

YASHAS SADANANDA

Bengaluru, India | +91 7022805441 | yxshas565@gmail.com | [LinkedIn](#) | [GitHub](#) | [Portfolio](#)

EDUCATION

PES University, Bengaluru

Aug 2024 - Present

B.Tech in Computer Science Engineering | GPA: 7.32 / 10

- Coursework: Data Structures & Algorithms, Operating Systems, Computer Networks, DBMS, OOP, Computer Architecture | 100+ LeetCode problems solved

EXPERIENCE

PESU Venture Labs (PVL) — ABC Program

May 2026 - Present

Technology Associate Intern | Founder-Track Startup Incubation | Bengaluru, India

- Selected for PESU Venture Labs' Accelerating Bold Concepts (ABC) program, a competitive founder-track incubation cohort, to build EdgeDaemon, an AI decision-provenance and auditability platform for regulated industries (healthcare, legal, finance)
- Architected the core technical thesis around a proprietary "Retain Score" benchmark to define a compliance-grade standard for on-device personalized AI, positioning EdgeDaemon in a structurally underserved market quadrant
- Built investor-facing pitch materials and go-to-market strategy for PVL ABC 2026 cohort review; led full-stack development across React, FastAPI/Node.js, and LLM-powered agentic pipelines for product prototyping
- AI safety angle: centered the platform's value proposition on auditability and traceability of model behavior for regulated AI deployment

RESEARCH & TECHNICAL PROJECTS

Autonomous ML Pipeline Orchestration Agent — GitLab AI Hackathon

Mar 2026

Python, LangGraph, Claude Sonnet, FastAPI, Docker, GitLab CI/CD

- Designed a 5-agent autonomous ML pipeline (Trigger → Dataset Analyst → Strategist → Code Generator → Reporter) with zero manual intervention; achieved F1 = 0.7616, ROC-AUC = 0.8454 on a held-out churn dataset with 131 automated tests passing
- Implemented typed LangGraph state graphs for inter-agent communication with structured JSON observability; deployed end-to-end on HuggingFace Spaces with Docker and a full GitLab CI/CD pipeline
- AI safety angle: built full inter-agent observability into the state graph so every automated decision in the pipeline is traceable and reproducible, not a black box

AutoTrade AI — AI Decision-Support Trading Platform

Apr 2026

Team Lead, Team SynapseX | Python, FastAPI, XGBoost, LSTM, Celery, Redis, Kafka, PostgreSQL

- Designed a six-layer agentic trading architecture (UI → Agent Core → Safety Stack → ML Signal Pipeline → Broker Adapter → Data/Memory) with a 30s perceive-reason-decide-execute-self-correct loop and multi-broker integration (Zerodha, Upstox, Angel One)
- Architected a five-layer safety stack — confidence gating, deterministic hard constraints, circuit breakers, 2FA consent, immutable audit trail — ensuring the agent has no direct path to live order execution
- AI safety angle: enforced human-in-the-loop consent before any high-stakes action, implementing controllability by design

ML Pipeline Agent — OpenEnv Hackathon (Meta × PyTorch)

Apr - May 2026

Python, FastAPI, Pydantic, Docker, HuggingFace Spaces, Q-Learning, Reward Shaping

- Built an RL-based ML pipeline agent (preprocessing, model selection, hyperparameter optimization) using Q-learning with epsilon-greedy exploration and Pydantic-validated state schemas; deployed via Docker on HuggingFace Spaces
- AI safety angle: constrained the agent's action space to prevent unsafe or invalid pipeline configurations, prioritizing predictable behavior over unconstrained exploration

MilletDataNet — ML Research System for Crop Disease Classification

Jun 2026 - Jul 2026

Python, PyTorch, ResNet18, FastAPI, React, MQTT | github.com/yxshas565/milletdatanet

- Built an end-to-end ML research system for millet leaf disease classification using a fine-tuned ResNet18 backbone with k-fold cross-validation and class rebalancing to correct severe class imbalance in the field-collected dataset
- Developed a FastAPI provenance service logging every inference request, model version, and confidence score, paired with a React dashboard and a simulated MQTT IoT layer modeling real-time sensor-to-inference data flow
- AI safety angle: instrumented inference-level provenance logging (input, model version, confidence, output) so every prediction is auditable after the fact

TECHNICAL SKILLS

- Languages: Python, C, C++, Java, JavaScript, TypeScript, SQL (PostgreSQL) | DSA: Arrays, Trees, Graphs, DP, Recursion, Vectors, Maps — 100+ LeetCode solved
- AI / ML: LangGraph, LangChain, Claude API, OpenAI API, Gemini API, HuggingFace, PyTorch, RAG, Agentic AI, LLMs, XGBoost, LSTM, Q-Learning, Prompt Engineering, Streamlit
- Full-Stack: React, Next.js, Tailwind CSS, Radix UI, Node.js, Express.js, FastAPI | Databases: MongoDB, MySQL, PostgreSQL, Supabase, Firebase
- DevOps / Tools: Docker, HuggingFace Spaces, GitLab CI/CD, Git, GitHub, VS Code, Vercel, Linux, MQTT

CERTIFICATIONS, ACTIVITIES & INTERESTS

- Anthropic — Claude Fluency in AI & Claude 101 | McKinsey Forward Learning Programme | Azure GenAI & Neural Networks | AI/ML, IIT Patna
- Python (Adv.) — HackerRank | C — CodeChef | Open Source Contributor: GSSoC, SSoC, HuggingFace