Adeep Hande

adeephande1000@gmail.com | Google Scholar | LinkedIn | Website | Washington, D.C.

Research Interests

My research centers on making large language models efficient and reliable enough to personalize at production scale. This spans three threads: (1) **LLM post-training**, including supervised fine-tuning, DPO, and knowledge distillation to build compact, deployable models from larger teachers; (2) **generative and sequential recommendation**, using LLM-generated synthetic data to address cold-start and long-tail recommendation; and (3) **efficient LLM inference**, building serving infrastructure that makes LLM-scale personalization affordable in production. I am interested in continuing this line of work during a PhD, with a particular focus on efficient post-training and inference methods.

Education

Indiana University Bloomington	M.S. in Data Science (GPA 3.8)	May 2024
Indian Institute of Information Technology, Tiruchirappalli	B.S. in Computer Science (GPA 3.3)	Apr 2022

Research & Industry Experience

ML Research Engineer, Comcast Applied AI Research — Washington, D.C.

Recommendation and Arbitration team · May 2024 – Present

- **LLM-driven audience segment discovery at scale.** Built and deployed a teacher-student LLM post-training pipeline (reasoning-aware SFT, DPO, and off-/on-policy evaluation) to generate and operationalize latent sports-audience segments across millions of users; achieved production-scale serving at 0.6% of enterprise-model cost through model distillation (NAACL 2027).
- **LLM data augmentation for cold-start recommendation.** Designed an LLM-based data-augmentation pipeline that injects synthetic user histories for cold-start items, achieving an 80% relative NDCG improvement for cold-start items while maintaining overall recommendation quality (CIKM 2026, EACL 2027).
- **Efficient inference at scale.** Built a self-hosted vLLM serving setup (4×A100) reaching ~326k tokens/sec for LLM-generated user-profile inference across 5–10M users, roughly 650× cheaper than managed-API inference at that scale.
- **Content discoverability simulation.** Architected a large-scale simulation using LLM-generated synthetic queries and historical clustering to evaluate the discoverability of upcoming content in the Xfinity TV ecosystem (RecSys 2025).
- **Forecasting infrastructure.** Developed campaign-level time-series forecasting models and ETL pipelines (Spark, Databricks, SQL) for Xfinity Campaign Manager across EU and NA markets, with automated dashboarding for weekly media-buy pacing and budget allocation.

Machine Learning Graduate Research Intern, Comcast — Washington, D.C.

May 2023 – Dec 2023

- **ASR evaluation.** Built an end-to-end speech evaluation framework benchmarking ASR quality (WER, embedding-based metrics) across diverse user demographics and acoustic conditions, scaling annotation throughput from 1K/day to 100K/week via automated proxy metrics (GEM @ ACL 2026).
- **Search quality evaluation.** Deployed a hybrid heuristic–LLM inference system for real-time moderation of 5M+ daily voice queries, flagging biased or policy-violating results; introduced a time-aware dynamic lexicon that reduced false positives by 35% (SIGIR 2025).

Publications

Under Review

1. “From Head to Tail: LLM-Augmented Sequential Recommendation and Trade-off Optimization” — **Adeep Hande**, Shaghayegh Agah, Neeraj Sharma, Yejin Kim, Karina H. Halevy, Sardar Hamidian — *EACL 2027*
2. “LLM-Driven Segment Discovery: Uncovering Latent User Intent at Scale” — Shaghayegh Agah*, **Adeep Hande***, Neeraj Sharma (*equal contribution) — *NAACL 2027, Industry Track*
3. “Candidate Grounding Matters: LLM-Guided Negative Selection for Production Recommenders” — Yejin Kim, **Adeep Hande**, Shaghayegh Agah, Neeraj Sharma, Sardar Hamidian — *NAACL 2027, Industry Track*

2026

1. “From Head to Tail: LLM-Augmented Sequential Recommendation and Trade-off Optimization” — **Adeep Hande**, Shaghayegh Agah, Neeraj Sharma, Karina Halevy, Sardar Hamidian — *GRAIL Workshop @ CIKM 2026*
2. “Evaluating ASR Quality at Scale on TV Entertainment Platforms” — **Adeep Hande**, Kishorekumar Sundararajan, Yidnekachew Endale, Akshatha Bapu KrishnaSwamy, Sachin Dabral, Dawn Reed, Michael Pereira — *GEM Workshop @ ACL 2026*

2025

1. “Simulating Discoverability for Upcoming Content on Entertainment Platforms” — **Adeep Hande**, et al. — *RecSys 2025, Prague*
2. “Content Moderation in TV Search: Balancing Compliance, Relevance, and User Experience” — **Adeep Hande**, et al. — *SIGIR 2025*
3. “Do Images Really Do the Talking? Analyzing the Significance of Images in Tamil Troll Meme Classification” — Siddhanth U. Hegde, **Adeep Hande**, Ruba Priyadharshini, Sajeetha Thavareesan, Ratnasingam Sakuntharaj, S. Thangasamy, B. Bharathi, Bharathi Raja Chakravarthi — *Advances in Computational Intelligence (Springer), 2025*

2022

1. “Multi-task Learning in Under-resourced Dravidian Languages” — **Adeep Hande**, Siddhanth U. Hegde, Bharathi Raja Chakravarthi — *Journal of Data, Information and Management (Springer)*
2. “How Can We Detect Homophobia and Transphobia? Experiments in a Multilingual Code-Mixed Setting for Social Media Governance” — Bharathi Raja Chakravarthi, **Adeep Hande**, Rahul Ponnusamy, Prasanna Kumar Kumaresan, Ruba Priyadharshini — *International Journal of Information Management Data Insights (Elsevier)*
3. “The Best of Both Worlds: Dual Channel Language Modeling for Hope Speech Detection in Low-resourced Kannada” — **Adeep Hande**, Siddhanth U. Hegde, Sivanesan Sangeetha, Ruba Priyadharshini, Bharathi Raja Chakravarthi — *LT-EDI Workshop @ ACL 2022*
4. “DE-ABUSE@TamilNLP-ACL 2022: Transliteration as Data Augmentation for Abuse Detection in Tamil” — Vasanth Palanikumar, Sean Benhur, **Adeep Hande**, Bharathi Raja Chakravarthi — *DravidianLangTech Workshop @ ACL 2022*
5. “Findings of the Shared Task on Multi-task Learning in Dravidian Languages” — Bharathi Raja Chakravarthi, Ruba Priyadharshini, C. N. Subalalitha, Sangeetha Sivanesan, Malliga Subramanian, Kogilavani Shanmugavadivel, Parameswari Krishnamurthy, **Adeep Hande**, Siddhanth U. Hegde, Roshan Nayak, et al. — *DravidianLangTech Workshop @ ACL 2022*
6. “Findings of the Second Shared Task on Troll Meme Classification in Tamil” — Shardul Suryawanshi, Bharathi Raja Chakravarthi, Mihael Arcan, Susan Levy, Paul Buitelaar, Prasanna Kumar Kumaresan, Rahul Ponnusamy, **Adeep Hande** — *DravidianLangTech Workshop @ ACL 2022*
7. “Findings of the Shared Task on Emotion Analysis in Tamil” — Sean Benhur, et al. (incl. **Adeep Hande**) — *DravidianLangTech Workshop @ ACL 2022*

2021

1. “Domain Identification of Scientific Articles Using Transfer Learning and Ensembles” — **Adeep Hande**, Karthik Puranik, Ruba Priyadharshini, Bharathi Raja Chakravarthi — *PAKDD 2021 (Springer)*
2. “Evaluating Pretrained Transformer-based Models for COVID-19 Fake News Detection” — **Adeep Hande**, Karthik Puranik, Ruba Priyadharshini, Sajeetha Thavareesan, Bharathi Raja Chakravarthi — *ICCMC 2021 (IEEE)*
3. “Offensive Language Identification in Low-resourced Code-mixed Dravidian Languages Using Pseudo-labeling” — **Adeep Hande**, Karthik Puranik, Konthala Yasaswini, Ruba Priyadharshini, Sajeetha Thavareesan, Anbukkarasi Sampath, Kogilavani Shanmugavadivel, Durairaj Thenmozhi, Bharathi Raja Chakravarthi — *arXiv preprint*
4. “IIITT@LT-EDI-EACL2021 – Hope Speech Detection: There Is Always Hope in Transformers” — Karthik Puranik, **Adeep Hande**, Ruba Priyadharshini, Sajeetha Thavareesan, Bharathi Raja Chakravarthi — *LT-EDI Workshop @ EACL 2021*
5. “UVCE-IIITT@DravidianLangTech-EACL2021: Tamil Troll Meme Classification: You Need to Pay More Attention” — Siddhanth U. Hegde, **Adeep Hande**, Ruba Priyadharshini, Sajeetha Thavareesan, Bharathi Raja Chakravarthi — *DravidianLangTech Workshop @ EACL 2021*
6. “IIITT@DravidianLangTech-EACL2021: Transfer Learning for Offensive Language Detection in Dravidian Languages” — Konthala Yasaswini, Karthik Puranik, **Adeep Hande**, Ruba Priyadharshini, Sajeetha Thavareesan, Bharathi Raja Chakravarthi — *DravidianLangTech Workshop @ EACL 2021*

7. “IIITT at CASE 2021 Task 1: Leveraging Pretrained Language Models for Multilingual Protest Detection” — Pawan Kalyan, Duddukunta Reddy, **Adeep Hande**, Ruba Priyadharshini, Ratnasingam Sakuntharaj, Bharathi Raja Chakravarthi — *CASE Workshop @ ACL-IJCNLP 2021*
8. “Attentive Fine-tuning of Transformers for Translation of Low-resourced Languages @LoResMT 2021” — Karthik Puranik, **Adeep Hande**, Ruba Priyadharshini, Thenmozhi Durairaj, Anbukkarasi Sampath, Kingston Pal Thamburaj, Bharathi Raja Chakravarthi — *LoResMT Workshop @ MT Summit 2021*
9. “Hypers at ComMA@ICON: Modelling Aggressive, Gender Bias and Communal Bias Identification” — Sean Benhur, Roshan Nayak, Kanchana Sivanraju, **Adeep Hande**, Subalalitha Chinnaudayar Navaneethakrishnan, Ruba Priyadharshini, Bharathi Raja Chakravarthi — *ICON 2021 (Shared Task)*
10. “Findings of the Sentiment Analysis of Dravidian Languages in Code-Mixed Text” — Bharathi Raja Chakravarthi, Ruba Priyadharshini, Sajeetha Thavareesan, Dhivya Chinnappa, Durairaj Thenmozhi, Elizabeth Sherly, John P. McCrae, **Adeep Hande**, Rahul Ponnusamy, Shubhanker Banerjee, et al. — *arXiv preprint*

2020

1. “KanCMD: Kannada CodeMixed Dataset for Sentiment Analysis and Offensive Language Detection” — **Adeep Hande**, Ruba Priyadharshini, Bharathi Raja Chakravarthi — *PEOPLES Workshop @ COLING 2020*

Patents

1. Systems for Content Filtering. Comcast, filed 2025.
2. Systems for Enhancing Content Discoverability in TV Systems. Comcast, filed 2025.

Academic Service

Workshop Organizing

- Co-organizer, Workshop on Multimodal Machine Learning in Low-Resource Languages (MMLow), co-located with ICON 2022.

Program Committee Membership

- Program committee member for shared-task workshops at EACL 2021 and ACL 2022 (DravidianLangTech / LT-EDI workshop series).

Peer Review

- 100+ peer reviews across academic venues, including ACL Rolling Review, ACL GEM, InterSpeech, ICASSP, SIGIR, COLING, DravidianLangTech, ACM WWW (main conference), ACM Multimedia (main conference), IEEE SLT workshops, and the KDD Datasets & Benchmarks Track.

Technical Skills

Machine Learning & NLP: Large Language Models, LLM Fine-Tuning (DPO, SFT, RLHF, Knowledge Distillation), Recommendation Systems, Ranking, Content Moderation, Retrieval-Augmented Generation, Bias & Fairness, ASR Evaluation, Information Retrieval

Frameworks & Libraries: PyTorch, TensorFlow, Hugging Face Transformers, vLLM, LangChain, OpenAI API, MLflow, Airflow

Data & Infrastructure: Spark, Databricks, AWS (SageMaker, Lambda, EC2), Docker, FastAPI, MongoDB, Kafka

Languages: Python, SQL

References

Available upon request.