## Heshan Devaka Fernando

Electrical, Computer, and Systems Engineering
Rensselaer Polytechnic Institute, Troy, NY 12180.
+1-518-961-5360 | fernah@rpi.edu | Google Scholar | GitHub | LinkedIn | heshandevaka.github.io


---

**Research Interests**

Optimization, Machine Learning, Multi-objective learning, Multi-modal Learning, Large Language Models

**Education**

**Renssealaer Polytechnic Institute (RPI)**, Troy, New York, USA *(Jan 2021 - Present)*

*Ph.D. in Electrical, Computer and Systems Engineering*

- Cumulative GPA (Current): 3.97/4.00
- Related Core Modules: Introduction to Optimization, Introduction to Deep Learning, Machine Learning from Data, Machine Learning and Optimization, Introduction to Stochastic Signals and Systems, Detection and Estimation Theory, Data Analysis in Economics and Finance

**University of Moratuwa**, Moratuwa, Sri Lanka *(Jan 2015 - Dec 2018)*

*B.Sc. Eng. (Hons.) in Electronic and Telecommunication Engineering*

- Cumulative GPA: 3.79/4.20 (First Class Honors)

**Research Experience**

**Multi objective learning (MOL)**

Rensselaer Polytechnic Institute, USA *(Dec 2021 - Apr 2024)*

- Investigated the non-convergence issues in stochastic non-convex multi-objective learning (MOL), and proposed a novel algorithm called MoCo for MOL that has provable convergence guarantees in convex and non-convex settings, one of the first of it's kind. The implementation of MoCo is available in the prominent open source multi-task learning library LibMTL.
- Analysed the interplay between optimization, generalization, and conflict aversion in MOL.
- Applied variance reduction methods used in classical single objective optimization in MOL for provably faster convergence.

**Efficient Post-training Framework for Large Language Models (LLMs)**

Rensselaer Polytechnic Institute, USA *(Aug 2024-Present)*

- Post-training of pre-trained LLMs typically consists of two stages: supervised fine-tuning (SFT) and preference learning, where SFT improves task-specific performance, while preference learning aligns models with human preferences for safety and effectiveness.
- The widely adopted standard approach for LLM post-training is sequential training, where SFT is followed by preference learning.
- Theoretically and empirically showed the sub-optimality of the standard approach for post-training, and propose a joint SFT and preference learning framework that is provably optimal, with minimal computation overhead compared to the standard approach.

**Improving Multi-modal Learning (MML) via Multi-objective Optimization (MOO)**

Rensselaer Polytechnic Institute, USA *(Dec 2023-Present)*

- Multi-modal learning (MML) integrates information from multiple modalities to achieve superior performance compared to single-modality learning, but studies reveal that vanilla MML can under perform compared to single-modality approaches due to imbalanced learning across modalities.
- Proposed a MOO-inspired modification to the MML objective and corresponding algorithm with theoretical convergence guarantees that achieve up to $\sim 3\times$ reduction in subroutine computation times empirically compared to baselines.

**Professional & Teaching Experience**

**Data Scientist Intern**, Amazon, New York, NY, USA *(May 2025 - Aug 2025)*

- Conducted exploratory data analysis on customer interactions with Alexa+.
- Built a customer action identification pipeline using unsupervised learning techniques and large language models (LLMs), and built a computationally efficient customer action detection pipeline that can detect identified customer actions at scale
- Demonstrated the efficacy of the proposed customer action identification and detection pipeline by creating interpretable customer behavior (e.g., retention, engagement, etc.) prediction models that use customer actions identified in the identification pipeline as features.

**Research Extern**, IBM Thomas J. Watson Research Center, NY, USA *(May 2022 - Aug 2022)*

- Conducted research in the area of hierarchical RL, specifically the option-critic framework.
- Investigated the possibility of designing a framework to discover sub-tasks in RL tasks by learning and maximizing some intrinsic reward, while maximizing the environment (extrinsic) reward, using a bi-level optimization approach. The underlying motivation is to discover sub-tasks which have better interpretability due to the learned intrinsic reward functions (sub-goals).
- Implemented the proposed framework in a grid-world RL navigation task.

**Graduate Assistant**, Rensselaer Polytechnic Institute, USA *(Jan 2021-Present)*

- Conduct research in the areas related to theoretical optimization and its applications.
- Teaching assistant for Machine learning and electrical engineering classes.

**Lecturer on Contract**, University of Moratuwa, Sri Lanka *(Jul 2019 - Dec 2020)*

- Conducted lab sessions and tutorials for signal processing and electronics classes

**Temporary Instructor**, University of Moratuwa, Sri Lanka *(Jan 2019 - Jul 2019)*

- Designed and supervised undergraduate electronics projects

**Consultant** (part time), Codify Lanka (Pvt) Ltd ,Colombo, Sri Lanka *(Apr 2019 - Apr 2020)*

- Built a robust Optical Character Recogniser (OCR) for information extraction from invoices and a commercial automated produce procurement platform for restaurants

## Awards & Achievements

**Notable top 5% Acceptance (Oral) and Travel Grant** at the Eleventh International Conference on Learning Representation (ICLR) 2023 *(May 2023)*

**Recognized on the Dean's List in semesters 1, 3, 6, 7, and 8** during undergraduate studies at the University of Moratuwa, Sri Lanka *(Jan 2015- Dec 2018)*

**Full Student Scholarship and Mahapola Higher Education Scholarship** for merit performance in university entrance exam by Ministry of Higher Education, Sri Lanka. *(Aug 2014)*

## Selected Publications

- **H. Fernando**, H. Shen, M. Liu, S. Chaudhury, K. Murugesan, and T. Chen. "Mitigating Gradient Bias in Multi-objective Learning: A Provably Convergent Approach." International Conference on Learning Representations (2023). [Paper] [Poster] [Presentation] [Code]
- **H. Fernando***, H. Shen*, P. Ram, Y. Zhou, H. Samulowitz, N. Baracaldo, and T. Chen. "Understanding Forgetting in LLM Supervised Fine-Tuning and Preference Learning - A Convex Optimization Perspective." ArXiv preprint (2024). [Paper] [Code]
- **H. Fernando**, P. Ram, Y. Zhou, H. Samulowitz, N. Baracaldo, and T. Chen. "Mitigating Modality Imbalance in Multi-modal Learning via Multi-objective Optimization." ArXiv preprint (2025). [Paper] [Code]
- L. Chen*, **H. Fernando***, Y. Ying, and T. Chen. "Three-way Trade-off in Multi-objective Learning: Optimization, Generalization and Conflict-avoidance." Advances in Neural Information Processing Systems 36 (2024). [Paper] [Poster] [Code]
- **H. Fernando**, L. Chen, S. Lu, P.-Y. Chen, M. Liu, S. Chaudhury, K. Murugesan, G. Liu, M. Wang, and T. Chen. "Variance Reduction Can Improve Trade-Off in Multi-Objective Learning." IEEE International Conference on Acoustics, Speech and Signal Processing (2024). [Paper] [Code]

## Skills

**Programming skills**: Python, Matlab, R, C
**Technical skills**: Pytorch, Keras, TensorFlow, Agentic AI, Git
**Languages**: English, Sinhalese

## Professional Service and Leadership

**Reviewer/Technical Program Committee** - Advances in Neural Information Processing Systems (NeurIPS), International Conference on Machine Learning (ICML), International Conference on Learning Representation (ICLR), International Conference on Artificial Intelligence and Statistic (AISTATS), AAAI Conference on Artificial Intelligence (AAAI), Transactions on Machine Learning Research (TMLR), IEEE Transactions on Signal Processing (TSP) *(Jan 2023 - Present)*

**Student Coordinator** - IEEE Communications Society - Sri Lanka Chapter *(Jul 2019 - Dec 2020)*

**Organizing Committee** - ENTC Careers Fair 2018, Department of Electronic and Telecommunication Engineering, University of Moratuwa, Sri Lanka *(Sep 2018)*

**Vice President (Membership)** - Catholic Students' Movement, University of Moratuwa, Sri Lanka *(Jan 2015 - Jan 2016)*