KAMAL POSHALA

Norman, Oklahoma | Open to Relocation | +1 (405) 430-2216 | kamalposhala.cs@gmail.com
LinkedIn | GitHub | Portfolio

F-1 Optional Practical Training | Available: July 6, 2026

SUMMARY

Software engineer with demonstrated project, research, and internship experience spanning full-stack web development, backend systems, and applied machine learning. Skilled in building scalable REST APIs, real-time distributed systems, and React/TypeScript front ends, alongside production-oriented ML pipelines for multilingual NLP and multimodal AI. Strong analytical and debugging ability, honed through diagnosing and resolving concurrency and performance bottlenecks in production-style systems. Comfortable across SQL and NoSQL databases, Git-based version control, and CI/CD pipelines, and quick to adapt to new frameworks and tools. Collaborative communicator, partnering with faculty, researchers, and cross-functional teams to turn ambiguous requirements into shipped features.

TECHNICAL SKILLS

Backend & Systems: Python, Java, JavaScript, Node.js, Express.js, FastAPI, REST APIs, Microservices, Celery, Redis, Socket.IO, Distributed Systems, Debugging & Troubleshooting, Docker, CI/CD, Linux, Git

Frontend: React.js, TypeScript, Tailwind CSS

Databases & Data Engineering: MongoDB, MySQL, Query Optimization, Indexing, Aggregation Pipelines, Data Pipeline Design, Data Validation

Machine Learning & AI: PyTorch, TensorFlow, Scikit-learn, Hugging Face Transformers, LoRA, PEFT, Fine-Tuning, Transfer Learning, mBERT, mT5, Multilingual NLP

Speech & Vision AI: OpenAI Whisper (ASR), Wav2Lip, GAN, Edge-TTS, FFmpeg

EXPERIENCE

Research Assistant — Advisor: Professor Farrokh Mistree — University of Oklahoma *June 2026 – Present*

- Architecting the SRL Academic Ancestry Platform in Python, building structured data validation pipelines that collect and visualize advisor-student research lineage across multiple generations of researchers.
- Engineering scalable data collection and validation workflows that consolidate multi-source researcher information into a unified, searchable dataset.
- Partnering with faculty and researchers to maintain data accuracy and improve the discoverability of academic genealogy records.

Software Engineer Intern — Persistent Systems *August 2023 – October 2023*

- Reduced Java module execution time by 25% through algorithmic redesign using graph traversal and heap-based scheduling, along with memory optimization techniques.
- Strengthened multithreaded application stability by diagnosing and resolving synchronization bottlenecks and CPU contention across concurrent Java modules.

PROJECTS

SyncSpace — Full-Stack Real-Time Collaborative Workspace | React, Node.js, Express.js, Socket.IO, MongoDB *2025*

- Architected a real-time collaborative platform with WebSocket-based concurrent editing, live presence tracking, and JWT-secured role-based access control; reduced MongoDB write load 60% under sustained edit bursts via debounced buffering.
- Designed room-based, event-driven communication with modular socket handlers; deployed backend on Koyeb and frontend on Vercel.

IPMS — Internship Program Management System | Node.js, Express.js, MongoDB, REST APIs, Cron *2025*

- Designed a distributed multi-role approval workflow with role-based access control and JWT authentication, implementing 15+ RESTful endpoints handling 1,000+ transaction events with centralized logging.
- Automated email reminders using scheduled cron jobs, reducing approval turnaround time by 40%.

CogniFetch — Full-Stack NLP Search & Retrieval Engine | Node.js, TF-IDF, MongoDB, React, OCR *2024*

- Built a retrieval engine with cosine similarity ranking over 10,000+ academic documents, achieving sub-50ms p95 query latency through inverted index design — a 45% latency reduction.
- Built an OCR and NLP ingestion pipeline transforming 2,000+ scanned PDFs into structured, searchable datasets.
- Paired the retrieval engine with a React front end for interactive document search.

SemEval 2026 — Multilingual Polarization Detection (22 Languages) | PyTorch, LoRA, mBERT, mT5, Scikit-learn *2025*

- Built a multilingual NLP classification pipeline across 22 languages using mBERT and mT5 with parameter-efficient LoRA fine-tuning, achieving Macro F1 of 0.87 with 60% fewer trainable parameters than full fine-tuning.

DeepDub — Automated Video Dubbing & Lip-Sync Platform | PyTorch, Whisper, Wav2Lip, FastAPI, Celery, Redis *2025*

- Designed an end-to-end multimodal AI pipeline (ASR, translation, TTS, GAN-based lip sync) with asynchronous FastAPI/Celery/Redis orchestration and modular fault isolation for videos up to 10 minutes.

EDUCATION

MS, Computer Science — University of Oklahoma May 2026

BTech, Information Technology — CVR College of Engineering | CGPA: 9.08/10.0 | May 2024