

APURV GAURAV

AI PRODUCT & PLATFORM LEADER | EVALUATION · DEVELOPER SYSTEMS · RELIABILITY

Camden, NJ | 412-523-8786 | apurvgaaurav@gmail.com | apurvgaaurav.ai | linkedin.com/in/apurv-gaurav | github.com/apurvgaaurav

EXECUTIVE PROFILE

AI product and platform leader combining 12+ years of release governance, systems reliability, technical execution, and field evidence with patent-backed applied AI product development. Inventor of U.S. Patent 12,670,085 B1 and creator of ReplayGuard, a working evaluation and authorization system for automated code remediation. Applied research spans provenance verification, human-review routing, safety measurement, and reliable AI execution.

SPECIALTY — Turning ambiguous, high-consequence problems into product requirements, metrics, evaluation plans, rollout controls, and cross-functional operating mechanisms.

SELECTED EVIDENCE

Issued U.S. Patent

US 12,670,085 B1 · issued June 30, 2026

GAISS 2026

Accepted paper · Industrial Advisory Committee member

Scaled Operating Foundation

Release governance and reliability within defined subsets of Comcast's publicly reported 75M-device platform

Working AI Product

ReplayGuard evaluator + public reference implementation

Peer-Reviewed Journal Research

IEEE Access + Journal of Systems Architecture manuscripts under review

Graduate AI Research

Penn State D.Eng. in Artificial Intelligence candidate

AI PRODUCT LEADERSHIP EVIDENCE

- Problem framing — Framed product-control problems across automated remediation, multimodal provenance, human-review routing, and heterogeneous inference runtimes.
- Product strategy & roadmap — Built an independent AI product portfolio spanning remediation authorization, decision replay, provenance, runtime controls, and evaluation workflows, with buyer/user framing, success metrics, integration hypotheses, and roadmap tradeoffs.
- Evaluation & metrics — Designed deterministic and reliability-aware evaluation approaches spanning code remediation, provenance verification, human-review routing, and deployment evidence.
- Rollout operating model — Helped shape stage-based Remote Feature Control rollout across Broadband, Video, Cameras, and xFi Pods, including readiness evidence, dashboards, deployment decisions, and escalation.
- Cross-functional execution — Worked across Product, development, QA, analytics, operations, and field stakeholders on prioritization, validation, technical evidence, and release decisions.
- Commercialization — Built ReplayGuard technical evidence, executive materials, integration framing, and acquisition/licensing pathways around an issued U.S. patent.

SELECTED AI PRODUCT & PLATFORM WORK

ReplayGuard | Independent Product Lead & Inventor | 2026–Present

- Translated an issued patent into a working AI developer-system product with requirements, evaluation workflows, public implementation, executive materials, integration framing, and commercialization pathways.

PCS-R | Reliability-Aware Provenance & Human-Review Routing | 2026–Present

- Framed which image-text provenance cases can advance automatically versus route to human review; designed reliability-aware evaluation logic and decision-routing criteria for generative-AI workflows.

Remote Feature Rollout Operating Model | Comcast | Jul 2018–2025

- Helped shape stage gates, readiness evidence, monitoring, feature approval, deployment priority, dashboards, communication, and issue response across Broadband, Video, Cameras, and xFi Pods.

PROFESSIONAL EXPERIENCE

Field Triage Engineer | Comcast | 2025–Present

- Investigate high-impact broadband and Wi-Fi issues using telemetry, device evidence, logs, controlled reproductions, and system state; separate plausible explanations from evidence strong enough to support a defect or release decision.
- Translate customer-impact signals and technical findings into decision-ready recommendations for owning Product, engineering, QA, platform, and operations teams.

Release Management & Remote Feature Rollout | Comcast | Jul 2018–2025

- Helped establish stage-based deployment and readiness practices for remotely configurable features; standardized change inputs, metrics pages, rollout tracking, and escalation paths with Product, engineering, QA, analytics, and operations.

EARLIER TECHNICAL EXPERIENCE

QA & Software Development | Comcast | Oct 2014–Jun 2018

- Worked across quality assurance, software development, release validation, defect analysis, and production-readiness workflows across connected-device environments; built a foundation in reproducibility, test evidence, dependencies, and release risk.

Software Engineer | Massachusetts Health Exchange Platform | Mar 2014–Oct 2014

- Worked with SOA, OSB, and OPA technologies on a high-availability public-sector platform supporting regulated healthcare workflows; built early experience in enterprise integration, software reliability, and operational discipline.

INDEPENDENT AI PRODUCT PORTFOLIO

- Deterministic Code Authorization — policy-controlled authorization for AI-assisted code changes before merge, release, or deployment.
- Decision Replay & Ledger-Verified Execution — reconstructable evidence paths for AI-influenced decisions.
- Cross-Modal Provenance — lineage and verification concepts across AI-generated and AI-transformed text, code, documents, and images.
- Runtime Execution Control — product-control concepts for approving, blocking, escalating, or constraining consequential AI-driven actions.

IP, RESEARCH & EXTERNAL VALIDATION

- U.S. Patent 12,670,085 B1 — "Deterministic Offline Code Remediation with Ledger-Verified Replay and Template-Based Patch Generation," issued June 30, 2026; applicant and inventor: Apurv Gaurav.
- Additional U.S. utility applications span deterministic code authorization, decision replay, cross-modal provenance, and runtime execution controls; pending claims and prosecution details remain outside public résumé scope.
- GAISS 2026 — accepted paper: "Deterministic Replay Verification for AI-Generated Code Remediation: An Offline Authorization Framework"; presentation scheduled October 2026.
- GAISS 2026 — Industrial Advisory Committee member; reviewer/service details will be updated only after official conference documentation.
- IEEE Access — PCS-R reliability-aware image-text provenance verification and human-review routing manuscript under review.
- Journal of Systems Architecture — "A Causal-Integrity Control Plane for Heterogeneous Inference Runtimes" — manuscript under review.

SELECTED PRODUCT ARTIFACTS

Strategy & requirements —

PRDs, problem statements, buyer/user framing, integration hypotheses, roadmap tradeoffs, and executive narratives.

Evaluation & rollout —

Success metrics, evaluation plans, evaluator workflows, rollout risks, readiness evidence, monitoring, and decision gates.

Commercial & diligence —

Public implementation, technical evidence, executive briefs, integration framing, and acquisition/licensing materials.

SELECTED WRITING & PRODUCT COMMUNICATION

Deployment Safety: Signals vs Evidence · AI PMs Should Start With the Workflow, Not the Model · Runtime Governance for Enterprise AI · Human Review vs Automation in AI Systems · Traceability and Replay in AI Systems

Public channels: LinkedIn articles · YouTube technical briefings · AI Product Thinking podcast

PRODUCT & TECHNICAL FLUENCY

LLM and agentic systems · distributed systems · evaluation and safety measurement · observability and telemetry · policy controls · human-in-the-loop workflows · Python · SQL · GitHub · NIST AI RMF · OWASP LLM Top 10

EDUCATION

Penn State World Campus — Doctor of Engineering, Artificial Intelligence | Aug 2025–Dec 2027 (expected)

University of Pennsylvania — M.S., Information Systems & Spatial Analysis | 2009–2011

Acropolis Institute of Technology and Research — B.Tech, Information Technology | 2005–2009