SIMRAN LALWANI

similalwani13@gmail.com · linkedin.com/in/similalwani · github.com/similalwani · simranlalwani.com · San Francisco, CA

EDUCATION

M.S., Data Science, Rochester Institute of Technology, Rochester NY, **GPA: 3.8 / 4.0** **August 2022 – December 2024**
Coursework: Neural Networks, Database Design Systems, Applied Statistics, Time series Forecasting, Applied Research

WORK EXPERIENCE

Machine Learning Engineer, TransferX, San Francisco **February 2026 – Present**
- Building a **data engine** on **PostgreSQL**: schemas and **ETL pipelines** for **multimodal data** at an early-stage EdTech platform.
- Built a **low-latency model-serving and retrieval-evaluation layer**, targeting **sub-50ms response** on **ranking tasks**.
- Engineered a **unit-testing-for-data layer** with **input-drift detection**. Cut **production model regressions by 25%**.

AI Systems Integration, Alta Potentia, Remote **March 2025 – November 2025**
- Deployed YOLOv8 on NVIDIA Jetson at **sub-100ms** latency for real-time object detection in retail and security applications.
- Quantized models with **TensorRT**, achieving **60% size reduction** while maintaining accuracy, enabling edge deployments that reduced infrastructure costs by **30% per device**.
- Built data pipeline processing **500K+ frames** in coordination with clients, reducing manual review workload by **16%** in retail loss prevention.

Machine Learning Engineer (CV Co-op), OWL Autonomous Imaging, Rochester NY **January 2024 – May 2024**
- Extended **MGNet** into a **thermal + RGB panoptic fusion** model with **depth estimation** on **Detectron2** and **Faster R-CNN**. Improving **20% mIoU** on **depth perception** for autonomous navigation.
- Developed an interactive **panoptic-labelling** UI with sliders, color-coded class picker, and **parallax correction**. **40%** faster annotation, scaled to **50K+ frames**.
  Built a **YOLOv5 data-quality pipeline** for **automated assessment**, **class balancing**, and **parallax correction**. **+15%** detection precision.
- Shipped real-time **super-resolution** and **denoising** for **16-bit thermal + RGB**. **Dockerized** into the existing **CI/CD pipeline**.

Graduate Research Assistant, RIT, Rochester NY **January 2023 – May 2023**
- Designed **Reinforcement Learning** driven RNN architecture search on **drone telemetry (NGAFID)** with Prof. Travis Desell, treating architecture mutations as policy actions optimized against flight-safety prediction reward .
- Co-developed **neuroevolution operators** for the **EXAMM** framework in **C++**: new **genome representations** and **crossover strategies**. Ran training across a **master-worker compute topology**.

Machine Learning Engineer, WhiteHat Education Technology, Mumbai **June 2020 – August 2021**
- Built robust **unit testing** frameworks (**PyTest)** for **300+** Python ML projects, drastically reducing production failures.
- Deployed automated **data drift detection** (**MLflow**, **Airflow**) to ensure stable model performance across diverse datasets.
- Analyzed marketing, sales, and student data, implementing algorithms that boosted **lead conversion**, **revenue**, **customer retention** and **learner engagement.**
- Led a **seven-person team** to design **data-driven A/B tests** and **predictive analytics workflows**, significantly improving **student success rates** and **overall program efficiency**.

SKILLS

- **Programming Languages**: Python, Java, C, C++, SQL, JavaScript (React), R
- **Python Libraries**: NumPy, Pandas, Streamlit, Plotly, NLTK, Matplotlib, OpenCV, TensorFlow, Keras, PyTorch, Seaborn, Albumenttaions
- **Tools & Frameworks**: LangChain, LlamaIndex, MySQL, vLLM, Transformers, Git, Docker, PowerBI, Tableau, SQLite, MongoDB

RESEARCH & PROJECTS

Cosmos-Overwatch: Synthetic Data Orchestration for Drone Perception | *NVIDIA Cosmos-Transfer2.5, PyTorch* **April 2026**
- Built an end-to-end **synthetic data pipeline** on **NVIDIA Cosmos-Transfer2.5** that domain-shifts real aerial footage into rain, fog, thermal, and night-fire.
- Implemented **Dual-ControlNet** (Edge + Depth) with **calibrated CFG** to preserve structure under **maximum perceptual domain shift**.
- Engineered an **inference orchestrator** that processes **48+ video sequences** on **A100-80GB instances**.
- Built **pre-flight validation** and **automated post-inference QC** catching black frames, static video, and structural dissolution. Validated **~50K MOT17-compatible annotations**.

Domain Generalization for ECG (Capstone) | *DANN, MMD, PyTorch | RIT* **December 2024**
- Implemented **DANN** and **MMD** on a proprietary clinical + simulation time-series dataset. **+30%** over the **VAE baseline**, **+18%** on unseen patients' ECG data.

VoiceBit Lite: Voice-Powered Restaurant Ordering | *FastAPI, Whisper, Groq, SQLite* **February 2026**
- Built an end-to-end **voice-ordering prototype**: **Whisper for transcription**, **Llama 3.3 for intent parsing**, and **REST API** with persistent storage. **Latency < 2s.**

WattsNext: Time-Series Forecasting for Energy Mix | *Python, LSTM, XGBoost* **August 2020**
- Developed an **LSTM + XGBoost** forecaster for **time-series** energy consumption: **92% accuracy**, **+30% sustainability**, **-20% energy costs**.