Jyotir Vinay Naram

+46-0768944415 | jv5102003@gmail.com | LinkedIn | GitHub | Portfolio
Karlskrona, Sweden | Authorized to work in Sweden | Open to relocation

PROFESSIONAL SUMMARY

Full-Stack Engineer and AI Developer with 2+ years of production experience building end-to-end AI-powered platforms, RAG pipelines, and neuro-symbolic systems. Demonstrated ownership as Technical Lead architecting multi-tenant social media platforms (Next.js, Node.js, PostgreSQL) integrated with Meta, LinkedIn, and OpenAI APIs at 99.9% uptime. Strong track record across agentic AI systems (Azure OpenAI, FAISS, LangChain, FastAPI), game AI development, and scalable backend engineering with Docker, CI/CD, and cloud infrastructure. Internship experience across Sweden and India in fast-paced agile teams; authorized to work in Sweden and open to relocation.

EDUCATION

- **Blekinge Institute of Technology** *September 2024 – September 2025*
  Bachelor Qualification Plan in Computer Science Karlskrona, Sweden
  - Specialization: Full-Stack Development, AI/ML Systems, Cloud Computing, Software Architecture, Agile Methodologies.
- **University College of Engineering Kakinada (JNTUK)** *November 2021 – June 2024*
  Bachelor of Technology in Computer Science & Engineering Kakinada, India
  - Specialization: ML, Data Structures & Algorithms, Database Systems, Web Technologies, Software Engineering.

LANGUAGES

Language	English	Swedish	Hindi	Telugu
Proficiency	C2 (Fluent)	A2 (Learning)	C2 (Fluent)	Native

CORE TECHNICAL SKILLS

Frontend: Next.js, React 19, TypeScript, Tailwind CSS, Vite, Angular, State Management, Server/Client Components
Backend / APIs: Node.js/Express, Python/FastAPI, RESTful APIs, OAuth 2.0, Microservices, Session Management, Rate Limiting
AI & ML: GPT & Azure OpenAI, LangChain, FAISS, RAG Pipelines, RoBERTa, Mistral 7B, Neuro-Symbolic Systems
Databases: PostgreSQL, Prisma ORM, Redis, Firebase, MongoDB, FAISS, JSONB, Query Optimization
Cloud / DevOps: AWS, Docker, CI/CD, GitHub Actions, Render, Netlify, Serverless Architecture
Other: Git/GitHub, Agile/Scrum, Jest, Pytest, API Design, Technical Documentation, Unreal Engine 5 (Blueprints, C++)

PROFESSIONAL EXPERIENCE

- **AiLandMedia** [🌐] *November 2025 – January 2026*
  Technical Lead & Full-Stack Engineer (Intern) Karlskrona, Sweden
  - Architected a multi-tenant AI-powered social media management platform end-to-end using Next.js (App Router), React 19, TypeScript, and Node.js/Express, delivering an MVP-quality production system in 1 month with 99.9% uptime and 90+ Lighthouse scores across all pages.
  - Integrated Meta Graph API v24.0, LinkedIn API, and OpenAI GPT-3.5-turbo with OAuth 2.0 flows, token refresh, and RAG pipeline for AI-assisted content generation; first-submission Meta App Review approval with zero delays.
  - Designed PostgreSQL (Neon) schema with Prisma ORM and JSONB analytics storage; implemented optimized indexing reducing query time by 35%, and set up CI/CD pipelines enabling zero-downtime deployments across 28 major features and 10 delivery phases.
  - Delivered real-time analytics ingestion pipeline with hourly cron-based metric collection and unified publishing for Facebook, Instagram, and LinkedIn (200 posts/batch via node-cron), reducing missed posts by 90% and reporting time by 90%.
- **Gaddr** [🌐] *May 2025 – September 2025*
  Full-Stack Developer Intern Stockholm, Sweden
  - Developed and maintained RESTful APIs for job postings, applicant tracking, and screening workflows using Next.js and Python/FastAPI, achieving 99.5% uptime and delivering features 3× faster than the team average by integrating AI tooling into the development workflow.
  - Integrated ChatGPT and Grammarly APIs directly into the job-posting platform to automate proposal drafting and review flows, reducing review cycle time from 24–72 hours down to 2–6 hours.
  - Refactored 3–6 legacy backend modules and hardened workflow validation, reducing API error occurrences by 30–50% across critical task and job-posting flows while maintaining Agile sprint cadence across a 10-sprint, 8-person team.

- **Realdini Studios [🌐]** *March 2025 – May 2025*
  AI Developer Intern Stockholm, Sweden
  - Developed production-ready AI-driven NPC behavior for an RPG shooter using Unreal Engine 5 Blueprints and C++, implementing NavMesh-based dynamic pathfinding and optimized state machine transitions to improve NPC navigation responsiveness.
  - Built core character systems including movement mechanics, input bindings, Animation Blueprints (ABP), and skeletal mesh integrations, then profiled AI performance and debugged behavior inconsistencies across multiple game environments to validate stability.
  - Collaborated via GitHub in sprint-based workflows (branching and PRs), contributing to QA, bug tracking, and internal playtesting for BarrelBound, a 2.5D RPG shooter delivered within a game jam timeframe.
- **Wipro Ltd [🌐]** *October 2023 – June 2024*
  Full-Stack Developer Intern Visakhapatnam, India
  - Developed responsive, component-driven web applications with Angular 19 and Python backend services, reducing page load times by 45% through targeted optimization and caching strategies.
  - Automated 15+ internal workflows using Python scripting, saving approximately 20 hours per week of manual effort; delivered 12+ production features across 8 Agile sprints maintaining 90%+ code quality.
  - Configured automated build and deployment workflows with GitHub Actions and Netlify, including production-grade routing, SPA fallback handling, and serverless function endpoints for demo API simulation.

KEY PROJECTS

- **Legora Trust-Architect [🐙] [🎮]** *February 2026*
  Neuro-Symbolic Legal AI Agent | Python 3.12, FastAPI, Azure OpenAI GPT-4, Pydantic, FAISS, React 19, TypeScript, Docker
  - Designed a hybrid neuro-symbolic architecture combining Azure OpenAI (GPT-4) with Pydantic-based symbolic logic for deterministic legal clause validation, eliminating LLM hallucinations through real-time constraint checking against hard-coded jurisdictional liability caps.
  - Implemented a 4-state agentic lifecycle (Retrieving, Drafting, Validating, Correcting) with recursive self-correction driven by programmatic symbolic engine errors, achieving 100% compliance with legal rules across all test cases.
  - Built a production-ready FastAPI backend with stateless JWT authentication, refresh token rotation, and SlowAPI rate limiting; integrated FAISS vector store for RAG-based legal precedent retrieval.
  - Configured Infrastructure-as-Code deployment via render.yaml and netlify.toml for one-click full-stack deployment; developed React 19 frontend with real-time agent state visualization and dual design modes.
- **Frostbite DevAssist [🐙]** *January 2026*
  RAG CLI Documentation Assistant | Python 3.11+, LangChain, FAISS, Azure OpenAI GPT-4o, Rich, NetworkX
  - Built an end-to-end RAG pipeline over Frostbite Engine Markdown documentation: Markdown ingestion, chunking, FAISS vector indexing, embedding search, and GPT-4o answer synthesis with grounded source citations.
  - Developed CLI workflows for single-query and interactive modes, optional knowledge-graph generation via NetworkX, and usage statistics; enabled natural-language querying over structured technical documentation.
  - Added observability by logging timestamp, query text, response time, and sources to JSON for post-hoc analysis; supported secure Azure OpenAI configuration with repeatable local indexing before first use.
- **Frostbite Agentic Orchestrator [🐙]** *January 2026*
  XML Asset Generator | Python, asyncio, Azure OpenAI GPT-4o, Stability AI, FastAPI, JSON Schema
  - Designed an agentic orchestration pipeline that converts natural-language prompts into validated Frostbite Asset Definitions (XML) and generates matching visual textures via a concurrent multimodal pipeline using asyncio and Stability AI.
  - Implemented hybrid neuro-symbolic validation against a strict Frostbite JSON schema with automated self-correction on validation failure; built cached Demo Mode for zero-latency demonstrations.
  - Delivered CLI, batch processor, and FastAPI web service with rich terminal UX; applied enterprise security patterns including .env isolation, config precedence rules, and sanitized logging to prevent credential leakage.
- **NPC Memory Dialogue System [🐙] [🎮]** *March 2025 – September 2025*
  Emotion-Aware AI Game Platform | Python, FastAPI, PostgreSQL, Mistral 7B, Groq API, RoBERTa, Docker
  - Architected an emotion-aware NPC dialogue framework using FastAPI and a PostgreSQL memory engine with JSON-based memory schemas, enabling dynamic, contextually coherent conversations that evolve over time.
  - Implemented RoBERTa-based sentiment analysis achieving 92% accuracy to influence NPC emotional state; integrated Mistral 7B (GGUF via llama-cpp-python) for local LLM inference and Groq API (Llama 3.1 8B) for scalable production dialogue generation.
  - Built production system with 85%+ test coverage, Docker containerization, OAuth 2.0 (Google/GitHub), JWT sessions with refresh token rotation, rate limiting, and automated Render.com deployment with environment variable management.

References available upon request. Visit portfolio to view more projects and test them out yourself.