REZA ASAD

rasad@sfu.ca

github.com/reza-asad (604)-715-4975 ca.linkedin.com/in/rezaasad

Research Interests

Designing and analyzing optimization algorithms for reinforcement and machine learning.

Education

Simon Fraser University **Jan 2020 – present**
PhD in Computer Science, Optimization in RL and ML.

Simon Fraser University **Sept 2018 – Dec 2019**
Master's in Computer Science, Visual Computing.

University of British Columbia **Sept 2013 – Apr 2015**
Master's in Applied Mathematics (3 credits short).

University of Toronto **Sept 2010 – Apr 2013**
Honours Bachelor of Science with High Distinction in Mathematics.

Publications

1. "Revisiting Actor-Critic Methods in Discrete Action Off-Policy Reinforcement Learning", **R. Asad**, R. Babanezhad, S. Vaswani. *NeurIPS Workshop on Aligning Reinforcement Learning Experimentalists and Theorists (ARLET), 2025.*
2. "Optimistic Actor-Critic with Parametric Policies: Unifying Sample Efficiency and Practicality", M. Lin, **R. Asad**, K. Tan, H. Ishfaq, S. Vaswani. *NeurIPS Workshop on Aligning Reinforcement Learning Experimentalists and Theorists (ARLET), 2025.*
3. "Fast Convergence of Softmax Policy Mirror Ascent", **R. Asad**, R. Babanezhad, I. Laradji, N. Le Roux, S. Vaswani. *International Conference on Artificial Intelligence and Statistics (AISTATS), 2025*.
4. "Surrogate Minimization: An Optimization Algorithm for Training Large Neural Networks with Model Parallelism", **R. Asad**, R. Babanezhad, I. Laradji, N. Le Roux, S. Vaswani. *Neurips Workshop on Optimization for Machine Learning (OPTML), 2023.*
5. "3DSSR: 3D Subscene Retrieval", **R. Asad**, M. Savva. *CVPR Workshop (Spotlight) on Structural and Compositional Learning on 3D Data, 2023.*
6. "Steiner Symmetrization Along a Certain Equi-distributed Sequence of Directions", **R. Asad**, A. Burchard. *Arxiv 2020.*
7. "CloudMaskGAN: A Content-Aware Unpaired Image-to-Image Translation Algorithm for Remote Sensing Imagery", S. Mohajerani, **R. Asad**, K. Abhishek, N. Sharma, A. van Duynhoven, P. Saeedi. *IEEE International Conference on Image Processing (ICIP), 2019.*

8. "Embedded Eigenvalues and the Nonlinear Schrödinger Equation", **R. Asad**, G. Simpson. *Journal of Mathematical Physics, 2011*.

Academic Research Experience

PhD Candidate, SFU **Jan 2020 – present**

- Developed a theoretically grounded policy gradient method, Softmax Policy Mirror Ascent (SPMA), bridging the gap between theory and practice in on-policy RL; published at AISTATS 2025.
- Identified a key limitation in the Discrete Soft Actor-Critic (DSAC) objective and developed a general actor-critic framework that subsumes DSAC and extends our SPMA to the off-policy setting; accepted at a NeurIPS workshop and under review for a major conference.
- Designed an optimization algorithm that reduces GPU idle time in large-scale model-parallel training while maintaining the convergence rate of gradient descent; accepted at NeurIPS workshop.
- Developed a self-supervised 3D point transformer for the task of 3D subscene retrieval. The paper received spotlight presentation at the CVPR workshop.

Master's Student, SFU **Oct 2018 – Apr 2019**

- Modified the CycleGAN loss to translate between snowy and non-snowy satellite images while preserving cloud pixels, significantly improving cloud segmentation using the generated synthetic data.
- Won the Best Project Award at SFU's AI Showcase; an extension of the work was accepted to IEEE ICIP 2019.

Recent Professional Experience

RL Scientist Intern at Microsoft Research, Montreal **Mar 2024 – Jun 2024**

- Investigated functional Nesterov acceleration in RL, building on prior work that interprets parametric Nesterov acceleration as a combination of primal (gradient descent) and dual (mirror descent) progress.
- Framed mirror ascent as a general optimization framework for on-policy RL, recovering popular objectives and leading to the design of Softmax Policy Mirror Ascent (SPMA); accepted at AISTATS 2025.

Computer Vision Scientist Intern at Terramera, Vancouver **May 2019 – Aug 2019**

- Implemented a deep crowd-counting model to estimate corn crops from drone imagery, achieving 99.96% test accuracy; integrated into Terramera's product line.
- Developed a plant phytotoxicity segmentation model, achieving ~95% accuracy; used by Terramera's biologists.

Data Scientist at Tamr, San Francisco Bay Area **Oct 2016 – Jun 2018**

- Built a classifier for categorizing manufacturing parts (~17M weekly records) at General Electric, developing automated pipelines with Spark, Elasticsearch, and PostgreSQL, resulting in $80M in estimated savings over 3 years.
- Implemented an active learning framework to optimize label acquisition, maximizing the efficiency of domain experts' time.

Honors and Awards

- CMPT Graduate Fellowship Award, 2020-2024 ($34000).
- Dean's Graduate Scholarship, 2020-2024 ($5000).
- British Columbia Graduate Scholarship (BCGS), 2020-2024 ($15000).
- Best project award at SFU's AI showcase, Simon Fraser University, Fall 2018.
- Undergraduate Student Research Awards (USRA) through Natural Sciences and Engineering Council of Canada (NSERC), University of Toronto, Summer 2012 ($6000).
- Exceptional Achievement Award, University of Toronto, 2011-2012 ($1000).
- Later Life Learning OSOTF Award, University of Toronto, 2011-2012 ($600).
- Fairfax Financial Award, University of Toronto, 2008-2012 ($20000).

Technical Skills

- **Programming Languages:** Python (Numpy, PyTorch, OpenCV, Pandas), C++ (OpenCV), R, MATLAB.
- **Tools:** PostgreSQL, Spark, Elasticsearch, Mongo DB, Kafka, Amazon EC2.

Services

- Reviewer: Neurips (2025), ICLR (2025), OPTML Workshop at Neurips (2023-2025), ARLET Workshop at Neurips (2025).
- Teaching Assistant: Technical Writing and Research Communication (2021, 2023).
- Mentored Undergraduate Students**:**
    - Grayson Lee, SFU. *Next: Master's student at SFU.*
    - Amirhossein Etaati, SFU. *Next: Software Developer at RBC.*
- Maintained SFU's Graphics Union Vision (GrUVi) website from 2020-2023.