

PUSHKARAJ BARADKAR

Los Angeles, CA | +1 (213) 275-9348 | pushkarajbaradkar1@gmail.com | [linkedin.com/in/pushks18](https://www.linkedin.com/in/pushks18) | github.com/pushks18

EDUCATION

University of Southern California, Viterbi School of Engineering GPA: 3.60 / 4.00 **Jan 2025 - Dec 2026**
Master of Science in Computer Science
Los Angeles, CA
Machine Learning for Data Science, Analysis of Algorithms, Web Technologies, Database Systems

University of Mumbai GPA: 9.13 / 10 **Jul 2020 - May 2024**
Bachelor of Engineering in Computer Engineering
Mumbai, India
Data Structures, Analysis of Algorithms, Operating Systems, Computer Networks, Machine Learning, Natural Language Processing

EXPERIENCE

Software Engineer Intern – Real-Time AI Systems | USC, Los Angeles, CA **Nov 2025 – Present**

- Built a real-time SLAM pipeline combining monocular VO and depth estimation, achieving ~20–25 FPS on noisy indoor data.
- Implemented keyframe selection and pose graph optimization, reducing drift to <0.8 m ATE over ~100 m trajectories.
- Integrated depth-based reconstruction to improve spatial consistency and reduce scale drift across long sequences.
- Debugged sensor noise, calibration drift, and VO failures, improving tracking stability and reducing dropouts by ~30%.
- Collected and validated multi-modal datasets, ensuring temporal alignment and robustness for downstream pipelines.

Software Engineer Intern – AI & Computer Vision | YMT Medical, India **Jan 2024 – Aug 2024**

- Engineered computer vision pipelines for image quality enhancement and denoising, improving robustness of AI-driven diagnostics.
- Trained and fine-tuned YOLOv8n on 2K dermatology images, supporting model validation and performance testing.
- Enhanced GPU inference and integrated clinical feedback to enhance model interpretability and reduce latency by 40%.
- Designed evaluation pipelines to measure model performance, interpretability, and latency tradeoffs in production-like settings.
- Worked closely with domain experts to refine model outputs and reduce failure cases in real clinical data.

Software Engineer Intern – Backend & APIs | Technoriya ERP Solution, India **Oct 2023 – Dec 2023**

- Developed and optimized backend services and APIs using Python and Node.js, supporting scalable data pipelines, system reliability, cloud-based deployment and processing layers for downstream analytics and ML workflows.
- Implemented secure authentication and access protocols via Firebase Auth and Firestore rules to ensure data integrity.
- Optimized RESTful APIs and Firestore queries to achieve 450 ms response times under 1,000 concurrent requests for ML integration.
- Streamlined deployment workflows using Vercel and GitHub Actions to improve delivery consistency and reliability.

PROJECTS

Godot MCP: AI-Native Game Development Interface | TypeScript, Node.js, MCP - [Link](#) **April 2026 – Present**

- Built an MCP server enabling LLM agents (Claude, Cursor) to control Godot projects via natural language; used by 10+ active users.
- Developed 30+ tools for scene, script, execution, and asset management with secure path validation and structured routing.
- Implemented CI/CD, auto-config detection, and multi-client integration for end-to-end 2D/3D workflow automation.

Documind: RAG & LLM Eval Platform | FastAPI, ChromaDB, Langfuse, PostgreSQL, Docker **April 2026 – Present**

- Built a production-grade RAG microservice using FastAPI + ChromaDB with query expansion, reranking, and grounding pipelines.
- Designed evaluation system using RAG, improving output reliability and reducing hallucinations via faithfulness metrics.
- Deployed Docker-based architecture with PostgreSQL + Langfuse for observability, tracking latency and cost across requests.

AgentPay: Autonomous Agent Economy | Solana, Avalanche, LangChain, Next.js - [Link](#) **April 2026 - Present**

- Built a multi-agent system where AI agents autonomously discover, hire, and pay other agents using on-chain payments and reputation.
- Designed smart contract system (Rust/Anchor) for agent registry, escrow, and staking/slashing with verifiable reputation tracking.
- Integrated LangChain-based orchestration with custom payment tooling (x402), enabling under 30s end-to-end agent interaction cycles.
- Developed real-time dashboard (Next.js, D3) visualizing live agent interactions and network dynamics.

Multimodal AI Interview System | Python, FastAPI, React, OpenAI/LLaMA, Whisper, OpenCV, Docker **March 2026 – Present**

- Built a real-time multimodal system using LLMs, Whisper, and OpenCV for automated technical interview evaluation workflows.
- Developed a multi-language code execution sandbox (Python/C++/Java) handling 200+ test cases with runtime and complexity analysis.
- Reduced transcription latency to <1.5s using streaming pipelines, enabling smooth real-time conversational interaction.
- Built a full-stack interface (React + FastAPI) with live transcripts, session tracking, and automated performance reports.

TECHNICAL SKILLS, PUBLICATIONS, CONFERENCES

- **Programming Languages:** Python, Java, C++, JavaScript, TypeScript, R, C#, PHP
- **AI / ML & Data Science:** Deep Learning, Generative AI, NLP, LLMs (OpenAI/GPT, Llama), Transformers, RAG, Agentic AI Systems, Autonomous Agents, Prompt Engineering, Computer Vision, Data Mining, Data Science, PyTorch, TensorFlow, LangChain
- **Frameworks & Development:** FastAPI, Flask, Django, Node.js, Express.js, React.js, Next.js, Angular.js, OAuth, OOP, Agile, .NET
- **Databases & Cloud:** MongoDB, SQL, PostgreSQL, Supabase, Redis, AWS, GCP, Azure, Docker, Kubernetes, Git, CI/CD, Gitlab, Linux environments (development, debugging, containerized services)
- **Software Engineering:** Object-Oriented Programming (OOP), design patterns, backend development, REST APIs, debugging, Agile, REST API, AWS, Lambda, EC2, Fargate, CloudWatch, SQS, SNS, SES, SageMaker