

EDUCATION

- Stanford University [GPA: 4.20/4.30, Advisor: Dr. Stephen Boyd] *Sept. 2015 – June. 2020 (expected)*
Ph.D. in Institute for Computational & Mathematical Engineering (ICME)
Main Interest: Optimization, Bayesian Learning, Data Mining, Deep/Machine Learning
- Peking University (PKU) [GPA: 3.82/4.00, Advisor: Dr. Zaiwen Wen] *Sept. 2011 – July. 2015*
B.S., School of Mathematical Sciences (SMS), Final Ranking: 1/181
Achievements: First Prize of Challenge Cup in PKU, 2015
Innovation Award in PKU (top 1%), 2014

PUBLICATIONS & PREPRINTS

- X. Guo, A. Hu, R. Xu and J. Zhang (alphabetical). *Arxiv preprint, submitted*
“Learning Mean-Field Games,”
- Presented a generalized mean-field game (GMFG) framework for simultaneous learning and decision making in stochastic games with entangled states and actions and a large population
 - Established the existence and uniqueness of a Nash equilibrium in GMFG under standard assumptions in the literature
 - Proposed a Q-learning based algorithm (GMF-Q) with carefully designed stabilization techniques, and derived its global convergence and complexity
 - Demonstrated the efficiency and robustness of GMF-Q compared to existing algorithms in multi-agent reinforcement learning on repeated Ad auction problems
- X. Guo, A. Hu, R. Xu and J. Zhang (alphabetical). *Arxiv preprint, submitted*
“On the Consistency and Computation of MLEs for Multivariate Hawkes Processes,”
Short version appeared in NeuIPS 2018 Workshop on Causal Learning
- Proved the first consistency result for the maximum likelihood estimators (MLEs) of general linear multivariate Hawkes processes (MHPs) under a short list of verifiable assumptions
 - Generalized the results to encompass regularization, approximation and missing data scenarios that have been widely encountered in practice without theoretical guidance
 - Proposed an Anderson accelerated alternating minimization algorithm for computing the MLEs of MHPs with global convergence to stationary points (which is also the first globally convergent Anderson accelerated algorithm for generic non-convex optimization)
 - Achieved the state-of-the-art performance on both synthetic and real world datasets
- J. Zhang, B. O’Donoghue, S. Boyd. *Arxiv preprint, submitted*
“Globally Convergent Type-I Anderson Acceleration for Non-Smooth Fixed-Point Iterations,”
- Proposed the first globally convergent variant of Anderson acceleration for general non-expansive non-smooth fixed-point problems by combining novel online pre-conditioning strategies with safe-guarding steps
 - Conducted extensive numerical experiments to show that many first order algorithms for convex optimization (ranging from proximal gradient descent, ADMM to value iteration) can be largely and consistently improved with the proposed algorithm
 - The algorithm is partly implemented in SCS 2.0, one of the default solvers used in the convex optimization parser-solver CVXPY 1.0
- A. Zanette*, J. Zhang*, M. Kochenderfer (*equal contribution). *European Conference on Machine Learning (ECML), 2018*
“Robust Super-level Set Estimation using Gaussian Process,”
- Proposed a global algorithm for active super-level set estimation via Gaussian Process
 - Balanced accuracy and proactivity to achieve robust performance with efficiency
 - Extensive numerical tests outperform the state-of-the-art (e.g. LSE and TruVaR)
 - Proved convergence in frequentist settings and connection to variance minimization

COURSE PROJECTS

- Mar. 2017 – May. 2017* **Ruminating Neural Networks for Sequence Modeling**
- Designed RumNet based on moving windows and auto-regressive attention units

- Achieved competitive perplexity with better parallelization efficiency for language modeling tasks compared to the state-of-the-art (e.g. QRNN and regularized LSTM)
- Implemented through TensorFlow and Microsoft Azure platform
- Final project prize winner in CS224N, 2017 Spring (3 out of 223 projects)
- Joint work with Hao Sheng and Jin Xie

Mar. 2017 – May. 2017

Information-Directed Sampling (IDS) for Reinforcement Learning (RL)

- Extended IDS for multi-armed bandits to solving general RL problems
- Proposed practical algorithms for both model-based and model-free settings
- Established a near-optimal regret bound by a novel state decomposition trick for reward learning under inverse RL settings
- Experiments show improved data efficiency compared to PSRL & UCRL2
- Joint work with Junyang Qian

INVITED TALKS

- From Extrapolation to Quasi-Newton: Stabilizing Type-I Anderson Acceleration for Line-Search Free, Memory-Efficient and Black-Box Acceleration
Operations Research Seminar, BICMR, June 14, 2019
- Learning Mean-Field Games
ByteDance AI Lab, June 13, 2019
- Globally Convergent Type-I Anderson Acceleration for Non-smooth Fixed-Point Iterations
ICME Departmental Seminar, Stanford, January 28, 2019
Large-Scale Conic Optimization, INFORMS, November 7, 2018
- Robust Bayesian Level Set Estimation via Gaussian Processes
Data Science and Deep Learning, INFORMS, November 6, 2018
- Balancing Proactiveness and Robustness in Online Super-level Set Estimation
Operations Research Seminar, BICMR, December 20, 2017

WORK EXPERIENCE

June. 2018 – Sept. 2018

Citadel LLC, AI Research (Seattle), Quantitative Research Intern

Project: Multi-Station Weather Forecast Bias Correction for Commodity Trading
Manager: Dr. Li Deng & Dr. Pusheng Zhang

June. 2017 – Sept. 2017

CloudBrain USA Inc., Deep Learning Research Intern

Project: Hybrid Neural GPU with Direct Encoder Feeding
Manager: Dr. Benyu Zhang

HONORS AND AWARDS

- Stanford Graduate Fellowship (highest honored) 2015 - 2020
- Team Silver Medal in Shing-Tung Yau Math Contest 2014
- Meritorious Award of 2014 ICM (top 15%) & 2013 MCM (top 17%)
- National Scholarship (highest honored) 2012, 2013 & 2014

PROFESSIONAL SERVICES

- Reviewer for Numerical Algorithms (NUMA), Operations Research (OR)
- Reviewer for NeuIPS (rated as the top 218 reviewers in 2018 with free registration), ICML (rated as the top 5% reviewers in 2019 with free registration), ICLR, UAI
- Session chair for *Efficient Learning and Optimization in Data Mining* in 2019 INFORMS annual meeting, Seattle
- Invited lecturer for *Convex Optimization Short Course* at 2019 International Conference on Statistical Optimization and Learning, Beijing
- Teaching assistant for CME 307 (2017-2019), EE 364A (2017) at Stanford

SKILLS

- **Programming & Software:** MATLAB, Python, R, SQL, LaTeX
- **Specialization:** Convex/Bayesian Optimization Algorithms, Data Mining, Deep learning, Algorithms for Parallel & Distributed Computation