AI Product Lifecycle

AI & Governance | Original applied framework | Portfolio version 1.0

A stage-and-gate lifecycle for moving an AI opportunity from workflow discovery to controlled production monitoring.

Problem it solves

AI initiatives frequently begin with model capability rather than a validated workflow problem, creating weak adoption, unclear accountability, and production risk.

When to use it

- A team is evaluating an AI-enabled workflow.
- Model output influences user decisions or consequential actions.
- Data readiness and evaluation criteria are uncertain.
- The product requires ongoing quality, latency, cost, and risk monitoring.

When not to use it

- A deterministic rule solves the problem more reliably.
- There is no meaningful workflow problem or user need.
- Data cannot be lawfully or responsibly used.
- The organization cannot support monitoring or human oversight.

Inputs

- Customer workflow and problem evidence
- Alternative non-AI solutions
- Data sources, ownership, and quality
- Risk and impact assessment
- Technical feasibility
- Product economics

Process

1. Opportunity - Define the workflow problem, value hypothesis, affected users, and alternatives.

2. AI suitability - Determine whether AI is appropriate compared with deterministic or manual approaches.

3. Data readiness - Assess availability, quality, privacy, representativeness, and access.

4. Prototype - Build the smallest testable capability and workflow integration.

5. Evaluation - Measure task quality, failure modes, fairness, robustness, latency, and cost.

6. Experience design - Design confidence, evidence, review, correction, escalation, and feedback.

7. Controlled release - Launch to a bounded audience with clear safeguards and support readiness.

8. Monitoring - Track quality, workflow success, adoption, drift, incidents, economics, and feedback.

9. Improve or retire - Scale, constrain, redesign, or retire based on production evidence.

Decision points

- Is AI materially better than available alternatives?
- Is the data fit for the intended context?
- Are evaluation thresholds tied to workflow consequences?
- Is human oversight appropriate and usable?
- Are production reliability, cost, and monitoring acceptable?

Outputs

- AI opportunity assessment
- Data-readiness decision
- Evaluation plan
- Human-oversight design
- Controlled-release plan
- Production monitoring dashboard

Success measures

- Workflow completion
- User acceptance and correction
- Escalation rate
- Quality by segment
- Latency
- Cost per workflow
- Incident rate
- Adoption and repeat use

Portfolio application

I apply this lifecycle when assessing enterprise AI opportunities, defining human-reviewed workflows, and translating quality, privacy, latency, cost, and monitoring needs into product requirements.

Common pitfalls

- Starting with a model instead of a workflow.
- Using one accuracy score as the definition of success.

- Leaving human review undefined.
- Ignoring production economics.
- Treating release as the end of evaluation.

Skills demonstrated

- AI product strategy
- Data readiness
- Evaluation design
- Human-in-the-loop
- Product economics
- Monitoring

Interview discussion prompts

- How do you decide whether AI is appropriate?
- How do you turn model evaluation into product requirements?
- Where should human review occur?
- What evidence would cause you to retire an AI capability?

This is a portfolio framework and is not legal, regulatory, compliance, financial, or medical advice.