Aranganathan Viswanathan, PgMP®

Senior AI Research and Data Program Manager• AI/ML Technology Product Operations
Santa Clara, CA | (513) 417-3374 | aranganathan.vishwanathan@gmail.com | LinkedIn Profile
160M+ Business Impact | 30M+ Labeled Assets | 9 AI Product Launches | 700+ Annotators Scaled | 50+ ML Programs

Summary

Senior AI technology program and product operations manager with 12+ years of experience translating product vision into operational execution - spanning data strategy, LLM evaluations for AI/ML products across research, safety, media, finance, payments, and autonomous vehicle industries.

I have driven technology program management and product operations for more than SO+ AI/ML initiatives, with VP-level visibility, contributing to more than $160M in revenue, embedding Responsible AI and safety principles into all phases of product development.

I have a reputation for turning chaotic AI development workflows into high-performance, repeatable systems that improve model quality, accelerate experimentation, and deliver measurable business results.

Professional Experience

Adobe Inc., San Jose, CA — May 2026 - Present

Senior AI Research and Data Program Manager - Brand Intelligence (ABI)

Joined Adobe's Brand Intelligence (ABI) organization as Senior Research Program Manager, partnering with ML, Product, and Forward Deployed Engineering (FDE) teams to help scale how AI-driven brand solutions are delivered to enterprise customers.

- Drove cross-functional alignment between ML, Product, and Forward Deployed Engineering (FDE) teams at Adobe's Brand Intelligence org, building direct working relationships across previously disconnected functions to close visibility and ownership gaps.
- Identified and diagnosed a core scaling gap in FDE operations, that customer-specific solutions were being duplicated separately by engineering instead of reused, through direct stakeholder interviews, then facilitated a structured review with FDE leadership to resolve ownership and next steps.
- Led a multi-round stakeholder review process for a proposed ML-FDE operating model, incorporating direct feedback from FDE leads to correct scope and avoid overstating findings, resulting in a leadership-aligned, evidence-based proposal rather than a top-down mandate.
- Scoped and championed an AI-assisted internal knowledge tool (RAG-based chat interface) to reduce time spent re-explaining customer context across teams, partnering with engineering and leveraging existing internal infrastructure.
- Served as the connective layer between program management and ML leadership, routing and scoping AI-related asks from cross-functional stakeholders so they reached the right team with the right context, reducing ad hoc escalations to senior ML leadership.

Amazon Web Services, Bellevue, WA

Senior Program Manager (L6) - AI/ML Evaluations & Product Operation — Jul 2022 - Present

As Sr. Program Manager {L6} in AWS's 300-person Computer Vision division - delivering computer vision software, platforms, and services - I owned the end-to-end AI/ML program portfolio across data strategy, model evaluation, governance, and go-to-market - managing a $17M+ annual budget across 50+ programs and partnering with Applied Scientists, Engineering, Data Operations and VP-level stakeholders to drive model releases and enterprise adoption.

- AI teams planned product, data, and model work independently, causing late dependency discoveries and delayed releases. I introduced cross-team governance with weekly ML program reviews, shared operational dashboards, and structured influence mechanisms that aligned research and engineering teams on model prioritization and sequencing. **Delivery delays dropped 30%, AI timelines improved 20%,** and the model became the standard operating framework.

- Our multi-modal AI data collection and annotation pipelines produced only Sk low-quality assets monthly, limiting model accuracy to 83% and driving churn and missed revenue targets. I built a new data platform with quality filtering - spanning text, image, video, audio, and LiDAR across 20+ languages and sensitive content domains - integrated synthetic data generation to fill collection gaps and onboarded premium vendors to unblock model training pipelines. **Throughput scaled lOx, model accuracy reached 94%, and revenue growth rebounded.**
- Our VI GenAI chatbot produced hallucinated responses and lacked reliable post-training evaluation, delaying launches and risking Responsible AI violations. I built a scalable evaluation framework combining LLM-as-a-judge, expert datasets, and human review. **Hallucinations dropped 50%,** iteration speed doubled, and improved models drove 23% customer usage.
- AI data programs ran without visibility into labeling, training, and evaluation capacity, causing late bottlenecks and unpredictable launches. I built a capacity planning and forecasting framework with capacity models, dashboards, and cross-team planning. **Release cycles fell 20%,** delivery became predictable, and faster launches drove **35% enterprise adoption.**
- Datasets were created late in the machine learning lifecycle, causing inconsistent labels, repeated relabeling, and stalled experiments that delayed product releases and hurt competitiveness. I transformed the ML data workflow by embedding dataset design into model planning and standardizing data pipelines. **Rework fell 30%, evaluation cycles sped up 25%,** and ML features reached customers faster.
- AI teams lacked standardized dataset lineage, versioning, and access controls, risking compliance exposure and unreliable model training. I led creation of a data governance framework - authoring the BRD, defining metadata standards, lineage tracking, and policy-based access management. **Dataset validation effort fell 40%,** strengthening compliance and confidence in AI training data.
- AWS's go-to-market for AI labeling was outdated, stalling deals and slowing sales cycles while we lost winnable enterprise opportunities. I led discovery, exposed product and messaging gaps, secured $3M funding, and revamped playbooks and feedback loops. **Sales accelerated 25%, revenue grew 25%+,** and stalled enterprise deals closed.
- AI launches relied on complex ML pipelines with risks surfacing late, and operations teams lacked launch readiness - causing delays, remediation costs, and low adoption. I built a structured risk and launch management framework with risk registers, readiness reviews, escalation protocols, and feedback systems. **Late-stage risks fell 35%, post-launch incidents dropped 30%,** and the governance model became the template for future launches.
- AI products required massive labeled datasets, but fragmented annotation teams couldn't scale, delaying model training and product launches. I built a 0-to-1 global ML data operations capability with 700+ annotators and structured workflows. The operation **delivered 21M high-quality assets,** enabled nine AI launches, and drove $64M impact.
- Annotation demand outpaced internal capacity and vendor performance was inconsistent, delaying datasets and slowing AI model development. I built a governed global vendor ecosystem with structured onboarding, pilots, and performance metrics. **Capacity increased 4x,** dataset bottlenecks cleared, and product teams accelerated model training cycles.
- Inconsistent annotation workflows caused labeling errors, dataset rework, and weaker model performance. I standardized annotation processes, implemented data quality frameworks with inter-annotator agreement tracking, and built operational dashboards. **Agreement exceeded 98% (above the 95% industry standard), productivity rose 28%, and efficiency improvements reduced costs by $6.8M.**
- Our content moderation models generated excessive false positives, forcing customers into manual reviews and eroding trust. I led a model quality overhaul, refining training datasets, tuning decision thresholds, and introducing evaluation metrics. **False positives dropped 3x,** our models outperformed competitors, and adoption doubled revenue with **$12M incremental growth.**
- Leadership was flying blind on ML data operations, relying on error-prone spreadsheets to track annotation throughput, quality, and delivery risks. I replaced manual reporting with SOL-powered dashboards and a unified OKR/KPI framework linking data operations to model delivery. Reporting errors were eliminated, **operational efficiency improved 35%,** dataset reliability reached 95%+ SLA compliance.
- Our ML data teams were growing fast but lacked clear career paths, risking loss of experienced annotators and operational knowledge. I led creation of a structured career ladder, skill tracks, and training programs with HR and leadership. **Turnover fell 4x, 40+ employees advanced into senior roles,** and operational expertise strengthened.

Core Competencies

LLM & **GenAI:** Data collection & annotation • Model evaluation • RLHF & SFT • Post-training programs • Foundation model platforms• RAG• Human-in-the-Loop • LLM-as-a-Judge • Responsible AI • Synthetic data

Product Operations: Roadmapping • Prioritization • PRDs/BRDs • ML Lifecycle Management • Capacity Planning • Data Pipeline Management• Data Quality Frameworks • Crowdsourcing Platforms (MTurk) • Technical Program Management• Agile/Scrum • PDLC •Stage-Gate• Jira & Confluence• SQL • Python • KPI Dashboards• Product analytics

Certifications

- PMI Program Management Professional (PgMP)®
- AWS Cloud Practitioner
- AWS AI Practitioner

Education

University of Madras, Chennai, India

- Master of Computer Applications (Computer Science)
- Bachelor of Science (Applied Physics)