

Jialin Song

jssong@caltech.edu
jialin-song.github.io

EDUCATION

- PhD Candidate** Caltech 2015 - Present
Program: Computing & Mathematical Sciences
GPA: 3.7 / 4.0
- Honours Bachelor of Science**, University of Toronto, St. George 2011 - 2015
Specialists: Computer Science Specialist and Mathematics Specialist
GPA: 3.95 / 4.0

RESEARCH EXPERIENCE

- Research Assistant** Sept. 2015 - Present
Advisor: Dr. Yisong Yue
- Model combinatorial optimization as sequential decision-making problems and conduct research into designing novel reinforcement learning and imitation learning algorithms to speed up solving NP-hard problems, including combinatorial auctions and risk-aware path planning. Our “learning to optimize” research has resulted in faster solvers than leading commercial solvers such as Gurobi for large scale integer programs.
 - Conduct novel research in different Bayesian optimization variants including multi-task and multi-fidelity settings. Our research has been applied successfully in material science applications in designing nanophotonic filters. We are also working on a sequence optimization problem in antibody designs.
 - Additional research topics include dueling bandits, coactive learning and adaptive submodular optimizations.
- Project in Computer Science** Sept. 2014 - June 2015
Supervisor: Dr. Faith Ellen
- Prove and implement new locking mechanisms for concurrent trees
- Summer research project at University of Toronto** May 2013 - Aug. 2013
Supervisor: Dr. Michael Brudno
- Improved medical diagnosis system for rare genetic disorders. Probabilistic graphical models and other models were used to conduct inference process which resulted in a better diagnosis system compared to previous work. The project was mainly done in Java.
- Research Opportunity Program** Sept. 2012 - Apr. 2013
Supervisor: Dr. V.Kumar Murty
- Researched on the diffusion property of cryptographic hash functions

PUBLICATIONS

- 2021: A. Alieva, A. Aceves, J. Song, S. Mayo, Y. Yue, Y. Chen, Learning to Make Decisions via Submodular Regularization, ICLR 2021, https://openreview.net/pdf?id=ac288vnG_7U
- 2021: N. Abcouwer, S. Daftry, S. Venkatraman, T. del Sesto, O. Toupet, R. Lanka, J. Song, Y. Yue, M. Ono, Machine Learning Based Path Planning for Improved Rover Navigation, IEEE Aerospace Conference 2021, <https://arxiv.org/abs/2011.06022>
- 2020: J. Song, R. Lanka, Y. Yue, B. Dilkina, A General Large Neighborhood Search Framework for Solving Integer Linear Programs, NeurIPS 2020, <https://arxiv.org/abs/2004.00422>
- 2020: S. Dai, J. Song, Y. Yue, Multi-task Bayesian Optimization via Gaussian Process Upper Confidence Bound, ICML 2020 Workshop on Real World Experiment Design and Active Learning https://realworldml.github.io/files/cr/35_Camera_Ready_RealML.pdf
- 2020: J. Song, J. Jiang, A. Yazdanbakhsh, E. Songhori, A. Goldie, N. Jaitly, A. Mirhoseini, Efficient Imitation Learning with Local Trajectory Optimization, ICML 2020 Workshop on Inductive Biases, Invariances and Generalization in RL https://biases-invariances-generalization.github.io/pdf/big_28.pdf
- 2019: J. Song, R. Lanka, Y. Yue, M. Ono, Co-training for Policy Learning, UAI 2019 (oral presentation),

<https://arxiv.org/abs/1907.04484>

2019: J. Song, Y. Chen, Y. Yue, A General Framework for Multi-fidelity Bayesian Optimization with Gaussian Processes, AISTATS 2019, <https://arxiv.org/abs/1811.00755>

2018: J. Song, Y. Tokpanov, Y. Chen, D. Fleischman, K. Fountaine, H. Atwater, Y. Yue, Optimizing Photonic Nanostructures via Multi-fidelity Gaussian Processes, NeurIPS 2018 Workshop on Machine Learning for Molecules and Materials, <https://arxiv.org/abs/1811.07707>

2018: J. Song*, R. Lanka*, A. Zhao, A. Bhatnagar, Y. Yue, M. Ono, Learning to Search via Retrospective Imitation, <https://arxiv.org/abs/1804.00846>

2018: C. Hawthorne*, E. Elsen*, J. Song*, A. Roberts, I. Simon, D. Eck, Onsets and Frames: Dual-Objective Piano Transcription, ISMIR 2018, <https://arxiv.org/abs/1710.11153>

2017: J. Song*, R. Lanka*, A. Zhao, Y. Yue, M. Ono, Learning to Search via Self-Imitation with Application to Risk-Aware Planning, NeurIPS 2017 Workshop on Learning with Limited Labeled Data, https://l1d-workshop.github.io/2017/papers/LLD_2017_paper_18.pdf

INDUSTRY EXPERIENCE

Research Intern at Google Brain

June 2019 - May. 2020

- Conducted research into understanding Monte Carlo Tree Search (MCTS) as a policy augmentation tool and how to improve the efficiency of running MCTS. We obtained both theoretical insights and empirical validations. This work was presented at the ICML 2020 Workshop on Inductive Biases, Invariances and Generalization in RL.
- Studied how to improve the sample complexity for interactive imitation learning algorithms through active learning. We designed a novel value-based active learning criterion and showed that it achieved superior performance under the same expert interaction budget.

Research Intern at Google Brain

June 2017 - Sept. 2017

- Completed an internship on Magenta team in Google Brain. Researched on improving polyphonic piano music transcription with convolutional neural networks. Designed a new multitask learning model and a novel loss function which together improved note-with-offset score by 100%. The work was presented at ISMIR 2018.

Software Engineering Intern at Google

June 2015 - Sept. 2015

- Researched and implemented a large scale machine learning pipeline to predict content ads final rankings after ad auctions, resulting in 20% accuracy improvement over the baseline method.

Software Engineering Intern at Yelp

May 2014 - Aug. 2014

- Implemented a brand-new framework to analyze A/B experiments for Search Quality team. The project is used by engineers and PMs from different teams including Search and Ads to evaluate experiment performances and make product decisions.

TEACHING EXPERIENCE

Teaching assistant:

- CS159 (Spring 2017, 2018, 2019, 2020): Selected Topics in Machine Learning
- CSC343 (Fall 2014): Introduction to Databases.
- CSC240 (Winter 2014): Enriched Introduction to the Theory of Computation.
- CSC108 (Fall 2013): Introduction to Computer Programming.

Research Mentor:

- Summer 2020 - now: ongoing mentorship for a SURF student at Caltech on Bayesian optimization.

- Fall 2019 - Spring 2020: ongoing mentorship for 6 undergraduates at Caltech on various research projects from accelerating integer program solvers to Bayesian optimization. One student's work was presented at the ICML 2020 Workshop on Real World Experiment Design and Active Learning.
- Fall 2017, Winter 2018: mentor a group of 4 undergraduates at Caltech on learning search heuristics for mixed integer solvers.
- Summer 2017: mentored a SURF student at Caltech, whose work was presented at the 2017 Southern California Machine Learning Symposium as a poster.

SCHOLARSHIPS AND AWARDS

Norman Stuart Robertson Scholarship in Mathematics, 2014
 Fay and David Masson Scholarship, 2014
 Galois Awards in Mathematics 2013 - 2014
 Samuel Beatty In-Course Scholarships, 2013 - 2014
 University of Toronto Scholar, 2011, 2012, 2014
 Dean's List, 2012 - 2014
 University of Toronto Excellence Award, 2013
 Alfred T. DeLury Scholarship in Mathematics, 2013
 Faculty of Arts & Science Alumni & Friends Undergraduate Scholarship, 2013
 Nicholas Martin Prize in Mathematics, 2012
 Samuel Joseph Granastein Scholarship, 2012
 Ted Mossman Scholarship, 2011

TALKS

- *A General Large Neighborhood Search Framework for Solving Integer Linear Programs*, INFORMS 2020, Nov. 2020
- *Learning to Optimize via Policy Learning*, Guest Lecture for CS 159: Data-Driven Algorithm Design, May. 2020
- *Learning to Search via Policy Optimization*, Google Tech Talk, Sept. 2019
- *A General Framework for Multi-fidelity Bayesian Optimization with Gaussian Processes*, SoCal Machine Learning Symposium, USC, Apr. 2019
- *Learning to Search via Retrospective Imitation*, INFORMS 2018 Joint Session OR Frontiers/Integer Programming/ICS: Machine Learning and Discrete Optimization I, Nov. 2018
- *Learning to Search via Retrospective Imitation*, Rice Computer Science Seminar, Rice University, Oct. 2018
- *Graph Theory and Diffusion Property of Cryptographic Hash Functions*, Ganita Seminar, University of Toronto, Nov. 2014
- *Diagnosing Rare Genetic Disorders*, UTEA Presentations, University of Toronto, Aug. 2013
- *Analyses on Incorrectly Implemented Schemes*, ENSICAEN, Aug. 2013

COMMUNITY SERVICE

| | |
|---------------------------------------------------------------------|----------------|
| Member of CMS Graduate Student Council | 2016 - Present |
| Member of Patents & Relations with Industry Faculty Board Committee | 2017 - Present |