# Haseeb Shah

*Edmonton, AB, Canada* hshah1@ualberta.ca www.haseeb.sh haseebs haseebs

## Education

**University of Alberta** Canada

PH.D. COMPUTER SCIENCE *Sep. 2024 - present*

**Supervisor:** Martha White

**Teaching Assistant:** AI Everywhere Course

**Lab Manager:** Martha White's Lab

**University of Alberta** Canada

M.SC. COMPUTER SCIENCE (THESIS-BASED) *Jan. 2021 - Jan. 2023*

**Supervisor:** Martha White **GPA:** 4.00/4.00

**Teaching Assistant:** Intermediate Machine Learning; Introduction to Computer Science

**Thesis:** Greedy Pruning for Continually Adapting Networks

## Publications [Google Scholar]

**$q$-Exponential Family for Policy Optimization** ICLR-25

L. ZHU*, **H. SHAH***, H. WANG*, M. WHITE | *EQUAL CONTRIBUTION Paper - Code

Offline + Online Reinforcement Learning | Policy Gradient Algorithms

We explore the effectiveness of $q$-exponential policies in policy optimization methods, finding that heavy-tailed policies ($q > 1$) are generally more effective and can consistently outperform the Gaussian policy.

**Scalable Real-Time Recurrent Learning Using Columnar-Constructive Networks** JMLR-23 [Journal]

K. JAVED, **H. SHAH**, R. SUTTON, M. WHITE Paper

Online Reinforcement Learning | Continual Learning | Neural Network Sparsity

Developed novel methods for making RTRL scale linearly with the number of parameters, without introducing noise or bias to gradient estimates. Demonstrated effectiveness by doing policy evaluation of pre-trained policies for Atari 2600 games.

**GVFs in the Real World: Making Predictions Online for Water Treatment** MLJ-23 [Journal]

K. JANJUA, **H. SHAH**, M. WHITE, E. MIAHI, M. C. MACHADO, A. WHITE Paper - Code

Offline + Online Reinforcement Learning | Time-series Prediction | Real-world System

Proposed a framework for making real-time predictions in water treatment plants using General Value Functions, addressing non-stationary high-volume systems.

**An Open-World Extension for Knowledge Graph Completion Models** AAAI-19 [conference - oral]

**H. SHAH**, J. VILLMOW, A. ULGES, U. SCHWANECKE, F. SHAFAIT Paper - Code

Language Models | Knowledge Graphs

Introduced a novel method enabling knowledge graph completion models to predict facts about open-world entities while demonstrating superior performance.

## Under Review and Workshops

**Deconstructing actor-critic: a large-scale empirical study of design components for practitioners** PNAS [Journal]

**H. SHAH**, L. ZHU, A. WHITE, M. WHITE *Under review*

Reinforcement learning | Algorithm reliability | Process control

We run and analyze over 33,000 experiments on a task derived from a real-world system to identify how actor-critic components affect the performance and stability in deployment, and highlight that many existing defaults are relatively more unstable.

**Stable learning in deployment while mitigating performance degradation in offline-to-online RL** — RLC-26 workshop

H. Wang, **H. Shah**, A. White, M. White — *Under review*

Offline-to-Online Reinforcement Learning

We introduce a new offline-to-online fine-tuning algorithm that gradually allows more exploration based on off-policy estimates of performance. We also propose new metrics to properly measure performance degradation.

**Symmetric Behavior Regularized Policy Optimization** — NeurIPS-26

L. Zhu, **H. Shah**, C. Zheng, N. Yukie, M. White — *Under review*

Offline Reinforcement Learning | Policy Gradient Algorithms

We study Symmetric Divergences for Behavior Regularized Policy Optimization and propose Symmetric $f$-Actor Critic, which avoids per-environment failures encountered by other offline Reinforcement Learning methods.

**Relation Specific Transformations for Open World Knowledge Graph Completion** — COLING-20

**H. Shah**, J. Villmow, A. Ulges — Paper - Code

Language Models | Knowledge Graphs

Introduced relation-specific transformations to substantially improve the performance of Open World Knowledge Graph Completion models. Additionally, proposed an approach for clustering of relations to reduce the training time and memory footprint.

**Distillation Techniques for Pseudo-rehearsal Based Incremental Learning** — Preprint

**H. Shah**, K. Javed, F. Shafait — Paper - Code

Continual Learning | Computer Vision | Generative Adversarial Networks

Demonstrated that existing methods to combat catastrophic forgetting using Generative Adversarial Networks are biased and proposed an approach to mitigate this bias.

## Research Visits

**Advanced Laser Light Source Laboratory, INRS** — Montreal, Canada

Reinforcement Learning Researcher — *Sept. 2025 - Oct. 2025*

- Designed a Reinforcement Learning algorithm for optimization of Laser WakeField Acceleration (LWFA)
- Deployed and tested on a live system at INRS Montreal with the help of physicists
- The deployed agent was able to successfully solve the multi-objective optimization problem without any offline pre-training

## Public Talks

**Online Feature Decorrelation for Neural Networks** — Amii - 2022

Tea Time Talks — Talk - Slides - Code

Online Reinforcement Learning | Continual Learning | Neural Network Sparsity

I present a new approach for decorrelating features in an online setting and demonstrate that it can effectively eliminate redundant features while producing a statistically significant performance improvement in the low-capacity function approximation setting.

## Work Experience

**RL Core Technologies** — Edmonton, Canada

Machine Learning Intern — *Jan. 2024 - Aug. 2024*

- Designed and implemented reinforcement learning-based solutions for optimizing and automating drinking water treatment plants in Alberta.
- Deployed the agents on the real physical system in production, resulting in significantly increased filter permeability recovery during backwashes and reduced costs during the chemical dosing process.

**Reinforcement Learning and Artificial Intelligence Lab (RLAI), University of Alberta** — Edmonton, Canada

Research Assistant with Dr. Martha White — *Feb. 2023 - Jan. 2024*

- Conducted research on the topics of neural network sparsity, recurrent learning and general value functions.
- Designed and deployed Reinforcement Learning-based control algorithms on a real-world drinking water treatment plant.
- Assumed the role of Lab Manager, facilitating the lab's operations and research activities.

**Learning and Visual Systems Group (LAVIS), Hoschule RheinMain** Wiesbaden, Germany

Research Assistant with Dr. Adrian Ulges *Jan. 2020 - Dec. 2020*

- Ported the original code from Tensorflow to Pytorch for semantic code search
- Designed and evaluated novel approaches for semantic code search, leveraging tree transformers and abstract syntax trees.
- Improved the state-of-the-art in open-world link prediction tasks.

**DCube Technologies** Islamabad, Pakistan

Machine Learning Engineer *May. 2020 - Aug 2020*

- Constructed knowledge graphs from raw legal case documents for the Supreme Court of Pakistan.
- Conducted statistical analysis for hospitals in Australia.

**Learning and Visual Systems Group (LAVIS), Hoschule RheinMain** Wiesbaden, Germany

Research Intern with Dr. Adrian Ulges *Jun. 2018 - Sep. 2018*

- Constructed a new benchmark for open-world link prediction.
- Implemented and evaluated novel approaches for open-world link prediction and compared them against the baselines.

**TUKL-NUST Research and Development Center, NUST** Islamabad, Pakistan

Research Assistant with Dr. Faisal Shafait *Jun. 2017 - Dec. 2017*

- Proposed an algorithm which improved state-of-the-art in retrieving visually similar fashion images.
- Proposed an algorithm for mitigating catastrophic forgetting in continual learning tasks using GANs.

## Honors and Awards

2024 **University of Alberta**: Ph.D. recruitment scholarship

2019 **AAAI**: Travel scholarship for Honolulu, Hawaii, to present my paper at the conference

2018 **DAAD**: Fully funded summer research exchange to Germany by the German Academic Exchange Service

2016 **National University of Sciences and Technology**: Merit-based scholarship for high-achievers (every semester)

2016 **ACM**: First place winner in the ACM digital design competition

## Miscellaneous

2025 **American Water Works Association**: Poster on ”Reinforcement Learning for Control of Water Treatment Plants”

2024 **Lifesaving Society Alberta and Northwest Territories**: Certified in Intermediate First Aid (CPR-C/AED)

2024 **University of Alberta Fencing Club**: Risk Management Executive

2023 **University of Alberta**: Created teaching and assessment content for the ”AI Everywhere course”

2023 **Alberta Machine Intelligence Institute**: Created case study on ”The Spectrum of Intelligence We See in Animals”

2019 **Eberhard Karls University of Tübingen, Germany**: Studied in M.Sc. Machine Learning program for a term

2018 **Pakistan Pattern Recognition Society**: Teaching assistant at PPRS’s first Autumn School on Deep Learning

2016 **AIESEC**: Assisted with the branding and communications by designing posters and artwork

## Relevant Skills

**Languages** Python · C++ · C

**ML Libraries** PyTorch · Tensorflow (1.0) · Jax

**Others** Latex · Git · MySQL/MariaDB · Vim · Unity Game Engine · Blender · Adobe Photoshop