

Keegan Choudhury

keeganarko@gmail.com | 469-777-0330 | [linkedin.com/in/keeganchou](https://www.linkedin.com/in/keeganchou) | github.com/keeganarko

EDUCATION

Texas A&M University, College Station, TX

Graduated May 2023

Bachelor of Science in Computer Science; Minor in Video Game Design

GPA: 3.96

EXPERIENCE

Capital One

McLean, VA

Product Manager

August 2025 - Present

- Lead **2 engineering teams** driving distinct initiatives on Graffiti 2.0, Capital One's card delivery platform processing **4.3M card requests per month**, migrating core infrastructure from batch file processing to real-time streaming
- Replacing TSYS, a **\$12.5M/year** third-party file processor, with in-house capabilities on Graffiti 2.0 — removing the largest external dependency from **Capital One's card issuance pipeline**
- Defining the platform onboarding pattern for a new co-branded premium credit card on Graffiti 2.0, establishing the reusable model that scales the platform from single-purpose to Capital One's centralized card issuance system
- Drive alignment across **4 partner organizations** (Operations, Engineering, Delivery, Product) on competing platform priorities by quantifying production impact, unblocking critical roadmap investments

Microsoft

Redmond, WA

Product Manager – Microsoft AI Development Acceleration Program (MAIDAP)

July 2023 - August 2025

- Delivered 6 zero-to-one LLM-powered initiatives across Copilot Studio, Excel, and internal AI tooling during MAIDAP, Microsoft's AI/ML rotational program — producing 2 filed patent applications, a KDD 2025 publication, and an open-source contribution to Microsoft Semantic Kernel
- Led an LLM-powered Excel formula repair initiative as primary PM for a 12-person engineering and research team; co-authored a filed patent application and a KDD 2025 paper introducing FoRepBench, a 618-sample benchmark dataset for context-aware runtime error correction (**GPT-4o, GPT-4.1, Phi-3, and Mistral**)
- Owned strategy for an agent framework enabling goal-oriented AI-human conversations; validated through 10+ user interviews and POC demos, open-sourced via Microsoft Semantic Kernel (28k+ stars), adopted by 3+ teams
- Shipped an LLM-powered support assistant for automated root-cause detection (projected ~\$1.07M annual savings); designed end-to-end UX and the manual LLM evaluation framework measuring response quality

INTERNSHIPS

Tesla

Austin, TX

Technical Program Manager Intern

September 2022 – December 2022

- Automated part verification on Tesla's Model 3 door assembly line, increasing production rate per shift by ~23% by deploying real-time installation-value validation at each workstation
- Built a factory KPI dashboard tracking throughput, inter-station cycle time, and production bottlenecks, driving leadership prioritization of automation investments on the most-clogged production lines

Microsoft

Redmond, WA

Program Manager Intern

May 2022 - August 2022

- Conducted competitive analysis of HoloLens 2 against MR/AR competitors (Meta Quest, Magic Leap, emerging Apple AR/VR), surfacing compute and compatibility gaps that informed a PRD and product strategy deck

AT&T Technologies

Dallas, TX

Technical Program Manager Intern

June 2021 - August 2021

- Defined product roadmap for an internal hiring tool that reduced engineering attrition by 12%
- Led a team of 3 interns to ship an internal PR-assignment tool (React/Django) **adopted by 10 engineers**

LEADERSHIP

Microsoft AI Lunch & Learn Series

December 2023 - May 2025

- Drove AI tool adoption and cross-industry AI literacy at Microsoft by curating a recurring Lunch & Learn series; recruited external and internal speakers showcasing emerging AI applications across multiple industries

SKILLS

AI/ML: LLM evaluation, prompt engineering, agent frameworks (Semantic Kernel), benchmark design, A/B testing

PM & Tools: Product strategy, customer research, GTM, technical PRDs | Figma, Power BI, SQL, Azure DevOps, Python