

# Shubham Goel

SENIOR MACHINE LEARNING ENGINEER · MULTIMODAL | RECSYS | COMPUTER VISION

☎ 217-979-1769 | ✉ shubhamgoel27@gmail.com | 🏠 shubh.io | 📱 shubhamgoel27 | 🌐 shubhamgoel27

## Professional Experience

### Senior Machine Learning Engineer

Menlo Park

META — INSTAGRAM ADS RANKING

Dec. '24 - May '26

- Built Instagram Ads' **Video OCR pipeline**: quality-ranked frame sampling with SSIM-based deduplication tailored to short-form video ads, plus a **domain-tuned OCR backbone** reaching **near-parity F1** with the production VLM (~1.5pp below Llama 4 Scout) at ~80x faster inference — delivering **0.06–0.08% online NE** (~\$30M+ annual revenue lift), ~0.12% on Reels.
- First author, **RecSys '26**, Industry (under review) — “*The Text on the Creative: An Under-Exploited Ranking Modality for Short-Form Video Ads.*”
- Designed and shipped a **multimodal LLM inference pipeline** generating ad-side ranking features (ad summary, persuasion technique, target audience, etc.) for the Instagram Ads Ranking Model on an off-the-shelf Qwen3 14B; ablated prompts and inference hyperparameters (temperature, sampling) against a per-field eval harness.
- Distilled Claude Opus 4.7 (teacher) labels into the base Qwen3 14B via SFT, lifting extraction accuracy ~**9pp**; final pipeline features delivered a **0.05% NE gain** on the IG Ads ranking model.
- Led a **cross-app data augmentation** effort integrating Facebook ad-click signal via auxiliary learning tasks, transferring click data across surfaces; shipped with a **0.04% NE gain** on the IG Ads ranking model (~\$20M+ annual revenue lift).
- Built **FeatureLens**, a feature-intelligence platform for ads ML — semantic search, near-duplicate detection, and cross-model gap analysis over the feature catalog; adopted by **80 ranking engineers** and powered shipment of ~**250 new features** across ranking models.

### Senior Machine Learning Scientist

Remote

ZEFR

Jun. '21 - Dec. '24

- Architected an internal **multimodal semantic search** platform on Qdrant over **1B+ image & text embeddings**, replacing keyword-based search for a 14-person labeling team and powering training data sourcing for CV/NLP fine-tunes across 12 brand-safety taxonomies.
- Fine-tuned **CLIP / ViT** models on ZEFR's content taxonomy for ad image classification; built a vector-search-based active learning system to scale sparse-class annotation in production — invited speaker, **Snap ML group (Aug '23)** on “*Labelling Sparse Data at Scale using Vector Search.*”
- Built a platform-agnostic ONNX inference pipeline serving fine-tuned ViT/LLM workloads at >1M images/day; designed the upstream training stack (auto-reports, quantization, model registry) and a Metaflow embedding pipeline over 100M+ videos.
- Drove ML strategy and roadmap for the team; led senior MLE hiring loops and mentored junior engineers.

### Research Assistant

UNIVERSITY OF ILLINOIS AT URBANA CHAMPAIGN

Jan. '20 - Dec. '20

- Led cyclone-damage prediction combining pre/post-disaster geospatial imagery with SBA parcel insurance data; Siamese change-detection network outperformed all baselines at IoU 0.804. Invited speaker, **GeoPython 2020**.

### Machine Learning Engineer - Computer Vision

RELIANCE JIO AI CENTRE OF EXCELLENCE

Oct. '18 - Jul. '19

- Built a Building Detection module for dense/cluttered geospatial imagery using Reverse Mask R-CNN, beating SOTA mIoU by 9.3%; shipped an imagery-acquisition tool that 10x'd throughput and replaced the team's external vendor tool across a 1,400-member GIS org.

### Software Engineer - Machine Learning

PITNEY BOWES

Oct. '16 - Oct. '18

- Built building-footprint segmentation (IoU 0.845) and rooftop-type classification (96.5%) in MXNet; implemented PSPNet in Caffe with a custom Sync-BatchNorm layer for multi-GPU training. First DL hire on the team; mentored subsequent joiners.

## Education

### University of Illinois at Urbana-Champaign

MASTERS IN STATISTICS

Aug. '19 - Dec. '20

- GPA **3.52**. Coursework: Statistical Machine Learning, Computer Vision, Machine Learning in NLP, Deep Learning, Autonomous Vehicles, Applied ML.

### Cluster Innovation Centre, Delhi University

BACHELORS IN TECHNOLOGY IN MATHEMATICS AND I.T.

Aug. '12 - Jul. '16

- Coursework: Machine Learning, Linear Algebra, Calculus, OOP, Computer Networking, Computer Graphics, Graph Theory.

## Skills

<b>Languages &amp; Core</b>	Python, SQL, Git, Unix, AWS
<b>ML / Deep Learning</b>	PyTorch, CUDA, ONNX, NumPy, Pandas, Scikit-learn, OpenCV
<b>LLMs &amp; Multimodal</b>	multimodal LLMs (text/image/video/audio), SFT / LoRA fine-tuning, teacher-student distillation, RQ-VAE / vector quantization, sentence-embedding training, LLM-as-judge eval, ViT
<b>ML Infra &amp; Retrieval</b>	Distributed inference, Qdrant, HNSW, Metaflow, Docker, AWS Batch, Seldon, Weights & Biases