

Manish Munikar

☎ (817) 917-1024 | ✉ munikarmanish@gmail.com | 🏠 munikarmanish.github.io | 📱 munikarmanish | 🌐 munikarmanish

Summary

I'm a Computer Science PhD candidate at UT Arlington (supervised by Dr. Jia Rao) researching on container network optimizations in Linux. I also have industry experience in applied machine learning. I am actively seeking a research scientist or software engineering role.

Education

University of Texas at Arlington, PhD in Computer Science	CGPA: 4.0	2020 – 2025 (expected)	Arlington, TX, USA
Tribhuvan University, Bachelors in Computer Engineering	Grade: 80%	2013 – 2017	Kathmandu, Nepal

Research Publications

- J. Lei, **M. Munikar**, H. Lu and J. Rao, "Accelerating Packet Processing in Container Overlay Networks via Packet-level Parallelism," in *37th IEEE International Parallel and Distributed Processing Conference (IPDPS 2023)*.
- M. Munikar**, J. Lei, H. Lu and J. Rao, "PRISM: Streamlined Packet Processing for Containers with Flow Prioritization," In *Proceedings of the 42nd IEEE International Conference on Distributed Computing Systems (ICDCS 2022)*.
- J. Lei, **M. Munikar**, K. Suo, H. Lu and J. Rao, "Parallelizing Packet Processing in Container Overlay Networks," In *Proceedings of the Sixteenth European Conference on Computer Systems (EuroSys 2021)*.
- M. Munikar**, S. Shakya and A. Shrestha, "Fine-grained Sentiment Classification using BERT," *2019 IEEE International Conference on Artificial Intelligence for Transforming Business and Society (AITB 2019)*.
- P. Dhakal, **M. Munikar** and B. Dahal, "One-Shot Template Matching for Automatic Document Data Capture," *2019 IEEE International Conference on Artificial Intelligence for Transforming Business and Society (AITB 2019)*.

Professional Experience

Applied Scientist Intern | Amazon, Bellevue, WA, USA | Python, ML, AWS, GNN, Label propagation May 2024 – Aug 2024

- Improved the F1-score of heavily imbalanced classification problems with limited labels by up to 15% without any more labeled data.
- Experimented with data-cleaning and semi-supervised learning algorithms, e.g., confident learning, graph-based label propagation.

Applied Scientist Intern | Amazon, Sunnyvale, CA, USA | Python, ML, Spark, LSH May 2022 – Aug 2022

- Developed ML models to detect anomalies in human-labeled datasets with over 80% precision.
- Experimented with multi-modal embeddings and approximate nearest neighbors algorithms such as locality-sensitive hash (LSH).

Applied Scientist Intern | Amazon, (Remote) | Python, ML, Spark May 2021 – Aug 2021

- Developed ML models using multi-modal embeddings to detect global inconsistencies in the Amazon catalog with up to 94% precision.
- Worked with massive Amazon product datasets using distributed frameworks like Spark.

Data Scientist | Docsumo, Kathmandu, Nepal | Python, SQL, ML, CV, NLP, Object-detection, OpenCV Jul 2018 – Dec 2019

- Developed computer vision object-detection models (Faster R-CNN, YOLO, SSD) and text-based models to identify key values in documents.
- Developed a novel template-based one-shot algorithm to extract structured information from document images with over 90% accuracy.
- Experienced all stages of ML project: data collection/annotation, model development, evaluation, and production-ready deployment.
- Developed predictive analysis systems to detect anomalies in time-series data.

Software Developer | LIS Nepal Pvt. Ltd., Lalitpur, Nepal | SQL query-optimization, ETL, Hadoop, Hive Oct 2017 – Jun 2018

- Developed business intelligence (BI) reports for global retail enterprises using large-scale optimized SQL queries.
- Wrote data integration (ETL) scripts using big data technologies (Hadoop, Hive, Sqoop, Flume).

Research & Project Experience

ECON: Expedited Container Overlay Network | C/C++, Linux kernel network stack, eBPF/XDP Ongoing

- Designed a fastpath in the Linux network stack to accelerate the processing of container overlay network packets in a backward-compatible way. ECON improves container throughput by up to 121%, reduces average latency by up to 61%, and reduces CPU usage for packet processing by up to 43%. ECON accelerates both TCP and UDP container flows, without requiring any application or hardware modifications.

PRISM: Priority-based Packet Scheduling for Overlay Network | C/C++, Linux kernel network stack, Priority differentiation 2022

- Implemented *priority-based* container overlay packet processing in the Linux kernel to improve the latency (by over 50%) and throughput (by up to 100%) of high-priority container flows in the presence of low-priority best-effort traffic.

FALCON: Parallelized Packet Processing for Container Overlay Network | C/C++, Linux kernel network stack 2021

- Added device-level packet steering in the Linux network stack to parallelize single-flow packet processing. FALCON also balances softirqs across cores based on load to improve container throughput by up to 300% and reduce tail latency by up to 53%.

Movie Review Mining and Recommendation System 📺 | Python, Django, SQL, ML, NLP, RNTN, Recommender system Aug 2017

- A web application that analyzes movie reviews' sentiments using deep learning (RNTN) and builds a collaborative-filtering recommender system on top of it. Users provide movie reviews and get personalized movie recommendations in return. Built using Python and NumPy.