

# Jasper Tan

JasperTan@utexas.edu | (908) 838-5638 | [Personal Website](#) | [LinkedIn](#) | [GitHub](#)

## EDUCATION

---

### The University of Texas at Austin

Master of Science in Electrical and Computer Engineering; Concentration: Machine Learning/Data Science

May 2026

GPA: 4.0/4.0

### The University of Texas at Austin

Bachelor of Science in Electrical and Computer Engineering; Minor in Business Administration

May 2025

GPA: 3.85/4.0

Coursework: 3D Deep Learning, Computer Vision, Systems in Generative AI, Reinforcement Learning, ML on Networks, Generative Models, Database Management, Data Science Lab, Comp Arch, OS, Matrices, Probability, Algorithms

## TECHNICAL SKILLS

---

**Programming Languages:** Python, SQL, C/C++, Java, Bash/Shell

**Frameworks & Tools:** PyTorch, SentenceTransformers, Diffusers, PEFT, Scikit-Learn, Weights & Biases, Streamlit

**Technologies:** LoRA/QLoRA, Snowflake, AWS (S3, EC2, SageMaker), Azure, Jenkins, Git, Linux, PostgreSQL, MongoDB, JIRA

## EXPERIENCE

---

### SyntheFY - Machine Learning Engineer; Austin, TX

Sep 2025 – Present

- Developing a multimodal diffusion foundation model to generate PPG time-series data based on heterogeneous metadata
- Architecting a Pytorch Lightning training pipeline with a custom meta data encoder for in-context example conditioning
- Developing a memory-efficient PyTorch dataloader with frequency resampling for diverse, large-scale PPG datasets

### AMD - Machine Learning Infrastructure Intern; Austin, TX

May 2025 – Aug 2025

- Built a scalable semantic retrieval pipeline using OpenAI embeddings to match 100+ hardware logs via vector similarity
- Compressed 200k-token documents by 97% using regex, chunking, and fingerprinting to enable long-context embeddings
- Improved ETL speed by 45x on 1.6M-row datasets via optimized Snowflake ingestion with Parquet and Apache Arrow
- Automated log ingestion from MongoDB and deployed interactive Streamlit dashboards to enable AI-driven diagnostics

### The University of Texas at Austin - Machine Learning Researcher; Austin, TX

Aug 2023 – May 2025

- Advanced a novel human activity classification model incorporating feature fusion with real-time acoustic and inertial data
- Adapted a MobileNet V2 architecture via transfer learning for feature extraction and fine-tuning based on IMU data
- Deployed a lightweight computer vision model for object localization on edge devices, optimizing inference via LiteRT

### Cvent - Software Engineer Intern; Tysons Corner, VA

Jun 2024 – Aug 2024

- Designed date-time modals in JavaScript with optimized state management using Redux, React Hooks, and mutations
- Implemented real-time data visualization and logging with Datadog to monitor GraphQL queries to a PostgreSQL database

### FirstParty - Applied Machine Learning Intern; New York, NY

Jun 2023 – Jun 2024

- Developed Python pipelines in SageMaker using GPT embeddings to compute string similarities across unstructured data
- Leveraged cosine similarities and Levenshtein distance algorithms to generate confidence scores for data stored in S3
- Engineered ML-based data classification pipelines using natural language processing, achieving a 95% accuracy rate

## PROJECTS

---

### Chatbot-Enhanced Recommender System

May 2025

- Built a conversational recommender integrating LLMs with DeepFM and BERT encoders to improve inclusion@10 by ~20%
- Fine-tuned Gemma, Llama 2, and Mistral using LoRA on MovieLens 20M to simulate user queries and generate dialogue
- Trained a dual-encoder model (BERT + item embeddings) with triplet loss and VICReg to align conversations with items

### Graph Reinforcement Learning for Semantic Segmentation

Jan 2025

- Implemented graph-based reinforcement learning to improve semantic segmentation of 2D images and 3D point clouds
- Developed graph convolutional networks with dueling Deep Q-learning, optimizing node classification and navigation
- Designed large-scale graph environments handling 30,000+ nodes using k-nearest neighbor for 2D and 3D data

### Fashion-Atlas

May 2024

- Devised a garment re-identification application aimed at localizing clothes from images to give tailored recommendations
- Leveraged YOLOv8 to train a real-time object detection and classification neural network to crop and identify images
- Trained a CNN on a ResNet 50 architecture with a triplet loss function and cosine similarity to generate feature embeddings

### Hindsight Experience Replay for Diffusion Models (HERD)

May 2024

- Fine-tuned a text-to-image diffusion model (Stable Diffusion) using Reinforcement Learning to generate prompted images
- Built a distributed training pipeline using Transformer RL, image reward, and policy gradient methods (DDPO, DPOK, DDPG)