



FREYA (YE) YU

Technical AI Product Manager · Responsible AI, Governance & Compliance

Seattle, WA · +1 206-850-3068 · yeyufreya@gmail.com · linkedin.com/in/yeyufreya · yeyufreya.com · github.com/yeyfreya

SUMMARY

Technical AI product manager building human-first AI systems people can trust. I turn governance, compliance, and safety requirements into product mechanisms — source provenance, explicit uncertainty, human review authority, and stated system limits — across product strategy, AI system architecture, and 0-to-1 delivery. Founder of an AI compliance product; joint inventor on a pending U.S. patent application.

CORE COMPETENCIES

AI Product — Product strategy · 0-to-1 delivery · Problem framing · Prioritization · User research and formative evaluation · Technical roadmap and handoff · Sponsor and stakeholder alignment

Responsible AI & Governance — AI risk and compliance analysis · Source provenance and primary-source verification · Human-in-the-loop review design · Release gating and claim discipline · EU AI Act, GDPR, HIPAA (applied research)

AI Systems & Technical — LLM orchestration and tool use · Web-retrieval pipelines · Multi-turn evaluation harnesses · Tracing and run-level provenance · Python · TypeScript · Claude API · Tavily · Streamlit · Supabase · Langfuse · OpenTelemetry · Firebase · Git

EXPERIENCE & SELECTED PROJECTS

Founder & AI Product Manager — AI Compliance Gap Analyzer

Seattle, WA · Feb 2026 – Present

Founder-led AI compliance research and decision-support product for small teams with ongoing obligations and no compliance function.

- Own product strategy, architecture, and AI-assisted delivery of a four-stage Python pipeline (scope, research, analysis, report) that maps regulatory territory across five dimensions — jurisdiction, profession, data, activity, vendor — runs official-domain-first web research, and returns a reviewable report with explicit assumptions, sources, and a mandatory human-review boundary.
- Made evidence integrity a release decision: when an AI-assisted audit across two identical-input v0.6 reports found **4 of 7 material claims wrong or misattributed**, held v0.6 and drove a scope-before-search redesign instead of prompting the model to sound more authoritative.
- Improved coverage in a small identical-input test: the governing professional body was absent from findings in 2 of 6 v0.6 runs, then present and cited to its own domain in 3 of 3 v0.7 reruns; instrumented run-level provenance and documented unresolved source-tier leakage and ranking instability.

Product Design Engineer (AI & Digital Experience) — Rose Creative Works

Part-time · Seattle, WA · May 2025 – Present

- Design and ship improvements to a client-facing web platform for a Seattle creative studio, translating ambiguous stakeholder requests into scoped product flows and clearer end-to-end user journeys.

AI Product Manager, Responsible AI — Responsible Innovation Labs

Seattle, WA · Sep 2024 – Mar 2025

Industry-sponsored capstone delivered through the University of Washington's Global Innovation Exchange, with RIL as sponsor and product stakeholder.

- Led product vision, research strategy, problem framing, prioritization, and end-to-end experience for a four-person team, reframing an ambiguous Responsible AI brief into a trust-centered transparency and decision-support prototype for early-stage AI teams.
- Moderated a seven-person formative evaluation and converted findings on sensitive-data handling, report provenance, input-to-output mapping, score clarity, and guidance discoverability into prioritized product changes; the round reported 85% task completion and 4.2/5 satisfaction.
- Defined the intended AI workflow and delivered a hosted low-code prototype (Webflow, Zapier, Firebase, rule-based scoring, model-assisted generation) with editable reporting and handoff documentation, while documenting that the assessment layer was not production-ready.

Product Manager & Product Designer — UW Global Innovation Exchange

Seattle, WA · Sep 2023 – Mar 2025

- Delivered interdisciplinary graduate product work across consumer, education, and affective-AI contexts; built and deployed a speech emotion recognition application end to end, integrating TensorFlow models with a custom Librosa audio pipeline on Azure and Hugging Face.

NemoSafe — Technical AI Product Manager, Evaluation & Orchestration

NVIDIA Spark hackathon · Aug 2026

- Led evaluation orchestration for a four-person, 48-hour team studying multi-turn relational risk in teen-facing AI; implemented the controller, staged scoring, reporting, packaged harness, tests, and configs, preserving unparseable evaluations as `insufficient_data` rather than passes.
- Team analysis of one 71-conversation, 2,840-turn synthetic run found 42 conversations with no stock per-turn taxonomy flags but at least one uncalibrated trajectory flag; kept 13 unparseable results visible and framed it as a measurement gap, not a safety verdict.

EDUCATION

Master of Science in Technology Innovation — University of Washington, Global Innovation Exchange · Sep 2023 – Mar 2025 · GPA 3.85 / 4.00

Master of Urban Planning; B.E. Urban & Rural Planning — Tongji University, Shanghai · 2014 – 2023

PATENTS, PUBLICATIONS & COMMUNITY

Joint inventor — “Auditing system for built environment of age-friendly street based on multisource big data,” pending U.S. application US20240362482A1 (filed Jan 2023), with CN114444941A and WO2023143000A1.

First author — “How Perceived Neighborhood Built-Environment Attributes Influence Parental License for Children’s Independent Mobility,” 56th ISOCARP World Planning Congress, 2021.

Co-author — “Stress Response and Network Resilience of Communities During the Public Health Emergency: A Study Using Photovoice Method,” 2020.

Founder, AI + Sun — Seattle gathering series for AI builders who value both innovation and wellbeing.

Community Steward, AIGovOps Foundation — Volunteer role with a practitioner community for AI governance and responsible AI.