Hamed Shirzad

Vancouver, Canada | shirzad@cs.ubc.ca | hamedshirzad.com | Google Scholar | github.com/hamed1375 | linkedin.com/in/hamedshirzad

Ph.D. candidate in Computer Science at UBC and Student Researcher at Google Research, working on reinforcement-learning post-training of large language models via self-play and self-improvement. First-author publications at NeurIPS, ICML, and AISTATS, and collaborative research in *Nature Biotechnology*. Research experience at Google Research, Valence Labs (Recursion), Borealis AI, Autodesk AI Lab, and the Max Planck Institute. Seeking Research Scientist roles in frontier AI.

Research Interests

Post-training of large language models (reinforcement learning, self-play, self-improvement); sparse and efficient transformers; graph machine learning; generative modeling; self-supervised and representation learning; reliable evaluation of learned representations; machine learning for scientific and drug discovery.

Education

University of British Columbia **2021 – Present**
Ph.D. in Computer Science — Advisor: Prof. Danica Sutherland Vancouver, Canada
Focus: efficient transformers, graph learning, and reliable evaluation.

Simon Fraser University **2019 – 2021**
M.Sc. in Computing Science — Advisor: Prof. Greg Mori Burnaby, Canada
Thesis: *Graph Generation Using Tree Decomposition*.

Sharif University of Technology **2014 – 2019**
B.Sc. in Computer Engineering & B.Sc. in Mathematics (double major) Tehran, Iran

Research Experience

Student Researcher — Google Research **Feb. 2026 – Present**
Remote

- Developing reinforcement-learning post-training methods in which language models self-improve through self-play, generating and learning from their own training signal.
- Fine-tuning Gemma 4 models up to 31B parameters with reinforcement-learning fine-tuning techniques, implemented in JAX, Flax, and the Tunix post-training library.

Research Intern — Valence Labs (Recursion), Virtual Cell **Mar. 2025 – Dec. 2025**
Montreal, Canada

- Co-developed knowledge-graph-based structured-prediction models of transcriptomic perturbation effects, contributing to TxPert (*Nature Biotechnology*, 2026).
- Developed SeedER, combining reinforcement learning (RL) and Graph Transformers to learn seed-and-expand retrieval policies for biological knowledge graphs; first-author paper under review at ICLR 2027.
- Collaborated in an interdisciplinary team connecting machine-learning methods with large-scale biological prediction problems in drug discovery.

Research Intern — Borealis AI (now RBC Borealis) **Sep. 2023 – Feb. 2024**
Vancouver, Canada

- Developed conditional diffusion models as self-supervised learning backbones for irregular event-sequence and time-series data; first-author paper at the Time Series for Health Workshop, ICLR 2024.

Research Intern — Autodesk AI Lab **Apr. 2021 – Sep. 2021**
Remote

- Developed contrastively learned evaluation metrics for graph generative models, resulting in a first-author NeurIPS 2022 paper.

Research Intern — Borealis AI **Apr. 2020 – Mar. 2021**
Remote

- Developed permutation-invariant graph generation via tree decomposition with Prof. Greg Mori and Dr. Hossein Hajimirsadeghi, leading to the first-author AISTATS 2022 paper TD-Gen.

Research Intern — Max Planck Institute for Intelligent Systems **Jul. 2018 – Sep. 2018**
Tübingen, Germany

- Worked on kernelized Stein discrepancy and its applications under Prof. Krikamol Muandet.

Selected Publications

(* equal contribution; full list on Google Scholar)

SeedER: Seed-Expand-Retrieve for Efficient Knowledge Graph Retrieval
H. Shirzad, F. Wenkel, D. Beaini, D. J. Sutherland, E. Noutahi — *Under review at ICLR 2027*
Combines reinforcement learning (RL) and Graph Transformers to learn query-conditioned graph-expansion policies for efficient knowledge-graph retrieval. [arXiv]

TxPert: Using Multiple Knowledge Graphs for Prediction of Transcriptomic Perturbation Effects
F. Wenkel, W. Tu, C. Masschelein, **H. Shirzad**, et al. — *Nature Biotechnology, 2026*
Structured prediction over multiple knowledge graphs for biological perturbation modeling. [paper] [GitHub] [blog]

Even Sparser Graph Transformers
H. Shirzad, H. Lin, B. Venkatachalam, A. Velingker, D. Woodruff, D. J. Sutherland — *NeurIPS 2024*
Sparse graph-transformer architecture for scalable learning on graphs. [arXiv] [GitHub]

Exphormer: Sparse Transformers for Graphs
H. Shirzad[*], A. Velingker[*], B. Venkatachalam[*], D. J. Sutherland, A. K. Sinop — *ICML 2023*
Linear-complexity graph transformers via expander graphs; featured on the Google Research blog. [arXiv] [GitHub]

Evaluating Graph Generative Models with Contrastively Learned Features
H. Shirzad, K. Hassani, D. J. Sutherland — *NeurIPS 2022*
Evaluation methodology for graph generative models using contrastively trained representations. [arXiv] [GitHub]

TD-Gen: Graph Generation Using Tree Decomposition
H. Shirzad, H. Hajimirsadeghi, A. H. Abdi, G. Mori — *AISTATS 2022*
Tree-decomposition-based, permutation-invariant graph generation. [arXiv]

A Theory for Compressibility of Graph Transformers for Transductive Learning
H. Shirzad, H. Lin, B. Venkatachalam, A. Velingker, D. Woodruff, D. J. Sutherland — *MLComp Workshop, NeurIPS 2024*
Theoretical analysis of graph-transformer compressibility. [arXiv]

Conditional Diffusion Models as a Self-supervised Backbone for Irregular Time Series
H. Shirzad, R. Deng, H. Zhao, F. Tung — *Time Series for Health Workshop, ICLR 2024*
Diffusion-based representation learning for irregular time series. [paper]

Global Communication or Graph-Specific Memory?
H. Shirzad, D. J. Sutherland — *Preprint, 2026* [arXiv]

Technical Skills

Languages: Python, C++, Bash, Java, MATLAB.
ML stack: JAX, Flax, Tunix, PyTorch, PyTorch Geometric, DGL, NumPy, SciPy, pandas, NetworkX.
Methods: LLM post-training, reinforcement learning (RL fine-tuning, self-play), transformers, graph learning, generative and diffusion models, self-supervised learning, model evaluation.

Awards & Recognition

2026	Top Reviewer: NeurIPS (complimentary registration); Learning on Graphs (LoG, 2024; monetary award).
2024	Fully Funded Trainee, Advanced Machine Learning Training Network (AML-TN); $32K stipend and $4K travel support.
2019	Graduate Fellowship, Simon Fraser University.
2014–17	National Elites Foundation Grant for top university entrance exam achievers and National Olympiad medalists.
2013	Bronze Medal, Iranian National Olympiad in Informatics.

Invited Talks

Keynote: *Scaling Graph Transformers with Sparse and Sparsified Attention*, Graph Signal Processing Workshop, 2025.
Tutorial: *Geometric and Topological Representation Learning*, Machine Learning for Signal Processing Conference, 2025.

Leadership & Service

- Reviewer (2022–2026): ICLR, NeurIPS, LoG, AISTATS, TMLR, and JMLR.
- Mentor, Research Experience Program (REX), UBC (2024): guided undergraduate students through a first research project.
- Co-founder, Homa Foundation (2023): helped raise $45K for 25 Iranian graduate students' travel to Canada.