

# Tushar Tyagi

Atlanta, GA | 678-444-7973 | [tushartyagi1811@gmail.com](mailto:tushartyagi1811@gmail.com) | [tushar-tyagi.github.io](https://tushar-tyagi.github.io) | [github.com/Tushar-Tyagi](https://github.com/Tushar-Tyagi) | [linkedin.com/tushartyagi01](https://linkedin.com/tushartyagi01)

## Education

Master of Science, Computer Science, Georgia Institute of Technology

GPA: 4.0/4.0

B.Tech, Computer Science and Engineering, Indian Institute of Technology Mandi

CGPA: 8.39 / 10.00

Teaching Assistant: IC152: Intro to Python and Data Science (500 students)

August 2024 - Present  
(Machine Learning Specialization)  
August 2017 - June 2021  
(Minor in Management Studies)

## Technical Skills

- **AI/ML Core:** Machine Learning, Deep Learning, NLP, Computer Vision, Generative AI, Agentic AI, Reinforcement Learning, Recommendation Systems, Time Series Forecasting, Statistics, Self-Supervised Learning, Quantization, Multimodal Agentic Systems, Model Deployment, Reasoning Fine-tuning LLMs.
- **Frameworks/Libraries:** PyTorch, TensorFlow, Keras, Hugging Face (Transformers), Scikit-learn, OpenCV, NLTK, SciPy, FAISS, Numpy, Pandas, Matplotlib, FAISS.
- **MLOps/Deployment:** MLflow, TFLite, ONNX Runtime, Docker, Kubernetes, TFX, TFExtended, TFServe, CI/CD.
- **Systems/Programming:** Python, C/C++, Multithreading, REST APIs, Bash, Git.

## Patent

[US20240404070A1](#): Precise book page boundary detection using deep ML model and image processing algorithms.

## Research Projects

- **Agentic Video Retrieval with Structured Verification:** Architecting a **compute-optimal** system to resolve the **ingest-vs-query cost bottleneck** via **structured verification**; developed **adaptive candidate pruning** and **uncertainty-aware tool selection** (Knapsack-style allocation) to maximize information gain per GPU dollar for complex multimodal queries.
- **VLM Context Engineering for Video Analytics:** Built fine-grained optimization system using task-type classifiers to dynamically select among five context retrieval complexities (baseline RAG to agentic search), achieving up to **10% accuracy improvement and 2x speedup** over static methods on Video-MME.
- **PROPEL:** Developed PROPEL, a training-free, explainable framework leveraging empirical bias profiling and inverse propensity weighting to mitigate position bias in LLM re-ranking, improving recommendation accuracy by up to 26.7% in NDCG@1 across real-world datasets.

## Technical Projects

- **DQN-HER:** Implemented **Hindsight Experience Replay (HER)** to solve sparse-reward BitFlip environments; integrated **Prioritized Experience Replay** using Sum Tree data structures, achieving a 100% success rate on 16-bit configurations.
- **LLM Fine-Tuning:** Implemented **Group Relative Policy Optimization (GRPO)** for LLM alignment; utilized group-mean rewards as control variates to stabilize policy updates and enforced **KL-divergence penalties** to maintain reference model proximity.
- **DeepPIR (ViT Adaptation):** Led fine-tuning of Swin-V2 for 2-channel images using **LoRA (Low-Rank Adaptation)** utilizing **Hugging Face PEFT module**; optimized training using progressive unfreezing and AdamW scheduling to achieve high-performance transfer learning from ImageNet.

## Experience

Machine Learning Engineer 2

July 2021 – July 2025

Adobe

Noida, India

- Engineered a real-time document classification pipeline (**90% accuracy**) with a custom downsampling algorithm achieving **10x speedup** over Lanczos; increased feature MAU by **50%**.
- Achieved **3x inference speedup** and **4x model size reduction** via quantization and pruning; eliminated 3-second GPU delays and boosted boundary detection by **25.6%**.
- Spearheaded a "Book Mode" segmentation feature for **5M monthly users**, improving dewarping accuracy from **70% to 80%** while doubling runtime performance.
- Architected a modular LLM-based actions pipeline for Acrobat AI Assistant; applied **Chain-of-Thought prompting** to deliver a **66% performance enhancement** and 40% query accuracy boost.

Product Development Intern

May 2020 – July 2020

Adobe

Noida, India

- Built image-cleaning tools using **OpenCV (Hough Lines)** and Keras; authored scripts to generate synthetic training data for segmentation models.

Product Development Intern

Jan 2021 – Feb 2021

Furrble

Bengaluru, India

- Designed and implemented scalable CRUD API endpoints using **GoLang** to enhance backend functionality.

Volunteer Developer

Aug 2018 – Nov 2018

EinsteinPy

Open Source

- Developed translation modules for physical coordinate systems to enhance interoperability for a Python library.