

Planning a Data Science Initiative: From Problem to Value



Planning discipline drives success more than modeling sophistication



Most failures are leadership-, data-, and process-driven, not algorithmic




Over 80% of AI initiatives fail;
~85% never leave experimentation




Course lifecycle: problem framing through governance to value realization



Diagnostic: which stage is weakest in your current initiative?



Course Map, Roles, and Working Definitions



Initiative: value-seeking effort, problem selection to value realization



Project: bounded delivery with defined scope and end date



Program: coordinated portfolio with shared data and platform assets



POC, pilot, production: distinct evidence requirements before funding continues



Agree decision rights: business question, scope changes, halt authority, post-launch owner



Running artifact: one-page initiative charter draft, refined each slide



Framing the Business Problem Before the Model

- Translate the sponsor's request into a decision: who, how often, and what they do now
- Separate business objective, analytical objective, and deliverable — “we need a model” is not an outcome
- Define success metric, guardrails, and baseline before any modeling starts
- Objectives are clearly defined only sometimes or rarely in most projects (Rexer Analytics, 2023)
- Teams that agreed and documented metric definitions before building shipped 81% of the time; those that skipped shipped 13%
- Deliver a one-page problem statement the sponsor can approve




Feasibility and Value Assessment

- Data feasibility: availability, access rights, history depth, labels, inference-time features
- Technical feasibility is one axis; organizational, legal, privacy, and vendor constraints can block
- Enterprise audits repeatedly find data-layer failures: undefined metrics, identity gaps, lineage gaps
- Value range against cost range, not a single point, so sponsors decide under uncertainty
- Time-boxed feasibility experiments prove one claim, with a pre-agreed kill criterion
- Scope a 30-day data diagnostic before committing to full remediation



Scoping into Value-Bearing Increments

- Convert approved problem statement into in-scope and out-of-scope work
- Slice scope into increments that each deliver a usable output
- Prioritize with sponsors: value, effort, risk, data readiness, dependencies
- Agree phased outcomes: what each phase boundary will and won't teach
- Name boundaries early: geography, population, channel, decision type, latency
- Document deferrals, not deletions, so the roadmap stays honest



Data Foundation, Governance, and Compliance Constraints



Public availability is not a lawful basis; assess expected privacy



Quality test: completeness, accuracy, timeliness, representativeness, fitness for purpose



Govern design, collection origin, labeling, assumptions, and bias mitigation (AI Act Art. 10)



Check early: high-risk AI duties, ADMT notice and opt-out, privacy impact assessments




Artifacts: data map, lineage, DPIA, model and data cards, known limitations



Data readiness is the top blocker: fund metric definitions, identity, lineage, feature governance

Data map



Lineage




DPIA



Model and data cards



Team, Roles, and Operating Model



Core roles: DS lead, analysts, engineers, domain expert, product owner, sponsor, post-launch owner



Insufficient skills is the most frequently cited barrier to data science project success



Model hiring and onboarding lead times; staffing is a schedule risk, not a formality



Build versus buy versus partner: weigh control, speed, cost, regulatory visibility, maintainability



Set engagement cadence with decision checkpoints and one escalation path



Name who owns AI risk management and who is accountable for outputs

Milestones, Estimating Under Uncertainty, and Risk Controls



Risk register
(categories)

Technical

Data

Organizational

Legal

Regulatory



Milestones mirror the lifecycle: framing, data readiness, modeling, validation, deployment, monitoring



Estimate in ranges with stated assumptions; define who can re-baseline



Risk register spans technical, data, organizational, legal, and regulatory categories, each owned



Agree stage gates and kill criteria with the sponsor before work starts




Lack of proactive planning is a top deployment impediment; plan rather than react



Build change management and post-launch ownership into the plan from day one

Evaluation, Deployment, and MLOps Readiness

FROM MODELS
TO LASTING IMPACT



Offline evaluation: validation strategy, acceptance thresholds, data slices, leakage checks



Kill training-serving skew: schema and feature logic must mimic training (golden rule)



Deployment modes: batch, real-time, embedded, human-in-the-loop; each needs its own rollback



Rollout: canary, shadow, blue/green, or A/B promotion with documented rollback path



Monitor data quality, drift, bias drift, latency, throughput; many teams run no tooling



Delayed ground truth: use proxy metrics and thresholds that trigger retrain or rollback



Version, lineage, model cards, and decision logs planned from the start



Risk register / control ledger

Make risks visible, define mitigations, and track decisions across the lifecycle.



Cost, Benefit, and the Business Case



Total cost: people, data acquisition and remediation, infrastructure, tooling, integration, change effort



Hidden costs: metric definitions, labeling, pipeline hardening, monitoring, ongoing retraining



Benefits differ by evidence standard: revenue, cost avoidance, risk reduction, capability building



Format finance expects: ranges, assumptions, sensitivity, payback window, what is not counted



Scope the ask to one use case: \$180K–\$420K in 8–14 weeks, not full data-layer remediation



Define attribution method and measurement window before launch

Total Cost



People



Data acquisition and remediation



Infrastructure



Tooling



Integration



Change effort

Benefits



Revenue



Cost avoidance



Risk reduction



Capability building

Governance, Risk Materiality, and Approval Package



Apply NIST AI RMF functions: govern, map, measure, manage—governance is cross-cutting



Risk materiality drives control intensity: higher-risk use cases get stricter scrutiny



Review gates: lawful basis, DPIA, model risk, fairness review, human oversight



Approval package: problem statement, feasibility, scope, plan, cost-benefit ranges



Anticipate sponsor objections; pre-negotiate scope, timeline, quality, control trade-offs



Build evidence once; reuse across AI Act, GDPR, and state frameworks

Risk Materiality



Control Intensity



Lawful basis



DPIA



Model risk



Fairness review



Human oversight

Approval Package



- Problem statement
- Feasibility
- Scope
- Plan
- Cost-benefit ranges

Value Realization and Post-Launch Ownership



Handoff design: name who owns the model, the metric, and alert response



Operational cadence: monitoring reviews, retraining triggers, incident response, periodic risk reassessment



Track value against original success and guardrail metrics, including segment-level checks



Budget retraining, monitoring, and support as recurring sustainment costs, not one-time



Fold data fixes, metric definitions, and docs into shared assets



Close the loop with the sponsor: report realized value and where estimates missed



Common Failure Modes and Planning Anti-Patterns

Learn to spot common pitfalls early — and plan for success.



Anti-pattern:
tech-led launch,
a solution looking
for a problem



Anti-pattern:
run as IT project,
not research-and-
iteration



Anti-pattern:
no upfront data
audit; fixes surface
after modeling



Anti-pattern:
no post-launch
owner or structured
monitoring



Anti-pattern:
unowned data
and unsettled metric
definitions



Countermeasures:
clarify decision
and metrics, fund
data readiness,
name owners, set
kill criteria

Initiative Charter



Workshop: Building the One-Page Initiative Plan

-  Assemble the artifact: problem statement, metrics, baseline, scope, feasibility, data readiness
-  Add plan, milestones, risk register, cost and **benefit** ranges, governance decisions, named owners
-  Peer review in mixed role groups; each role pressure-tests the others' sections
-  Red-team: name your three likely failure modes and each one's earliest warning signal
-  Close: define the next three decisions with dates, decision-makers, and required evidence
-  Two-minute sponsor pitch: lead with the decision and value range, not the technique



Red-team

Name your three likely failure modes and each one's earliest warning signal.



Next Three Decisions

Define the next three decisions with dates, decision-makers, and required evidence.

PersonWise.

PersonWise • AI digital-human course platform



Scan the code or click anywhere to watch this course live, with voice Q&A

personwise.ai/t/97a2cd17/public