Tiffany Peil

AI Governance, Assurance and Accountability · Engineering Risk and Technical Program and Operations Research Leadership

San Francisco, CA · linkedin.com/in/tiffanypeil

Summary

13+ years across regulated engineering, safety-critical verification and validation, operational risk, and identity and authorization program management. Experience includes verification and validation for FDA-regulated medical devices, QA and operational-risk evaluation for autonomous-vehicle and robotics prototypes, engineering risk management across four Blue Origin business units, and program management for identity infrastructure at Adobe.

Since 2025, focused on AI assurance, including agent-deployment assessment, authorization boundaries, record integrity, decision lineage, and independent review and contestability. Designed an assessment method for AI agent deployments (v0.1) and directed implementation of a Python scorer applied to two internal systems. A v0.2 extension proposes an independent review architecture (design stage).

Experience in empirical test design, requirements-based validation, operational-domain coverage assessment, controlled fault testing, and evidence-based interpretation of technical findings.

Selected research and technical work

Assessment method and scorecard for AI agent deployments `v0.1, Oct 2026`

An impact assessment method and scorecard at the deployment and agent layer, using points of relation and points of control (PoR/PoC) as the working unit of analysis: 56 metrics in 10 domains, maturity levels capped by evidence grade, release gates that fail closed, and no aggregate score. A Python scorer and a companion spreadsheet return the same release decisions on every case tested (one example assessment, three variants and both internal assessments). Not validation of the metrics or thresholds; not yet a benchmark or certification.

Public proxy and independent review architecture `v0.2 design, Oct 2026`

Proposes that publicly governed AI proxies could make independent assurance and contestability more accessible, given a lawful mandate, protected access and human accountability. Separates roles for representing affected people, reviewing claims and enforcing law. Specifies an access protocol from mandate to revocation, proposed metrics (including reviewer independence, evidence noninterference, effective challenge and remedy enforceability), a threat model covering reviewer capture and data exfiltration, and proposed validation studies. AI prepares and tests; humans and institutions keep legal decision authority. PoR/PoC and proxy oversight remain research hypotheses; thresholds are uncalibrated, and nothing has been piloted or validated.

Agent authority and spec-to-behavior evaluations `Design stage`

Specified evaluations for recommendation-versus-execution boundaries, delegation scope, provenance across handoffs, and refusal and escalation, plus a pipeline that turns published model specifications into traceable behavioral tests. Not yet implemented or run against baselines.

Working paper `In progress`

Information architectures as algorithmic governance systems: how the structure of information systems conditions decisions and encodes values at scale.

Provisional U.S. patent application

SHD-CCP, co-inventor. Application No. 63/876,451, filed Sept 5, 2025 (Cooley LLP). Covers context preservation, verification, and topology-preserving representation in AI systems.

Experience

Polycentria (Emotivity Inc.) · San Francisco, CA · Co-founder

Research-stage company developing traceability, verification, and decision-record systems for consequential AI.

Head of AI Governance & Operations Strategy — Apr 2025 – Present

- Concluded that the evidence for product readiness wasn't there, and moved the company from near-term commercialization to research and governance.
- Directed the AI coding agents that built the scorecard tooling, and applied the scorecard to two internal systems: the company's site and its AI agent pipeline, and a record-integrity mechanism.
- In the agent pipeline, agents held permissions for high-impact infrastructure actions that their written rules never authorized, and there were gaps in agent identity and injection testing. Nothing measured whether review of hundreds of agent commits caught errors.
- Because the first two internal assessments used a non-independent AI assessor working for the owner, and the agent pipeline was mostly AI-written, made independent re-scoring a release-blocking gate. Both assessments currently fail it. This led to the v0.2 independent review design.
- Designed and conducted seeded-alteration tests of a record-integrity mechanism, comparing detection of accidental edits with owner-controlled record rewrites. Interpreted the observed failures and narrowed the supported assurance claim.
- Conducted a preliminary coverage review of selected frameworks, benchmarks and vendor approaches at the deployment layer, including frontier lab safety frameworks, the EU GPAI Code of Practice, AgentDojo and AgentHarm. Mapped each metric against relevant provisions and controls in NIST AI RMF, the EU AI Act (Regulation (EU) 2024/1689 as amended by Regulation (EU) 2026/1744), ISO/IEC 42001 and 42005, and the OWASP Top 10 for Agentic Applications 2026. The mapping doesn't establish legal conformity or certification.
- Reviewed precedents for independent audit and protected data access (including EU DSA Articles 37 and 40, EU AI Act Article 74 and the FTC's 2019 Facebook order) and mapped possible routes for mandating independent review. This isn't legal analysis and hasn't had counsel review.
- Set the company's evidence standard: technical documentation, investor materials and demos label each claim as a demonstrated result, a hypothesis, planned work, a limitation or an open question.

Founding Chief Operating Officer — Apr 2023 – Apr 2025

- Built the internal AI agents and API pipelines that ran the company's research workflows.
- Built legal, financial, fundraising and research operations: corporate formation, outside counsel management and patent strategy, and investor materials and fundraising infrastructure.
- Ran customer discovery and market research across healthcare, financial services and regulatory compliance.

Udemy · San Francisco, CA

Senior Technical Program Manager (Contract) — Jul 2022 – Apr 2023

- Ran program operations for six product launches in nine countries over seven months, coordinating Product, Engineering, Trust & Safety, Legal, Localization and Go-to-Market (50+ stakeholders). Delivery velocity rose 35% with no added headcount.

Adobe · San Francisco, CA

Engineering Program Manager, Identity Platform (Contract) — Feb 2022 – Jul 2022

- As program manager, ran the program for identity infrastructure (authentication, authorization, account management and security) serving 26M+ users. Reported to the VP of Software Engineering.
- Built the risk-assessment, escalation and incident-response frameworks for a monolith-to-microservices migration. The migration was completed without major incident.

Blue Origin · Kent, WA

Senior Technical Program Manager; Interim PMO Director Oct 2020 – Aug 2021

- Owned the engineering risk-management process across four business units (New Shepard, New Glenn, Engines, Advanced Development), including the HLS/Artemis proposal. Reported to the SVP of Safety & Mission Assurance.
- Mapped failure modes, ownership gaps and inconsistent risk processes; designed a single risk process, secured CEO approval, and piloted it in all four units within six months. It was then adopted as the standard across the organization, which had more than 3,000 employees.
- Analyzed risk-record quality and operational metrics across four business units to identify inconsistent practices, ownership gaps and evidence deficiencies; represented findings through structured reporting to support process design and leadership decisions.
- Managed six technical program managers as Interim PMO Director.

Fresh Consulting · Bellevue, WA

Integration Specialist, QA Engineer, and Program Manager (Contract) Apr 2019 – Nov 2019

- QA, systems integration and operational-risk evaluation for autonomous-vehicle and robotics prototypes, covering field readiness and solar-site deployments.
- Translated autonomous-system behavior and perception requirements into field-data collection specifications, labeling criteria and operational-design-domain coverage requirements; assessed metadata quality, scenario coverage and edge-case gaps before subsequent model training and evaluation.
- Analyzed distributions of observations across operating conditions and scenario classes to assess dataset representativeness, identify underrepresented conditions and interpret model-evaluation findings against the intended operational domain.

Philips Healthcare, Volcano Corporation, and NeuroVision Medical

System Test Engineer Aug 2012 – Oct 2017

- System test, verification, validation and quality engineering for FDA-regulated cardiovascular and ophthalmic devices under FDA design controls (21 CFR 820): requirements traceability, validation evidence, failure analysis, change control and release. Also worked on regulatory submissions and audits.
- Executed requirements-traceable verification and validation of FDA-regulated medical devices, documenting test outcomes, failure analysis, acceptance criteria and change-control evidence for release and regulatory review.

Other role: Apptio, Solutions Engineer and Account Manager (Jan 2020 – Oct 2020).

Education and certification

- B.S., Management Information Systems, University of San Francisco, 2016, summa cum laude
- Certified TBM Executive, Technology Business Management Council

Additional education

- PhD, Organizational Psychology and Behavior, University of Otago: admitted, deferred
- MBA, Sustainable Entrepreneurship, Meridian University: admitted, deferred