

Jae Park

(224) 433-0362 | jaepark.jbp@gmail.com | jaepark.com | [linkedin/jaepark](https://linkedin.com/in/jaepark) | [github/jaepark23](https://github.com/jaepark23)

EDUCATION

University of Illinois Urbana-Champaign

B.S. in Computer Science and Statistics

Expected May 2027

GPA: 3.91/4.00

WORK EXPERIENCE

AMD

Software Engineer Intern

Austin, TX

May 2025 – August 2025

- Spearheaded a CPU trace system for EPYC processors, accelerating post-silicon validation and debugging by 30%
- Built a Rust decoder for binary CPU trace data, achieving 33% faster parsing than legacy tools
- Delivered a full-stack trace visualization platform in FastAPI, React, and Docker, allowing post-silicon debug engineers to collaboratively analyze EPYC traces
- Cut AWS S3 storage costs by 70% via Parquet compression on trace files
- Wrote a Python tracing library with Poetry that abstracted trace capture complexity and enabled automation

printpal.io

Software Engineer

Chicago, IL

May 2023 – Present

- Developed a social platform for 100,000+ 3D printing users using Flask, JavaScript, SQL, and Docker, supporting model upload, sharing, and marketplace sales
- Implemented lazy-loading with Redis caching for infinite scrolling, reducing API overhead by 30%
- Utilized Celery for distributed task processing including Redis cache updates, health checks, and automated cleanup
- Built platform wide search using Elasticsearch for accurate retrieval of users and models
- Spearheaded a Swift mobile app for AI generated 3D models from text or images, with in-app model preview rendering

FANUC America

AI Engineer Intern

Hoffman Estates, IL

May 2024 – August 2024

- Architected an on-premises RAG chatbot using LangChain that reduced technical document lookup time by 80%, achieving 90% retrieval accuracy on internal benchmarks
- Developed a PDF chunking pipeline in C# with iterative evaluation to optimize chunk size and retrieval quality
- Converted real production queries from engineering and domain experts into evaluation benchmarks used to validate system performance

PROJECTS

Schrodinger | *Python, TensorFlow/Keras, FastAPI, React, AWS*

January 2026 – March 2026

- Deployed a FastAPI backend on Dockerized EC2 serving live predictions with real-time NBA stats and injury data
- Architected an AWS pipeline (DynamoDB, S3, EventBridge + Lambda) that generates predictions and resolves them against real outcomes
- Trained a Keras neural network on 20 years of NBA data (150+ features per team) with PCA and standard scaling, achieving an accuracy of 70%
- Built a React and Material UI frontend on S3 behind CloudFront with prediction history and accuracy tracking views

Customer Discovery Platform (Disruption Lab) | *React, TypeScript, Next.js, OpenAI*

Sep 2025 – Dec 2025

- Built an AI-driven platform to help founders validate customer pain points through structured, multi-turn simulated interviews
- Implemented a conversation pipeline using the OpenAI API to model realistic discovery flows with repeatable, consistent qualitative feedback
- Developed a responsive React frontend to manage interview sessions, state, and real-time responses

SKILLS

Languages: Python, Rust, JavaScript/TypeScript, SQL, C#, Swift, C++

Frameworks/Libraries: FastAPI, Flask, React, Django, Next.js, Pandas, WebAssembly, SwiftUI, LangChain, Keras

Databases: MySQL, MongoDB, DynamoDB, Firebase

Tools: Git, Docker, AWS, Redis, Elasticsearch, Celery