

Hardik Manish Runwal



Mathematics & Computing Researcher · Machine Learning · Model Behavior & Interpretability
hardikrunwal123@gmail.com | +91-9405184640 | GitHub | LinkedIn | Hugging Face

Research Interests

Model Behavior & Interpretability · Representation Learning · Generative Models · Computer Vision · AI Safety

Research Experience

Lead Student Researcher — Dr. Ujjwal Verma *Aug 2025 – Present*

Manipal Institute of Technology *Satellite-Imagery Research*

- Leading a faculty-supervised project on transformer-based detection and segmentation for satellite imagery.
- Owned the end-to-end pipeline for a **10+ GB** dataset — annotation, cleaning, preprocessing.
- Designed experiments across **11+ detection and 6 segmentation models**, tuning hyperparameters and losses.
- Coordinated a **7-student team** through weekly reviews, result analysis, and reruns.
- Analyzed **50–100 experiments** (accuracy, precision, recall, F1, IoU, mIoU, mAP); best model: **91% accuracy, 0.73 mIoU**.
- Reviewed literature to guide model selection; contributing to the report and manuscript prep.

Research Projects

AI-Powered Document QA (RAG System) *LangChain, FAISS*

- Built an end-to-end retrieval-augmented generation pipeline covering document ingestion, semantic chunking, FAISS-based vector search, and LLM answer synthesis, achieving sub-second retrieval over multi-hundred-page corpora.
- Decomposed the system into four independently benchmarkable components (ingestion, indexing, retrieval, generation) and measured retrieval precision at top-5 along with answer faithfulness using automated evaluation scripts.
- Designed explicit extension points for multi-document grounding and hallucination detection, and validated the architecture on a financial knowledge base with domain-restricted context control.

RL Experimentation Framework *PyTorch, Stable-Baselines3*

- Implemented Q-Learning and PPO agents across discrete and continuous Gymnasium environments; PPO converged to a stable policy in CartPole-v1 within $\sim 100K$ steps, outperforming the Q-Learning baseline.
- Automated model checkpointing, replay-video generation, and deployment to the HuggingFace Hub with reproducible evaluation scripts; 3+ trained agents made publicly available.
- Built reusable training and evaluation utilities that cut experiment setup time from hours to minutes across different environment families.

Denoising Diffusion Probabilistic Models *PyTorch*

- Implemented DDPM from scratch by working through the original paper: forward/reverse diffusion, a cosine noise schedule, and a U-Net denoiser, producing visually coherent samples on MNIST, CIFAR-10, and a reduced CelebA-HQ subset.
- Switched from a linear to a cosine noise schedule, improving FID by 18%, and ran ablations over timestep counts (100–1000) to quantify the quality-speed tradeoff.

Word2Vec — CBOW & Skip-gram *NumPy, PyTorch*

- Implemented CBOW and Skip-gram with negative sampling first in NumPy to understand the underlying gradient computations, then reproduced the same models in PyTorch for efficient training on a custom corpus.
- Validated embedding quality using semantic analogy tests with cosine similarity, and visualized the learned embedding space to confirm meaningful latent structure.

Variational Autoencoder

PyTorch

- Implemented a VAE from its mathematical formulation, combining the reparameterization trick with a joint reconstruction and KL-divergence objective, and trained it on MNIST and CIFAR-10.
- Ran ablations over latent dimensionality (16–256) and inspected latent-space interpolations, observing smooth semantic transitions between classes as capacity changed.

XGBoost Classification Pipeline *scikit-learn, XGBoost*

- Built a feature-engineered classification pipeline with 5-fold cross-validation to get a robust estimate of generalization performance.
- Tuned hyperparameters including max depth, learning rate, and subsample ratio, achieving **91% F1** on held-out test data.

LambdaLab — HF Trade Arrival Modeling *FastAPI, React, Plotly*

- Modeled tick-level trade arrivals with self-exciting Hawkes processes, whose intensity increases after each event to capture clustering in market activity; MLE-fitted the parameters on 1M-event market data, giving a 34% log-likelihood improvement over a Poisson baseline.
- Exposed the fitted model through a FastAPI back-end paired with an interactive React/Plotly dashboard, letting users tune excitation parameters in real time and fit the model to their own uploaded CSV datasets.

Education

Manipal Institute of Technology *2024 – 2028 (expected)*
B.Tech. in Mathematics and Computing

Coursework: Data Structures & Algorithms[†], Operating Systems[†], Database Systems[†], Linear Algebra, Probability & Stochastic Processes, Fundamentals of ML[†], AI[†], Time Series Analysis[†], Numerical Methods[†]

[†]In progress **Upcoming:** Linear/Non-Linear Optimization, Game Theory, Math for Finance

Technical Skills

Languages: Python, Java, C, C++, SQL, JavaScript/TypeScript
ML & DL: PyTorch, TensorFlow, Keras, scikit-learn, XGBoost, Stable-Baselines3

GenAI / LLMs: LangChain, HuggingFace Transformers, FAISS, RAG pipelines, diffusion models, VAEs

DL Concepts: Transformers, Computer Vision, Object Detection, Semantic Segmentation, Representation Learning, NLP

Scientific Python: NumPy, SciPy, Pandas, Matplotlib, Plotly
Tools & Cloud: FastAPI, Streamlit, React, Git, Docker, TensorBoard, Linux, AWS, GCP

Competitive Programming & Open Source

LeetCode (Jul 2025 – Present): 225+ problems solved (Python/Java/C); rating 1576; DP, D&C, backtracking, binary search, hashing.

GitHub: Open-source portfolio spanning LLM systems, generative models, and RL pipelines, with methodology notes and reproducibility scripts.

Certifications

Intro to Deep Learning, Intro to Machine Learning, Computer Vision (Kaggle) · Python & OS (Google/Coursera) · Supervised ML: Regression & Classification (Stanford/DeepLearning.ai) · AWS Cloud Practitioner, AI Practitioner, ML Specialty — Associate (In progress)