

# Panambur Pawan Bhandarkar

Redwood City, CA

Email : pawan@bhandarkar.me

Mobile : +1(650)537-7341

Portfolio: [bhandarkar.me](https://bhandarkar.me)

Github: [bhandarkar.me/github](https://github.com/bhandarkar)

Linkedin: [bhandarkar.me/linkedin](https://linkedin.com/in/bhandarkar)

## Summary

---

Software engineer with 4+ years of experience building production AI infrastructure. At xAI, owned and scaled the model debugging and RL observability infrastructure as well as the human data platform that directly accelerated Grok model iterations.

## Experience

---

**xAI** Palo Alto, CA  
*Member of Technical Staff* Dec 2024 - Present

- **Nebula Platform (Tier 0)** - Built a suite of internal tools for end-to-end model development and debugging, including file viewers, specialized trace viewers, benchmark queuing and automation services, production data debuggers, model deployment automations etc - xAI now spends over 7000 engineering hours per week on Nebula
- **RL Data App** - Subsystem of the Nebula Platform delivering high-performance aggregation and analytics tools to track and debug RL training and evaluation metrics—now integral to all RL workflows at xAI.
- **Starfleet** - Owned and scaled xAI's human-data platform powering human evals and high-quality data generation for post-training Grok, with features including multi-model arenas, QA pipelines, real-time analytics and performance management, auto-graders, HRIS-style operations management, and more.
- **Research Agent** - Built xAI's AI assistant with cursor-style UX hosted on Nebula Platform, enabling large-scale LLM-based analytics and data processing on training and evaluation data from RL runs using an MCP server with specialized model-debugging tools.
- **Devbots** - Led a team of 3 to design and ship xAI's Remote Agents infrastructure, boosting engineering productivity by enabling all users to interact with agents via Slack App, Web UI, PR reviews, and more, to spin up sandboxed coding environments for feature development.
- **Nebula Notebooks** - (Solo) Designed and shipped an in-house version of Google Colab with managed Jupyter kernels, custom UI/UX, and bespoke cell renderers specialized for visualizing agent traces, grader results, and more.
- **Infrastructure & Observability** - Built infrastructure and observability solutions such as Grafana dashboards and alerts to maintain high reliability and uptime of all critical platforms mentioned above.

**Converge Lab** Sunnyvale, CA  
*Founding Engineer* Jan 2024 - Dec 2024

- Built and launched an LLM-powered hardware co-designer with generative UI, human-in-the-loop feedback, and tiered access control using Next.js, React, and TypeScript. Helped secure a partnership with Nvidia.
- Developed a FastAPI RAG backend with Pinecone, Elasticsearch, and LangChain for semantic retrieval across electronic component datasheets, with DynamoDB for persistence.
- Built Python pipelines for extracting and parsing PDF datasheets into structured data for LLM context windows.

**TuneIn Radio** San Francisco, CA  
*Software Engineer Intern, Platform* Jun 2023 - Aug 2023

- Engineered real-time data pipelines for **7.5M monthly users**, processing 2M+ messages/day at sub-30ms latency using WebSockets, Apache ActiveMQ, and Redis Streams. Automated deployments with Terraform and Kubernetes

## Skills

---

**Languages & Frameworks** : TypeScript, Python, React, Next.js, Node.js, Go, FastAPI, GraphQL

**AI/ML Infrastructure** : RLHF pipelines, model evaluation, LangChain, RAG, Pinecone, Elasticsearch

**Data & Observability** : PostgreSQL, Redis, ClickHouse, DynamoDB, IndexedDB, Sentry, Axiom, D3/Recharts

**Infrastructure** : Kubernetes, Docker, AWS, GCS, gRPC, WebSockets, CI/CD, Terraform

## Education

---

**Carnegie Mellon University** Mountain View, CA  
*Master's in Software Engineering* Dec 2023

**Nitte Mahalinga Adyanthaya Memorial Institute of Technology** Nitte, India  
*Bachelor's in Computer Science and Engineering* May 2020