

Tharindu Cyril Weerasooriya

✉ thacyrweer@gmail.com • 🌐 cyrilw.com • in cyrilw • 🌐 cyrilw
🔑 P6zvClgAAAAJ

Research Highlights

Generative AI, Agents, and Reasoning

Focus: Developing agentic frameworks for Text-to-SQL and code generation, and optimizing LLM reasoning via multi-source genetic algorithms (MultiGA).

Publications: NeurIPS '25 Workshop [3], Enterprise Analytics Agents (FALM) [7], and AI Refinery.

Fairness in Human Feedback & Disagreement Prediction

Focus: Building socially representative datasets and predicting human disagreement in subjective tasks (hate speech, offensive language) to improve model fairness.

Publications: AAAI '25 [1], EMNLP '24 [8], EMNLP '23 [9], ACL '23 [11, 12], ECAI '20 [15].

Interests: Large Language Models, AI Agents, Human Feedback (RLHF), Label Distribution Learning, Crowdsourcing, NLP, and Computational Social Science.

Education

Rochester Institute of Technology

Ph.D. in Computer Science, GPA 3.9

Machine Learning, Natural Language Processing, Fairness in AI, Data Science, and Computational Social Science.

Thesis: Fairness in Human Annotation. **Advisor:** Dr. Christopher M. Homan

Rochester, NY

Aug 2018–May 2024

University of Kelaniya

B.Sc. in Computer Science

Thesis: KeyXtract Twitter Model - An Essential Keywords Extraction Model for Twitter. **Advisor:** Dr. Sidath Liyanage

Sri Lanka

Mar 2013–Mar 2017

Experience

Accenture - Center for Advanced AI

Research Scientist, CA

Mountain View

April 2024 - Current

Collaborative GenAI Research & Development

- **Fortune Analytics (FALM):** Co-developed a domain-specific code generation LLM (Mixtral) with a cross-functional team of researchers and engineers to enable natural language analytics for Fortune 500 data [7].
- **AI Refinery (NVIDIA Collab):** Partnered with NVIDIA and internal marketing teams to architect "AI Refinery," leading the implementation of self-reasoning Llama 3.1 agents to automate complex enterprise workflows.
- **Safety & Guardrails:** Led a safety initiative across the engineering organization, designing and integrating automated prompt evaluation pipelines to enforce strict topic guardrails for agentic systems.
- **MultiGA (NeurIPS '25):** Directed research strategy and co-authored "MultiGA" with intern and senior research collaborators, introducing a multi-source seeding framework (GPT-4o, DeepSeek) for evolutionary algorithms [3].

Technical Leadership & Community

- **Internship Program Lead:** Coordinated the Summer 2025 hiring pipeline in collaboration with leadership, screening 600+ applicants and facilitating host matching for PhD and undergraduate cohorts.

Amazon Ads

Applied Scientist Intern, NY

New York

May 2023 - Aug 2023

Project: Contextual Relevancy & LLM-Assisted Evaluation

- Designed novel methods for assessing annotator quality across multiple languages to improve Ground Truth data for ad relevance.
- Built fine-tuning pipelines (PyTorch, LangChain) incorporating varied levels of human feedback to align LLMs for generative AI tasks.
- Integrated annotator quality metrics into contextual relevancy pipelines, enhancing model robustness against noisy labels.

Meta (Facebook)

Machine Learning Research Engineer Intern, CA

Menlo Park

May 2022 - Aug 2022

Team: Facebook Creators Wellbeing

- Developed a Multi-Label Multi-Task model using PyTorch to assist page administrators in automated comment management.
- Outperformed existing production models by **40%** (ROC and PR AUC) by leveraging real-time user interaction data.
- Successfully deployed and evaluated the model in production via A/B testing on 4% of global Facebook traffic (millions of daily actions).

Rochester Institute of Technology

Research Assistant, NY

Rochester

2018 - 2024

Research: Fairness in AI & Human-in-the-loop

- Led large-scale human subjects research (budget \$8,000+) managing a workforce of 750+ annotators to build socially representative datasets.
- Developed methods for predicting human disagreement in subjective tasks (hate speech, offensive language) using BERT and probabilistic modeling.
- Published research at EMNLP, ACL, AAAI, and ECAI on creating privacy-centric, fair machine learning pipelines [1, 9].

LKEMR

Tech/Infrastructure Lead

Sri Lanka

2018-2022

Project: Cloud-based Electronic Medical Record (Open-EMR)

- Architected a web-based EMR system (PHP/MySQL on AWS) currently managing records for 1,000+ patients across Sri Lanka.
- Deployed the system as a critical tool for identifying potential COVID-19 patients during the pandemic; recognized as the primary widely-used EMR in the region.

EACL 2021 (Virtual Conference)

Infrastructure Lead

Virtual

Jan 2021 - July 2021

- Engineered the AWS cloud infrastructure for hosting the fully virtual EACL conference.
- Implemented Single-Sign-On (SSO) integration between external registration systems and AWS Cognito (OAuth2).

Boston College

Research Assistant Intern

Virtual

May 2021 - Aug 2023

- Investigated fairness in Automatic Speech Recognition (ASR) models for low-resource languages.
- Mentored REU students on fundamental Machine Learning concepts and research methodologies.

IBM AI Horizon (RPI)

Research Assistant, NY

Troy

May 2019 - Aug 2019

- Developed NLP pipelines (SpaCy, Doc2Vec) to aggregate nutritional data from 1M+ crowdsourced recipes.
- Constructed a food knowledge graph capable of generating FDA-compliant nutritional labels from unstructured recipe text.

Skills

Data Science: Tableau, Jupyter Notebooks (Colab), SQL (Presto) and Excel

Programming: Python, R, Java, C, and PHP

Infrastructure: Google Cloud, AWS, and DigitalOcean

ML/NLP: PyTorch, Keras, TensorFlow, SpaCy, NLTK, HuggingFace (BERT) and sklearn

Crowdsourcing: Qualtrics and Amazon Mechanical Turk

Design: Photoshop, Illustrator, and Lightroom

Conference Publications

- [1] Sujan Dutta, Deepak Pandita, **Tharindu Cyril Weerasooriya**, Marcos Zampieri, Christopher M. Homan, and Ashiqur R. KhudaBukhsh. "ARTICLE: Annotator Reliability Through In-Context Learning". In: *AAAI 2025*. Feb. 2025.
- [2] Shanilka Haturusinghe, **Tharindu Cyril Weerasooriya**, Christopher M Homan, Marcos Zampieri, and Sidath Ravindra Liyanage. "Subasa - Adapting Language Models for Low-resourced Offensive Language Detection in Sinhala". In: *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics (Student Research Workshop)*. Albuquerque, USA: Association for Computational Linguistics, Apr. 2025.

- [3] Isabelle Diana May-Xin Ng, **Tharindu Cyril Weerasooriya**, Haitao Zhu, and Wei Wei. “MultiGA: Leveraging Multi-Source Seeding in Genetic Algorithms”. In: *The First Workshop on Efficient Reasoning at NeurIPS 2025*. 2025.
- [4] Deepak Pandita, **Tharindu Cyril Weerasooriya**, Ankit Parag Shah, Isabelle Diana May-Xin Ng, Christopher M. Homan, and Wei Wei. “ProRefine: Inference-Time Prompt Refinement with Textual Feedback”. In: *The First Workshop on Efficient Reasoning at NeurIPS 2025*. 2025.
- [5] Mandira Sawkar, Samay U. Shetty, Deepak Pandita, **Tharindu Cyril Weerasooriya**, and Christopher M. Homan. “LPI-RIT at LeWiDi-2025: Improving Distributional Predictions via Metadata and Loss Reweighting with DisCo”. In: *Proceedings of the The 4th Workshop on Perspectivist Approaches to NLP*. Nov. 2025.
- [6] Jinghan Zhang, Fengran Mo, **Tharindu Cyril Weerasooriya**, Yeyang Zhou, Xinyue Ye, Dongjie Wang, Yanjie Fu, and Kunpeng Liu. “Blind Spot Navigation in Large Language Model Reasoning with Thought Space Explorer”. In: *The 5th Workshop on Mathematical Reasoning and AI at NeurIPS 2025*. 2025.
- [8] Deepak Pandita, **Tharindu Cyril Weerasooriya**, Sujan Dutta, Sarah K. Luger, Tharindu Ranasinghe, Ashiqur R. KhudaBukhsh, Marcos Zampieri, and Christopher M. Homan. “Rater Cohesion and Quality from a Vicarious Perspective”. In: *Findings of the Empirical Methods in Natural Language Processing: EMNLP 2024*. Nov. 2024.
- [9] **Tharindu Cyril Weerasooriya**, Sujan Dutta, Tharindu Ranasinghe, Marcos Zampieri, Christopher Homan, and Ashiqur KhudaBukhsh. “Vicarious Offense and Noise Audit of Offensive Speech Classifiers: Unifying Human and Machine Disagreement on What is Offensive”. In: *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Dec. 2023.
- [10] **Tharindu Cyril Weerasooriya**, Sarah Luger, and Christopher M. Homan. “Offensiveness as an Opinion: Dissecting Population-level Label Distributions”. In: *Proceedings of the 11th International Conference on Learning Representations*. 2023.
- [11] **Tharindu Cyril Weerasooriya**, Sarah Luger, Saloni Poddar, Ashiqur KhudaBukhsh, and Christopher M. Homan. “Subjective Crowd Disagreements for Subjective Data: Uncovering Meaningful CrowdOpinion with Population-level Learning”. In: *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Toronto, Canada: Association for Computational Linguistics, July 2023.
- [12] **Tharindu Cyril Weerasooriya**, Alexander G. Ororbia II, Raj Bhensadadia, Ashiqur KhudaBukhsh, and Christopher Homan. “Disagreement Matters: Preserving Label Diversity by Jointly Modeling Item and Annotator Label Distributions with DisCo”. In: *Findings of the Association for Computational Linguistics: ACL 2023*. Toronto, Canada: Association for Computational Linguistics, July 2023.
- [13] Christopher Homan, **Tharindu Cyril Weerasooriya**, Lora Aroyo, and Chris Welty. “Annotator Response Distributions as a Sampling Frame”. en. In: *Proceedings of the 1st Workshop on Perspectivist Approaches to NLP @LREC2022*. European Language Resources Association, 2022, p. 10.
- [14] **Tharindu Cyril Weerasooriya**, Alexander G Ororbia, and Christopher M Homan. “Improving Label Quality by Joint Probabilistic Modeling of Items and Annotators”. en. In: *Proceedings of the 1st Workshop on Perspectivist Approaches to NLP @LREC2022*. European Language Resources Association, 2022, p. 5.
- [15] **Weerasooriya, Tharindu Cyril** and Liu, Tong and Homan, Christopher M. “Neighborhood-based Pooling for Population-level Label Distribution Learning”. In: *Proceedings of the 24th European Conference on Artificial Intelligence*. 2020.
- [16] Kumara Mendis, **Tharindu Cyril Weerasooriya**, Supun Withana, Prabath Liyanage, Aruni Weerakoon Silva, Rajitha Wickramasinghe, and Chaminda Weerabaddana. “Cloud-Based Open Source Primary Care Electronic Patient Record System for Sri Lankan Citizens”. In: *2019 National Information Technology Conference, NITC 2019*. 2019, pp. 8–10.
- [17] **Tharindu Weerasooriya**, Nandula Perera, and S. R. Liyanage. “KeyXtract Twitter Model - An Essential Keywords Extraction Model for Twitter Designed using NLP Tools”. In: *10th KDU International Research Conference*. Ratmalana, 2017. arXiv: 1708.02912.

- [20] **Tharindu Weerasooriya**, Nandula Perera, and S. R. Liyanage. "A Method to Extract Essential Keywords from a Tweet using NLP Tools". In: *16th International Conference on Advances in ICT for Emerging Regions, ICTer 2016 - Conference Proceedings*. 2016, pp. 29–34.

Journal Papers

- [18] Nandula Perera and **Tharindu Weerasooriya**. "The 'Racecourse' of Then and Now: Evolution of the Sri Lankan English Vocabulary Over Two Generations of SLE Speakers". In: *VISTAS Journal* 10 (2016), pp. 1–23.
- [19] **Tharindu Weerasooriya**. "Automated Railway Ticketing System: Replacing the Paper based Ticket with the Electronic National Identity". In: *ICT for Development Working Paper Series* 6.1-2 (2016), pp. 67–77.

White Papers

- [7] Yujia Bao, Ankit Parag Shah, Neeru Narang, Jonathan Rivers, Rajeev Maksey, Lan Guan, Louise N. Barrere, Shelley Evenson, Rahul Basole, Connie Miao, Ankit Mehta, Fabien Boulay, Su Min Park, Natalie E. Pearson, Eldhose Joy, Tiger He, Sumiran Thakur, Koustav Ghosal, Josh On, Phoebe Morrison, Tim Major, Eva Siqi Wang, Gina Escobar, Jiaheng Wei, **Tharindu Cyril Weerasooriya**, Queena Song, Daria Lashkevich, Clare Chen, Gyuhak Kim, Dengpan Yin, Don Hejna, Mo Nomeli, and Wei Wei. *Harnessing Business and Media Insights with Large Language Models*. 2024.

Teaching

2023: Teaching Assistant for CSCI 736 - Neural Networks and ML (22 students)	<i>RIT</i>
2020: Adjunct Lecturer for CSCI 635 - Introduction to Machine Learning (~ 30 students)	<i>RIT</i>
2020-2021: Teaching Assistant for CSCI 235 - Mechanics of Programming (3 Semesters, ~ 120 students)	<i>RIT</i>
2018: Teaching Assistant for Programming in C (~ 130 students)	<i>University of Kelaniya</i>
2017: Teaching Assistant for Object Oriented Programming (~ 70 students)	<i>University of Kelaniya</i>
2017: Teaching Assistant for Data Structures and Algorithms (~ 70 students)	<i>University of Kelaniya</i>

Mentorship

2025: Isabelle Ng, BS Student (Intern Manager)	<i>UC Berkeley</i>
2025: Deepak Pandita, PhD Student (Intern Manager)	<i>Rochester Institute of Technology</i>
2024: Shanilka Haturusinghe, BS Student	<i>University of Kelaniya</i>
2022: Saloni Poddar, MS Student	<i>Rochester Institute of Technology</i>
2022: Yu Liang, MS Student	<i>Rochester Institute of Technology</i>
2022: Raj Bakulbhai Bhensadadia, MS Student	<i>Rochester Institute of Technology</i>
2021: Abhay Rajendra Dixit, MS Student	<i>Rochester Institute of Technology</i>
2020: Anuj Kulkarni, MS Student	<i>Rochester Institute of Technology</i>

Service

Paper Reviewer

2023-2025: ACL	2021: AISTATS '21
2021,2023,2024: EMNLP	2023: FAccT '23
2022: ICWSM '22	2019,2022: SocInfo '19,'22
2021-2023: IEEE Transactions on Neural Networks and Learning Systems	2022: WWW '23

Student Organizations.....

Doctoral Student Association **Rochester Institute of Technology**
President, NY *2020 - 2023*

- Represented the Ph.D. student population at sessions for the Board of Trustees of Rochester Institute of Technology.
- Co-led projects for fundraising and distribution of awards for student conference participation.
- Aided in advocating for student concerns, managing the budget and events for the organization, and the infrastructure for it.

HacKLN 2017 **University of Kelaniya**
Faculty Representative, SL *2017*

- Advised Students in organizing and fundraising for the hackathon.
- Guided the students as the teaching assistant in charge for the Department for Vidya 50, the 50th Celebratory Exhibition of Faculty of Science.

Gavel Club University of Kelaniya **University of Kelaniya**
Web Lead, SL *2014-2017*

- Organizing Committee Member of Inter University Best Speaker Contest 2014 - 2017.
- Designed, developed and maintained the official Club Website, Contest Website and the online account management system for the executive committee.

Hiring Committees.....

2022: Search Committee for Dean for Graduate School *Rochester Institute of Technology*

2022: Search Committee for Dean for College of Computing *Rochester Institute of Technology*

2021: Search Committee for Assistant Dean of Graduate School *Rochester Institute of Technology*

Invited Talks

2025: VOICE: Dissecting the Opinion of the Politcal Stance *University at Albany*

2023: CrowdOpinion: Dissecting the Opinion of the Crowd *Rochester Institute of Technology*

2023: CrowdOpinion: Dissecting the Opinion of the Crowd *Internal Amazon HCI Workshop*

2023: Fairness in Human Annotation *Rochester Institute of Technology*

2022: Auditing Machine Learning Classifiers on Vicarious Offense *Internal Google Disagreement Workshop*

Awards

2021: AWS Startup credit for LKEMR (\$1000) *AWS*

2020: Google Cloud Research Credit (\$1000) *Google*

2020: Crowdfunding Campaign for DSA (raised \$1000) *Crowdfunded*

Certificates

2018,2023: Social and Behavioural Research Certificate *Citi Program*

2017: Competent Gavelier Leader *Toastmasters International USA*

2017: Competent Gavelier Communicator *Toastmasters International USA*