Hamed Damirchi

Adelaide, Australia | hamdamirchi@gmail.com
hamed-d.github.io | GitHub | Google Scholar | LinkedIn

Machine learning researcher and engineer: LLM reasoning, parameter-efficient fine-tuning, model editing and multi-agent RAG systems; applied work in energy forecasting, healthcare AI and robotics. First-author ACL 2026 oral paper, Best Paper Award nominee. PhD at Adelaide University (AIML), expected 2026.

Technical skills

LLM methods RAG; multi-agent LLM systems; LLM judges; parameter-efficient fine-tuning (LoRA); model editing.
Programming and development Python; C/C++; CUDA; MATLAB/Simulink; Git; Docker.
Machine learning PyTorch; TensorRT; distributed training; low-bandwidth distributed training.
Robotics and embedded SLAM; ROS/ROS2; GTSAM; STM32; embedded Linux.

Doctoral research

Doctoral Researcher
2023–present

Adelaide University | Australian Institute for Machine Learning (AIML)

- **Oral presentation | Best Paper Award nominee** (ACL 2026, first author). Developed methods to assess LLM reasoning from internal representation trajectories, studying how changes across layers distinguish correct and incorrect responses.
- First-author follow-up preprint combines residual-stream motion with coarse regions and state directions, improving answer selection on unseen reasoning benchmarks and transferring to factual tasks.
- Developed Partial-LoRA, reducing trainable parameters by up to 87% while maintaining or improving accuracy in experiments across 8 vision and 12 language tasks.
- Developed task-vector decomposition for more targeted model editing; experiments reported 47% lower toxicity while preserving performance on general-knowledge tasks.
- Co-authored work on CLIP backbone ensembling, vision-transformer feature fusion, and representation trajectories for out-of-distribution detection.

Applied ML and engineering experience

Research Officer (RO4)
2024–2025

Australian National University | Centre for Public Awareness of Science

- Designed, implemented and deployed an LLM-based multi-agent framework for AI policy deliberation and risk mitigation, using RAG context; established evaluation and benchmarking frameworks.
- Four-person team with a social scientist, a psychology student and a tech lead. Designed personas with the psychology student, capturing stakeholder interests, needs, constraints and goals; designed simulations and preprocessed and formatted data with the social scientist.
- Used LLM judges to oversee deliberations and prevent repetitive loops, supporting human judges and stakeholder-team or independent representatives. Applied to Australian government AI governance scenarios (about 400 national and international stakeholders) and Murray River water disputes.

Casual Academic Researcher
Jan–Jun 2025

Adelaide University | Electricity Demand Forecasting

- Designed and implemented the forecasting algorithm and its evaluation and benchmarking, combining NeuralProphet, Gaussian processes and probabilistic model fusion.
- Forecast local demand two weeks ahead and the state's top-10 annual peaks to inform battery charging, discharge and solar use.
- Four-person team with two software engineers and a tech lead. Set milestones with the client and tech lead, who set final goals and managed the project and schedule; the startup's engineers collected data and implemented the algorithm interface.

Tech Lead
2021–2022

ARAS Labs | Healthcare AI: Corneal-disease detection

- Led technical and research work, directing three engineers on algorithm implementation and the pipeline from data processing to predictions.
- Coordinated data collection with ophthalmic surgeons and annotation by ophthalmologists; designed label aggregation to derive a single ground-truth label from multiple experts.
- Worked with the hospital to anonymise clinical data and establish an on-site data-processing environment, ensuring identifiable data remained on hospital premises.
- Proposed a weakly supervised learning approach for corneal-disease detection, including keratoconus.

Researcher 2022–2024
University of South Carolina | AISys

- Designed and implemented a compositional learning method based on independent causal mechanisms, automatically extracting generative factors from vision datasets without predefined factor assignments.
- Presented Independent Modular Networks at ICRA 2023, Workshop on effective Representations, Abstractions, and Priors for Robot Learning (RAP4Robots).

Researcher 2018–2022
ARAS Labs | Robotics: Localisation and sensor fusion

- Seven-person team spanning mechanical, control, electronics and mechatronics engineering. Fused position and orientation estimates from camera and inertial sensors with uncertainty-aware robot kinematics for cable-robot localisation.
- With another engineer, designed and implemented pose-estimation algorithms for each sensor array, and designed and manufactured the robot's end effector, data-acquisition (DAQ) system and joints.
- Proposed an uncertainty-aware localisation subsystem for long- and short-term odometry pipelines for autonomous-vehicle navigation.

Teaching and tutoring

Secondary Instructor, Large Language Models and Applications (COMP6012) Jun 2026–present
Adelaide University | Target audience: Master's students | 120 students taught

- Delivered 30 hours across 10 practicals/workshops on LLM foundations, training, fine-tuning, RAG and evaluation; updated materials and graded student work.

Teaching Assistant, Statistical Machine Learning Jul–Nov 2025
Adelaide University | Target audience: Master's and bachelor's students | 50 students taught

- Delivered workshops and practicals explaining algorithms, implementation and method selection for unfamiliar problems; graded student work.

Education

PhD in Computer Science | Adelaide University 2023–2026 (expected)
Australian Institute for Machine Learning (AIML) | Centre for Augmented Reasoning
Advisors: Javen Shi (PI); Yuhang Liu (Co-Advisor); Zhen Zhang (Co-Advisor)
Thesis: Towards Compositional Learning: Modifying Conceptual Mechanisms in Neural Networks
Focus: Enabling learning and inference to solve problems by identifying parts and how they interact to form a whole.

MSc, Mechatronics Engineering K. N. Toosi University of Technology 2018–2021

BSc, Mechanical Engineering Tabriz University 2013–2018

Selected publications

Truth as a Trajectory: What Internal Representations Reveal About Large Language Model Reasoning
H. Damirchi, I. Meza De la Jara, E. Abbasnejad, A. Shamsi, Z. Zhang, J. Q. Shi.
ACL 2026 | Long Papers. **Oral presentation | Best Paper Award nominee**

Reasoning Errors Have a Region and a Direction in the Residual-Stream Trajectory of LLMs
H. Damirchi, I. Meza De la Jara, D. Ranasinghe, Y. Liu, J. Shi.
Preprint, 2026 | arXiv:2608.05660.

The Quest for Winning Tickets in Low-Rank Adapters
H. Damirchi, C. Rodriguez-Opazo, E. Abbasnejad, Z. Zhang, J. Shi.
Preprint, 2025 | arXiv:2512.22495.

Decomposing Task Vectors for Refined Model Editing
H. Damirchi, E. Abbasnejad, Z. Zhang, J. Shi.
Preprint, 2025 | arXiv:2512.22511.

Synergy and Diversity in CLIP: Enhancing Performance Through Adaptive Backbone Ensembling
C. Rodriguez-Opazo, E. Abbasnejad, D. Teney, **H. Damirchi**, E. Marrese-Taylor, A. van den Hengel.
ICLR 2025.

A Consistency-Based Loss for Deep Odometry Through Uncertainty Propagation
H. Damirchi, R. Khorrambakht, H. D. Taghirad, B. Moshiri.
ICRA 2023 | IEEE International Conference on Robotics and Automation.

Graph-Based Visual-Kinematic Fusion and Monte Carlo Initialization for Fast-Deployable Cable-Driven Robots
R. Khorrambakht, **H. Damirchi**, M. R. Dindarloo, A. Saki, S. A. Khalilpour, H. D. Taghirad, S. Weiss.
IROS 2023 | IEEE/RSJ International Conference on Intelligent Robots and Systems.

Full publication list: Google Scholar.