

Rohit Kumar Salla

Education

Virginia Polytechnic Institute and State University - Blacksburg, Virginia
Master of Science in Computer Engineering August 2025 - May 2027

SRM Institute of Science and Technology - India
Bachelor of Technology in Computer Science August 2021 – May 2025
GPA: 3.8/4.0

Publications

Conference Papers

- **AURORA: Neuro-Symbolic Continual Indexing for Evolving RAG Systems**
Association for Computational Linguistics (ACL) 2026 — Accepted
- **CrossQ: Task-Aligned Cross-Token Conditional Quantization for Late Interaction Retrieval**
The International Conference on Machine Learning (ICML) 2026 — Accepted
- **Learning When to Stop: Prefix-Optimal Dynamic Diffusion Policies for Continuous Control**
Reinforcement Learning Conference 2026 (RLC) — Accepted
- **When Can Pure Exploitation Succeed in Linear RL? Decoys and Self-Identifiability for Greedy LSVI**
Reinforcement Learning Conference 2026 (RLC) — Accepted
- **Multimodal Patent Retrieval: Integrating Text and Images with Long-Tail Optimization Techniques**
ICCCNT 2025 — Accepted

Workshop Papers

- **Quantifying Cross-Query Contradictions in Multi-Query LLM Reasoning**
ICLR 2026 Workshop on LLM Reasoning — Accepted
- **Multiscale Diagnostics of Visual Language Models**
ICCV 2025 Workshop (CV4DC) — Accepted
- **Curate, Don't Augment: Budgeted Multimodal Intervention Design for Invariant Vision-Language Learning**
CVPR 2026 Workshop — Accepted
- **Beyond Hallucinations: A Composite Score for Measuring Reliability in Open-Source LLMs**
AAAI 2026 Workshop (AIR-FM) — Accepted

Research Experience

Indian Institute of Science (IISc), *Computer Vision Intern* Mar 2025 – June 2025
Advisor: Prof. Suresh Sundaram

- Researched transformer-based object detection and developed **Hi-DETR**, a hierarchical dense supervision detector for improved feature aggregation and small-object sensitivity.
- Implemented a hybrid encoder with multi-scale deformable attention and task-aligned cross-attention to enhance detection robustness in cluttered AGV environments.

- Demonstrated improved small-object recall and localization accuracy over **RT-DETRv3** and **SO-DETR** baselines on internal benchmarks, with higher robustness in cluttered scenes.
- Submitted the Hi-DETR manuscript to **IEEE Expert Systems** for publication consideration (first-author submission).
- Built end-to-end perception modules for autonomous ground vehicles (AGVs), integrating LiDAR, radar, and camera streams for detection, tracking, and 3D scene understanding.
- Deployed optimized PyTorch pipelines to embedded systems using ROS and OpenCV, achieving real-time inference under on-board compute constraints.

National Chung Cheng University, Taiwan — *Research Intern*
 Advisor: Prof. Chih-Yi Chiu

Dec 2024 – Mar 2025

- Developed a Diversity-Aware Recommendation System (Image + Text) under TEEP@AsiaPlusProgram.
- Integrated ViT for image and BERT for text feature extraction, enhancing multi-modal recommendations.
- Utilized FAISS for fast similarity retrieval and implemented diversity & fairness metrics to improve recommendation quality.
- Applied re-ranking techniques to ensure diverse, novel, and unbiased food item suggestions.
- Challenges Overcome: Addressed data alignment issues between image and text features, and fine-tuned model parameters for optimal diversity in recommendations.

Taiwan-India Big Data Analytics Lab, SRM University — *Research Intern* **Aug 2023 – Sep 2024**

- Implemented a CLIP-based multimodal patent retrieval system linking patent drawings to abstracts.
- Integrated long-tail optimization, novelty, and fairness metrics for high-precision patent search.
- Researched techniques like "LLM Patent Retrieval," "PatentGPT," and "DeepPatent2" to improve patent retrieval by linking drawings with abstracts for better contextual understanding.

Selected Projects

Backprop NEAT with Increased Network Complexity

- Integrated NEAT with backpropagation using JAX, enabling simultaneous evolution of network architectures and weight training.
- Validated on three 2D classification tasks (XOR, Circle, Spiral), demonstrating robust performance across diverse data distributions.
- Employed advanced mutation operators to incrementally add nodes and connections, increasing network complexity while effectively managing overfitting.
- Leveraged JAX for high-performance numerical computing, facilitating rapid experimentation and prototyping of evolving neural architectures.
- Documented improved decision boundaries and enhanced classification accuracy across various datasets.
- Achieved perfect classification on the XOR dataset and notable performance gains on the Circle and Spiral tasks, illustrating the evolved networks' optimal balance between complexity and predictive accuracy.

Multiagent Debate for Visual Reasoning

- Adapted the multiagent debate strategy for vision models to improve factual accuracy and reasoning in visual tasks. The idea is inspired by the paper "Improving Factuality and Reasoning in Language Models through Multiagent Debate".
- Deployed multiple instances of state-of-the-art vision models (e.g., Vision Transformers/CLIP) to debate image interpretations.
- Designed a multi-round debate protocol that aggregates diverse perspectives to converge on the most consistent and accurate visual prediction.
- Validated the approach on visual reasoning datasets (e.g., VQA/GQA), achieving significant improvements over single-agent baselines.

Skills

Languages: Python, C++, Bash, JavaScript

Frameworks: PyTorch, TensorFlow, JAX

Vision Tools: OpenCV, Open3D

Systems: Git, Docker, MLflow, DVC

Cloud/Compute: AWS, Google Colab, NVIDIA Jetson

Awards

- Qualified Round 2 of Google Code Jam
- Awarded TFP Winter Scholarship from Virginia Tech to attend AAAI 2026.