

Tarun Kumar

✉ kumartaruncse@gmail.com

☎ 9080696732

PROFILE SUMMARY

Senior Research Scientist and technologist with a Ph.D. in Computer Science (IIT Madras, Intel India Ph.D. Fellow) and experience in applied research spanning Generative AI, Large Language Models, agentic AI systems, and NLP. Author of 12+ peer-reviewed papers at venues including ICML, ACL and AAAI, and inventor on 18 filed patents covering Generative AI, Gen AI infrastructure, LLM agent guardrails, and retrieval-augmented generation. Track record of defining research direction and translating novel algorithms into production-grade, enterprise-scale AI systems at HPE Labs.

TOP SKILLS

- **Generative AI & LLMs:** RAG architecture & optimization, prompt orchestration, foundation model evaluation/selection, foundation model virtualization
- **Agentic AI Systems:** policy-driven guardrails, autonomous agent evaluation, reliable tool use, MCP-based infrastructure agents
- **NLP:** embedding models, hallucination mitigation, knowledge graph completion, fairness-aware data generation
- **Data-Centric & Trustworthy AI:** dataset valuation & lineage, data discovery, model explainability
- **Time-Series & Multimodal AI:** time-series foundation model adaptation & anomaly detection, vision-language generation
- **Research Leadership:** research roadmap definition, cross-functional collaboration, mentoring, patent portfolio development

PROFESSIONAL EXPERIENCE

Senior Research Scientist: Hewlett Packard Enterprise

Jan 2025 - present

- Working on Agentic AI systems and foundation models for enterprise and scientific workloads, with an emphasis on Data organization, LLM agent guardrails, and model selection pipelines.

Research Scientist: Hewlett Packard Enterprise

Dec 2021 - Dec 2024

- Worked on data discovery in the realm of Data-Centric and Trustworthy AI.

Project Associate: Robert Bosch Center for Data Science AI, IIT Madras

June 2018 - Dec 2021

- Network Representation Learning in collaboration with Intel Labs

Research Intern, Landis & Gyr (Toshiba Corp.)

July 2014-May 2015

- Implementation of Java based DCW interpreter
- Porting EmbOS and EmFile on Volta board

EDUCATION

Indian Institute of Technology Madras

CGPA – 8.5/10

*Ph.D. (**Intel PhD Fellowship**), Computer Science and Engineering*

2016 – 2022

Best Ph.D. Thesis in Data Science, IIT Madras

Advisors: Balaraman Ravindran and Manikandan Narayanan

Thesis: Algorithms For Rich Graph Structures
Media Coverage: Times of India, India Today, Business Standard

Shiv Nadar University

M.Tech., Department of Computer Science and Engineering

CGPA – 9.07/10

2013 – 2015

Distinguished Alumni Award, 2026

Advisor: Rajeev Kumar Singh

Thesis: Analysis of Compiler Optimization Techniques for Multi Core Architecture using Feature Mining

Punjab Technical University

B.Tech., Department of Computer Science and Engineering

Percentage – 85.3/100

2009 – 2013

PATENTS (Filed)

2026

20. An activation geometry-based trap detection and hierarchical escalation mechanism in foundation models
19. LLM-powered Policy-based Guardrails for MCP-Driven Infrastructure Management
18. Autonomous Evaluation Framework for AI Agents in Complex Infrastructure Management Environments
17. CompactRAG: Criticality-Aware Tiered Compression of Vector Embeddings for Retrieval-Augmented Generation Systems
16. Cross-dataset meta-learning for metric forecasting with limited historical data
15. ConceptRAG: Semantically Enriched Nearest Neighborhood Clustering for Retrieval Augmented Generation
14. VERA-Lead: Verifiable Early Risk Anticipation and Driver-to-Impact Diagnosis for Multivariate IT Operations
13. REVID-AD: Graph-Ordered Evidence Videos for Range-Aware Multivariate Time-Series Anomaly Detection

2025

12. A novel architectural paradigm for virtualizing Foundation Models inspired by Operating System principles
11. Adaptive Semantic Hierarchical Clustering with User-Guided Partial Hierarchies for Object Stores
10. A system to add hierarchical navigation in an object storage platform using concept graphs
9. System and Method for Co-optimizing the recommendations and evaluations of Foundation Models (FMs) using metadata characteristics and performance benchmarks
8. A framework to create hierarchical navigation in object store by connecting the real-world categories to extracted domain concepts
7. Augmenting Time Series Foundation Models with Pattern and Context-Aware Adaptation

6. Forecast2Anomaly (F2A): Adapting Multivariate Time Series Foundation Models for Anomaly Prediction
5. Pre-Action Reasoning and Prompt-Based Control with Metaprompting for Reliable Tool Use in LLM Agents
4. Dual Space VisionTS: Vision-Inspired Dual-Space Modeling for Multivariate Few-Shot Time Series Forecasting

2024

3. Graph-based dataset valuation to solve AI problems
2. AI Pipeline Metadata Knowledge Graph Completion using Pretrained Large Language Models
1. Demystifying Black Box Time-Series Forecasting: Post Hoc Shape-based Perturbation Analysis for Explainability and Insights

PUBLICATIONS (Google Scholar) ([J]ournal, [C]onference, [W]orkshop)

o 2026

- [C8] Suparna Bhattacharya, **Tarun Kumar**, Cong Xu, Satish Kumar Mopur, Jiahao Li, Ashish Mishra, Aalap Tripathy, Annmary Justine Koomthanam, Martin Foltin and Ian Foster, “It is Time to Virtualize Foundation Models with a Self-evolving Operating System Layer” accepted to be published at 43rd International Conference of Machine Learning (ICML) 2026.
- [C7] [Top 50% accepted] Ashish Mishra, **Tarun Kumar**, Arpit Shah, Suparna Bhattacharya, Martin Foltin “CEBC: Conformal Evidence-Bounded Control for Low-Hallucination Vision–Language Generation” accepted for publication in ACL 2026.
- [C6] **Tarun Kumar**, Aalap Tripathy, Gayathri Saranathan, Suparna Bhattacharya, Martin Foltin, Scott Hinchley, Donald M Bahls, David Brookshire, ““InfrastructureSentinel: Policy Enforced Guardrails for Secure MCP-driven Infrastructure Agents.” In Proceedings of the AAAI Conference on Artificial Intelligence, vol. 40, no. 47, pp. 40295-40301. 2026.”
- [W9] Gayathri Saranathan, Aalap Tripathy, Tarun Kumar, Scott Hinchley, Martin Foltin, Christopher L Holmes, David Brookshire, Donald M Bahls, Cong Xu, Robert W. Wisniewski, Larry Kaplan, Suparna Bhattacharya, “Beyond Curated Benchmarking: Automated Evaluation of LLM Agents for Safe and Reliable IT Infrastructure Management” (AAAI 2026 Workshop on Agentic AI Benchmarks and Applications for Enterprise Tasks)
- [J6] Fernando Perez-Cruz, Briland Hitaj, Giovanna Maria Dimitri, Oishi Deb, Amine M’Charrak, Shikhar Srivastava, Christopher Kanan, Anup Kumar Gupta, Zefang Liu, **Tarun Kumar**, Álvaro Galisteo Bermúdez, Lisa Alazraki, Vasilios Mavroudis, Yuyang Xue, Brian Matejek, Swati Rajwal, Avinash Kumar Pandey, Ayyüce Begüm Bektaş, Anil B Murthy, Feng Chen et. al., “The Unreasonable Benchmark”, Journal of Data-centric Machine Learning Research (DMLR), 2026.

o 2025

- [C5] [Best Theme Paper award winner] Gayathri Saranathan, Cong Xu, Mahammad Parwez Alam, **Tarun Kumar**, Martin Foltin, Soon Yee Wong, and Suparna Bhattacharya, “SubLIME: Subset Selection via Rank Correlation Prediction for Data-Efficient LLM Evaluation”, accepted for publication at ACL 2025 (Oral presentation + panel invitation – top 1% in the accepted papers).
- [J5] Revathy Venkaratamanan, Aalap Tripathy, **Tarun Kumar**, Annmary Justine, Arpit Shah, Sergey Serebryakov, Martin Foltin, Suparna Bhattacharya, Paolo Faraboschi and Amit Sheth

“Constructing a Metadata Knowledge Graph as an atlas for demystifying AI Pipeline optimization.” *Frontiers in Big Data Journal*, 2025.

- [W8] Gayathri Saranathan, Martin Foltin, Cong Xu, Suparna Bhattacharya, Annmary Justine, **Tarun Kumar**, Arpit Shah, Aalap Tripathy, Rishab Sharma, Shreyas Kulkarni, Maxim Ziatdinov, Ayana Ghosh, Kevin Roccapiore, Paolo Faraboschi, Ian Foster, “Advancing Autonomous Microscopy Agents with domain guided dynamic retrieval in a Virtual Foundation Model OS” (accepted for presentation in Harnessing the Power of various forms of Generative AI for Science and Engineering: The Trillion Parameter Consortium (TPC), 2025)
- [W7] Mahammad Parwez Alam, Gayathri Saranathan, Cong Xu, Javier Aula-Blasco, Martin Foltin, **Tarun Kumar**, Soon Yee Wong, and Suparna Bhattacharya “SubLIME*: Data Efficient Foundation Model Evaluation across Modalities, Languages and Benchmarks.”, ICLR 2025 Workshop on Navigating and Addressing Data Problems for Foundation Models.
- [W6] **Tarun Kumar**, Cong Xu, Arpit Shah, Baradji Diallo, Martin Foltin, Suparna Bhattacharya “Co-optimizing Recommendation and Evaluation for LLM Selection”, ICLR 2025 Workshop on Foundation Models in the Wild
- [W5] Shubhadip Nag, Srinjoy Das, Agniva Saha, Anushree Ghosh, Soumi Das, **Tarun Kumar**, Suparna Bhattacharya, Sourangshu Bhattacharya, “MPSelectTune: Prompt Selection for Fine-tuning improves Concept Unlearning in LLMs”, Reliable ML workshop at NeurIPS 2025
- [W4] Ashish Mishra, Tarun Kumar, Satish Kumar Mopur, Sergey Serebryakov, Suparna Bhattacharya, Martin Foltin, Ramanagopal Vogety, Phanidhar Koganti, “Adaptive Refinement of Time Series Foundation Models via Pattern and Context-Awareness” BERT2S workshop at NeurIPS 2025

○ 2024

- [W3] Ashish Mishra, Gyanaranjan Nayak, Suparna Bhattacharya, **Tarun Kumar**, Arpit Shah, Martin Foltin, “LLM-Guided Counterfactual Data Generation for Fairer AI” at DCAI workshop at The Web Conference (WWW 2024), DMLR workshop at ICLR 2024.
- [C4] Cong Xu, **Tarun Kumar**, Martin Foltin, Annmary Justine, Sergey Serebryakov, Arpit Shah, Ashish Mishra, Gyanaranjan Nayak, Suparna Bhattacharya, Paolo Faraboschi “System for Recommendation and Evaluation of Large Language Models for practical tasks in Science.” CUG Conference 2024.

○ 2023

- [W2] **Tarun Kumar**, Arpit Shah, Ashish Mishra, Suparna Bhattacharya, Arun Mahendran, Ted Dunning, Glyn Bowden “From Roots to Fruits: Exploring Lineage for Dataset Recommendations.” Data Economy workshop (2023), SIGMOD.
- [C3] Naganand Yadati, **Tarun Kumar**, Deepak Maurya, Partha Talukdar, and Balaraman Ravindran. “HEAL: Embedding Multi-layer Hypergraphs.” *Learning on Graphs (LOG)* 2023.

○ 2022

- [W1] Naganand Yadati, **Tarun Kumar**, Deepak Maurya, Partha Talukdar, and Balaraman Ravindran. “HEAL: Embedding Multi-layer Hypergraphs.” *Deep Learning on Graphs: Method and Applications (DLG-AAAI’22)* workshop at AAAI 2022.
- [J4] Aditya Jadhav, **Tarun Kumar***, Mohit Raghavendra*, Tamizhini L., and Manikandan Narayanan. “Predicting cross-tissue hormone-gene relations using balanced word embeddings.” *Bioinformatics* (2022).

- [J3] **Tarun Kumar**, Ramanathan Sethuraman, Sanga Mitra, Balaraman Ravindran and Manikandan Narayanan. (2023) “MultiCens: Multilayer network centrality measures to uncover molecular mediators of tissue-tissue communication.” (PLOS Computational Biology, 19(4), e1011022.)

○ **2020**

- [C2] **Tarun Kumar***, K Darwin*, Srinivasan Parthasarathy, and Balaraman Ravindran. 2020. “HPRA: Hyperedge Prediction using Resource Allocation.” In 12th ACM Conference on Web Science (WebSci 2020). Association for Computing Machinery, New York, NY, USA, 135–143.
- [J2] **Tarun Kumar**, Sankaran Vaidyanathan, Harini Ananthapadmanabhan, Srinivasan Parthasarathy, and Balaraman Ravindran. “Hypergraph Clustering by Iteratively Reweighted Modularity Maximization.” Applied Network Science 5, no. 1 (2020): 1-22.

○ **2019**

- [J1] **Tarun Kumar**, Manikandan Narayanan, and Balaraman Ravindran. “Effect of Inter-layer Coupling on Multilayer Network Centrality Measures.” Journal of the Indian Institute of Science: 99, 237–246 (2019).
- [C1] **Tarun Kumar***, Sankaran Vaidyanathan*, Harini Ananthapadmanabhan, Srinivasan Parthasarathy, and Balaraman Ravindran. “A New Measure of Modularity in Hypergraphs: Theoretical Insights and Implications for Effective Clustering.” In International Conference on Complex Networks and Their Applications. pp. 286-297. Springer, Cham, 2019.

AWARDS AND ACHIEVEMENTS

- Distinguished Alumni Award, Shiv Nadar University, 2026
- HPE Labs Star Award, 2026, 2025.
- IP Day awards: Most prolific inventors list for filing 11 patents in a year
- HPE Hybrid Cloud and Office of CTO Champion award, 2025
- “Innovator of the year” award by HPE, 2024.
- **Best Ph.D. thesis award** in Data Science at IIT Madras, 2023.
- Awarded **Intel India PhD fellowship** for year 2020-2021.

TALKS/PRESENTATIONS

- irtual Foundation Model as an OS for Generative AI, HPE TechCon 2026.
- “LATTICE: A data intelligence platform for Private Cloud for AI”, HPE TechCon 2025
- LLM recommendation and Evaluation for Scientific use cases (CUG 2024, Australia)
- LLM-Guided Counterfactual Data Generation for Fairer AI (DCAI 2024, Singapore)
- “Holistic Evaluation and Recommendation of LLMs” HPE TechCon 2024.
- “MultiCens: Multilayer network centrality measures to uncover molecular mediators of tissue-tissue communication,” IMSc Network Biology Day, 2023
- “Surprising Data Scientists by Data Discovery.” HPE TechCon, 2023.

- “Hypergraph clustering by modularity maximization,” Poster presentation and a short talk, RBCD-SAI’s international summit on Data Science and AI, June 18, 2020.
- “Algorithms for hypergraphs: clustering and hyperedge prediction,” University of Texas Dallas Outreach Event, May 20, 2020.
- “Centrality in multilayer networks,” 6th IBSE workshop, IIT Madras, Nov 2, 2019.
- “Hypergraph Clustering: A Modularity Maximization Approach,” RBCDSAI Webscience trust conference on data science and AI, IIT Madras, Feb 25, 2019.
- “Representation Learning in Hypergraphs,” Second RBC-DSAI workshop on Recent Progress in Data Science and AI, IIT Madras, April 18, 2018.

SERVICES

- Co-organizer for AgenticOS workshop at NeurIPS 2026
- Co-organizer for Graphs and more Complex structures for Learning and Reasoning (GCLR) workshop series at AAAI 2021, AAAI 2022, AAAI 2023, AAAI 2024, AAAI 2026.
- Area Chair: CODS-COMAD 2026
- Journals reviewer: Nature Scientific Reports, PLOS ONE, IEEE Transactions on Network Science and Engineering, Complexity, Applied Network Science, Journal of Statistical Mechanics: theory and experiment, Social Network Analysis and Mining journal, Current Bioinformatics, Artificial Intelligence Review, The Journal of Supercomputing
- Program Committee member: NeurIPS (2022-26), ICML (2024-26), ICLR (2022, 2024-26), AAAI (2025-27), KDD (2025)
- Reviewer: HPE TechCon 2025, 2026.
- Co-Editor: (Special Issue) Focus on Learning and Reasoning in Generalized Graphs, Journal of Physics: Complexity, 2023.
- Member (2024 -): Industry and University Relations co-ordinator at India R&D Innovation Stream (IRIS) at HPE.
- Membership of professional bodies: Association for the Advancement of Artificial Intelligence, ACM India Special Interest Group on Knowledge Discovery and Data Mining