

Hammad Rizwan

✉ hammad.rizwan@dal.ca | [in](#) [hammad-rizwan](#) | [Google Scholar](#)

EDUCATION

Dalhousie University

PhD Computer Science

Advisor: Hassan Sajjad

Halifax, Canada

Sept. 2023 –

Lahore University of Management Sciences (LUMS)

MS Computer Science

Advisor: Asim Karim

Lahore, Pakistan

Sept. 2018 – Jun 2020

Thesis Hate-Speech and Offensive Language Detection in Roman Urdu

RESEARCH INTERESTS

Model editing, machine unlearning, interpretability, and self-supervised learning.

PUBLICATIONS

Peer-reviewed

- **Hammad Rizwan** Domenic Rosati, Ga Wu, Hassan Sajjad “Resolving Lexical Bias in Model Editing” *Proceedings of the 42nd International Conference on Machine Learning (ICML), 2025*. [Paper Link] (26.9% acceptance rate)
- Ubaid Azam, **Hammad Rizwan**, Asim Karim “Exploring Data Augmentation Strategies for Hate Speech Detection in Roman Urdu.” *Proceedings of the Thirteenth Language Resources and Evaluation Conference (LREC), 2022*. [Paper Link]
- **Hammad Rizwan**, Muhammad Haroon Shakeel, Asim Karim “Hate-Speech and Offensive Language Detection in Roman Urdu.” *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2020*. [Paper Link] (24.5% acceptance rate)

Preprints

- Muhammad Umair Haider, **Hammad Rizwan**, Hassan Sajjad, Peizhong Ju, AB Siddique “Neurons Speak in Ranges: Breaking Free from Discrete Neuronal Attribution”, 2025 (Under Review) [Paper Link]
- **Hammad Rizwan***, Mahtab Sarvmali*, Hassan Sajjad, Ga Wu “Instance-Level Difficulty: A Missing Perspective in Machine Unlearning”, 2025 [Paper Link]
- **Hammad Rizwan***, Ubaid Azam*, Faizad Ullah, Asim Karim “Towards Explainable Hate Speech Detection in Roman Urdu” [TBD Link]

ACADEMIC EXPERIENCE

Teaching Assistant

- Dalhousie University:
Introduction to Robotics (CSCI-1108, Fall and Winter 2023,2024), Deep Learning Applications (CSCI-5501, Summer, 2025)
- LUMS:
Natural Language Processing (CS-5316, Fall 2019), Introduction to Programming (CS-200, Winter 2018)

HyperMatrix Lab at Dalhousie University

Sept. 2023 –

Doctoral Researcher

- Conducting research on weight-preserving techniques for efficient and interpretable model editing.
- Designing and evaluating algorithms for certified machine unlearning, under the supervision of Dr. Ga Wu.

Robotics and Intelligent Computing (RICE) Lab at LUMS

Nov. 2020 – June 2022

Research Associate

Digital Forestry

- Built a monocular vision-based deep learning system for automated tree diameter estimation, improving forest monitoring efficiency.
- Created Pakistan's first annotated dataset for tree classification and semantic segmentation, used by WWF and Forestry Pakistan.
- Delivered scalable, low-cost AI solutions for national forestry management initiatives.

Knowledge and Data Engineering (KADE) Lab at LUMS

Sept. 2019 – June 2020

Research Assistant

Hate Speech and Generative Models

- Created the first open-source Offensive and Hate Speech Detection dataset for Roman Urdu and developed a custom model that outperformed existing models by 3 Macro F1 points.
- Explored differentiable surrogate objectives for GANs based paraphrase generation

INDUSTRY EXPERIENCE

Upbieng.ai

Feb. 2024 – Nov. 2024

Mitacs Accelerate Intern

Mental Health Applications

- Designed and implemented a dialogue system for the Upbeing mental health application, contributing to research in AI-driven mental health support.

Arbisoft

Feb. 2023 – July 2023

Senior ML Engineer

Social Media Applications

- Developed a personalised recommendation system for a dining-focused social media application.
- Enhanced entity extraction accuracy on noisy short-text inputs, improving downstream search relevance and recommendation quality.

ISSM.ai

Oct. 2022 – Feb. 2023

NLP Engineer

Automation in Banking Processes

- Designed and deployed Urdu/Roman Urdu transliteration and semantic similarity models to power multilingual banking automation workflows.
- Led chatbot development for United Bank Limited (UBL) and Allied Bank, resulting in faster query resolution and reduced support load.

Eighty.ai

Jan. 2021 – Dec. 2021

NLP Engineer

Document Processing

- Developed a commodity classification recommendation system based on harmonised system (HS) codes across 97 categories and 1,244 subcategories.
- Built high-precision document labelling and extraction system by developing custom NER and embedding models for the financial domain.

Ai XPRT (AuditXPRT)

July 2020 – Dec. 2020

NLP Engineer

Document Processing

- Created a system to detect high-risk individuals/corporations for anti-money laundering efforts at AIA Insurance.
- Enhanced ComplyXPRT, a regulatory compliance system for Barclays Bank, integrating lifelong learning for adaptive modelling and achieving a 30% reduction in inference time through model distillation and multithreaded deployment.

AWARDS AND SCHOLARSHIP

Doctoral Research Scholarship

2023 -2027

Dalhousie University

Undergraduate Merit Scholarship

2014 -2017

COMSATS

50% Tuition Waiver

INVITED TALKS

Model Interpretability and Model Editing <i>Deep learning for Natural Language Processing - Dalhousie University</i>	2025
Exploration of Neuronal Activations <i>AI Cluster - Dalhousie University</i>	2024
Computer Vision for Digital Forestry <i>National Agriculture Robotics Workshop - LUMS</i>	2022

PROFESSIONAL SERVICES

- Co-Organiser: Machine Learning Theory Seminar - Dalhousie University. 2024,2025
- Conference Reviewer: NeurIPS 2024,2025. ICML 2025. ICLR 2025,2026.
- Journal Reviewer: LRE, IPM, NLE