Alex Boyd

alexjamesboyd2@gmail.com

Education

University of California, Irvine — Irvine, CA
Ph.D. Statistics — Sep 2018 - Feb 2024 *(est.)*
Advisors: Dr. Padhraic Smyth & Dr. Stephan Mandt
Overall GPA: 4.00 out of 4.00

California Polytechnic State University — San Luis Obispo, CA
B.S. Software Engineering — Sep 2014 - Jun 2018
Interdisciplinary Data Science Minor
Advisor: Dr. Dennis Sun
Overall GPA: 4.00 out of 4.00

Research Interests

I am interested developing techniques that lie at the intersection of probabilistic graphical models and deep learning that enable interpretability and flexibility when modeling complex processes (specifically involving data that is sequential in nature, such as time-series, discrete event data, and natural language).

Publications

A. Boyd, Y. Chang, S. Mandt, and P. Smyth, 'Inference for Mark-Censored Temporal Point Processes,' *Proceedings of the 39th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2023.

E. Wong-Toi, A. Boyd, V. Fortuin, and S. Mandt, 'Understanding Pathologies of Deep Heteroskedastic Regression,' *Preprint (currently under review)*, 2023

A. Boyd, Y. Chang, S. Mandt, and P. Smyth, 'Probabilistic Querying of Continuous-Time Event Sequences,' *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics (AIStats)*, 2023.

A. Boyd*, S. Showalter*, S. Mandt, and P. Smyth, 'Predictive Querying for Autoregressive Neural Sequence Models,' *Advances in Neural Information Processing Systems 35 (NeurIPS)*, 2022.

A. Alexos*, A. Boyd*, and S. Mandt, 'Structured Stochastic Gradient MCMC,' *Proceedings of the 39th International Conference on Machine Learning (ICML)*, 2022.

A. Li, A. Boyd, P. Smyth, and S. Mandt, 'Detecting and Adapting to Irregular Distribution Shifts in Bayesian Online Learning,' *Advances in Neural Information Processing Systems 34 (NeurIPS)*, 2021.

P. Putzel, H. Do, A. Boyd, H. Zhong, and P. Smyth, 'Dynamic Survival Analysis for EHR Data with Personalized Parametric Distributions,' *Proceedings of the 6th Machine Learning for Healthcare Conference (MLHC)*, 2021.

A. Boyd, P. Smyth, R. Bamler, and S. Mandt, 'User-Dependent Neural Sequence Models for Continuous-Time Event Data,' *Advances in Neural Information Processing Systems 33 (NeurIPS)*, 2020.

A. Boyd, R. Puri, M. Shoeybi, M. Patwary, and B. Catanzaro, 'Large Scale Multi-Actor Generative Dialog Modeling,' *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020.

Professional Experience

Applied Machine Learning Research Intern Cupertino, CA
Apple Jun 2021 - Sep 2021

- Researched new techniques to automate processing and tagging of information in documents (Document AI) to improve efficiency of laborious business processes
- Leveraged few-shot learning approaches (e.g., prototypical networks) to allow for Document AI to be used across many situations with few ground-truth labels

Deep Learning Research Intern for Program Synthesis Redmond, WA
Microsoft Research Jun 2020 - Sep 2020

- Researched new methods for contextual time extraction from natural language utterances (emails) via program synthesis to help automatically schedule meetings
- Designed a domain specific programming language and associated interpreter to accurately describe complex sets of time intervals
- Created a model training regime that utilized both direct supervision with teacher forcing and indirect supervision with REINFORCE and a novel reward based on Wasserstein distance

Applied Deep Learning Research Intern Santa Clara, CA
NVIDIA Jun 2019 - Jan 2020

- Researched new dialog modeling techniques with transformers for the purposes of better response generation, longer conversational contexts, and more fine tuned control of an agent's "personality"
- Extracted millions of multi-turn conversations with millions of different users from Reddit to ensure natural sounding training material as well as a diverse amount of talking points

Machine Learning Intern Pleasanton, CA
Workday Jun 2018 - Sep 2018

- Researched neural network architectures for a natural language interface to databases to support non-technical users in generating business intelligence reports
- Improved upon state-of-the-art deep learning models by utilizing OpenAI's transformer language model as preprocessing for a hierarchical recurrent pointer network
- Augmented previously developed model to operate as an agent in a graphical system, thus allowing intent parsing, parameter assignment, and action execution in a conversational environment
- Engineered both use cases to be served in a stateless web application powered by the AWS ecosystem for scaling purposes when the research transitions from ideation into production

Data Science Intern Pleasanton, CA
Workday Jun 2017 - Sep 2017

- Employed custom Markov Chain recognition algorithm on unlabeled product usage data to uncover underlying process model definitions on a per tenant basis
- Analyzed users through multilevel clustering involving hierarchical clustering and probabilistic mixture models on raw usage patterns to identify key individuals that trigger large amounts of product consumption
- Implemented and designed a deep recurrent neural network to monitor raw log data to predict the future load of the main product within the next four hours for preventative server allocation

R Shiny App Developer Remote Locations
Contracted by Dr. Roxy Peck Mar 2017 - Sep 2017

Software Engineering Intern Goleta, CA
Toyon Inc. Jun 2016 - Sep 2016

Teaching Experience

Head Teaching Assistant for Intro to Machine Learning Irvine, CA
School of ICS - UCI Sep 2023 - Dec 2023

Teaching Assistant for Intro to Statistics Irvine, CA
School of ICS - UCI Sep 2018 - Dec 2018

Supplementary Mathematics Workshop Leader San Luis Obispo, CA
Student Academic Services - Cal Poly Jan 2016 - Jun 2018

Introduction to Data Science Teaching Assistant San Luis Obispo, CA
Statistics Department - Cal Poly Jan 2017 - Mar 2017

Study Session Leader San Luis Obispo, CA
Student Academic Serivces - Cal Poly Jan 2015 - Jan 2016