

# Kamakshya Nanda

224-591-8916 | [knanda3001@gmail.com](mailto:knanda3001@gmail.com) | [linkedin.com/in/kamakshya-nanda](https://linkedin.com/in/kamakshya-nanda) | [github.com/kamakshyan](https://github.com/kamakshyan)

## EDUCATION

---

**Illinois Institute of Technology**  
*Master of Science in Computer Science*

Chicago, IL  
Aug. 2024 – May 2026

## PROJECTS

---

- Kosh AI** | *Python, FastAPI, DuckDB, QLoRA, Gemma 4, Ollama, React* Feb 2026 – Present
- Fine-tuned **Gemma 4** via **QLoRA** (Unsloth) on curated text-to-SQL datasets, exporting to **GGUF** for local inference via **Ollama** on a self-hosted NUC – achieving sub-second query generation with zero cloud dependency.
  - Designed a **semantic schema pruning** layer that embeds schema metadata and retrieves only contextually relevant fields per query, cutting prompt token overhead by ~60% and significantly reducing LLM hallucination rate.
  - Architected a unified analytics engine using **DuckDB** to virtualize heterogeneous data sources (PostgreSQL, MySQL, CSVs) into a normalized SQL layer, enabling cross-format natural language querying for non-technical users.
- PHInsight** | *Python, FastAPI, Presidio, scispaCy* | *ScarletHacks 2026 – Largest Market Opportunity* Mar 2026
- Built a **drop-in FastAPI proxy** that intercepts outbound LLM API calls, scanning and redacting **PHI** in real time using **Microsoft Presidio** and **scispaCy** before payloads leave the network perimeter.
  - Designed an **audit trail** backed by PostgreSQL, logging every detection event with entity type, confidence score, redaction action, and timestamp – enabling compliance reporting without storing raw patient data.
  - Architected the full proxy pipeline (request interception → PHI scan → selective redaction → forwarded completion → response reconstruction) as a **zero-config drop-in** requiring only an endpoint swap.
- Privacy-Preserving Medical AI** | *Python, PyTorch, Flower (flwr), Federated Learning* April 2025 – Present
- Developed a distributed brain tumor detection system using **PyTorch**, training on CT/MRI scans across **10+ simulated client nodes** via the **Flower** federated learning framework - raw patient data never leaves local nodes.
  - Implemented custom classifier heads per client with a shared global aggregation strategy, improving global model accuracy by **12%** over the federated baseline despite non-IID data distribution across nodes.

## EMPLOYMENT

---

- Software Engineer** Feb. 2024 – July 2024  
*C-BIA Consulting*
- Built a cross-platform mobile app in **.NET MAUI** for an education client; refactored core modules to reduce binary size and improve cold-start performance on iOS and Android.
  - Replaced a manual Excel-based settlement process for a banking partner by building a **SQL ingestion tool** that parses **XML** transaction logs – eliminating data entry errors across thousands of weekly records.
  - Extended virtual classroom tooling with **Vue.js**, building a Zoom plugin enabling teachers to overlay images and video onto their live camera feed in real time.

- Software Developer** May 2023 – Aug. 2023  
*Proteus Technologies*
- Automated data preparation workflows for the analytics team using **Python**, integrating with **Dremio** to eliminate hours of manual ETL work per week.
  - Modernized a legacy BI dashboard using **Angular** and **D3.js**; optimized chart rendering logic to handle large datasets without blocking the UI thread.

## TECHNICAL SKILLS

---

**Languages:** Python, JavaScript, TypeScript, SQL, C++, C  
**AI / ML:** PyTorch, QLoRA, Unsloth, Ollama, Presidio, scispaCy, Flower (flwr), LLM fine-tuning (GGUF)  
**Frameworks:** FastAPI, Next.js, React, Angular, Vue.js, Node.js, Express.js, .NET 6  
**Infrastructure:** Docker, Tailscale, Caddy, Cloudflare DDNS, GitHub Actions, Vercel, Azure  
**Databases:** PostgreSQL, DuckDB, MySQL, MongoDB  
**Certifications:** AWS Academy – Cloud Foundations