ning, and Claudiu Dinicu. “Inference in Cluster-Based High Dimensional Graphical Models”. *Journal of Machine Learning Research*, 2020.
- [6] Carson Eisenach, Zhuoran Yang. “Natural Policy Gradient for Exponential Families”. Technical Report, 2018.
- [7] Carson Eisenach, Zhaoran Wang, and Han Liu. “Nonparametrically Learning Activation Functions in Deep Neural Nets”. Technical Report, 2016.

Selected Talks and Presentations

- “Topics in Multi-Agent Reinforcement Learning”. Reinforcement Learning Seminar, Lehigh University, Bethlehem, PA. November 2018.
- “Marginal Policy Gradients for Complex Control”. ICSA Symposium, New Brunswick, NJ. June 2018.
- “High Dimensional Inference for G-Block Latent Variable Graphical Models”. Cornell Day of Statistics. September 2016.

SELECTED PROJECTS

Torchkit: A Toolkit for Deep RL 2017-2019
Key features are an auto-differentiation engine for vector valued functions and tools to manage and run batches of RL experiments, built on top of PyTorch. Algorithms implemented include natural policy gradient [4], angular policy gradient [2], and TRPO. By using a factored form of the KL-divergence and directly computing derivatives of vector valued functions, fisher-vector products are computed much faster than in standard implementations. URL: <https://github.com/ceisenach/torchkit>

GFORCE: An R Package 2016-2017
Implemented an R package (written mainly in C) providing clustering algorithms and inferential procedures from [1,3]. The highlight is a novel first-order SDP solver for convex relaxations of clustering problems, as well as hypothesis testing and FDR control for latent variable graphical models. Implementations of traditional clustering methods like Lloyd’s algorithm (with kmeans++ initialization) and hierarchical clustering are also provided. URL: https://github.com/ceisenach/R_GFORCE

TEACHING EXPERIENCE

Princeton University, Assistant Instructor Princeton, NJ
 Fundamentals of Statistics (ORF 245) Fall 2015; Fall 2016; Spring 2019
 Optimization (ORF 307) Spring 2016; Spring 2017
 Time Series and Regression (ORF 405) Fall 2018

Williams College, Teaching Assistant Williamstown, MA
 Tax Policy in Emerging Markets (ECON 514) Spring 2014
 Real Analysis (MATH 305) Spring 2012
 Data Structures and Advanced Programming (CSCI 136) Fall 2011

TECHNICAL SKILLS

Languages: C (proficient), Python (proficient), R (proficient), JavaScript (proficient), TypeScript (proficient), C++ (familiar), Java (familiar), Fortran (familiar)
Frameworks and Tools: Pytorch, MXNet, MATLAB

HONORS AND AWARDS

Sigma Xi: 2014. Nominated and invited to join as a member.

Herschel-Smith Fellowship: 2014, declined. Full tuition and stipend for two years of graduate study at Cambridge University.

Morgan Prize in Applied Mathematics: 2013, 2014.

Other Awards: Class of '60 Computer Science Scholar (2012), Dean's List at Williams College (2010-2014), National Merit Scholar (2010), AP National Scholar (2010)

INTERESTS

Skiing (expert skier and former ski instructor), travel and nature photography, squash