FELENTRA / FOUNDATIONAL WHITE PAPER

The Decision Architecture Framework

Designing the journey from uncertainty to responsible action

Felix Olusola Abayomi

Felentra · First edition · August 2026

“Every important decision deserves thoughtful architecture.”

PUBLICATION NOTE

How to use and cite this paper

A field-defining proposal should be judged by whether it is clear enough to test, useful enough to apply, and open enough to improve.

Recommended citation

Abayomi, F. O. (2026). The Decision Architecture Framework: Designing the Journey from Uncertainty to Responsible Action. Felentra. First edition.

Status and scope

This paper proposes Decision Architecture as an emerging integrative discipline. It synthesizes established work in decision science, behavioral economics, human factors, systems thinking, design, governance, and responsible artificial intelligence. The seven-stage Felentra Framework is an original organizing model; it is offered as a practical and researchable proposition, not as a claim that prior disciplines are obsolete.

The three cases in this paper are Felentra product and research applications. They illustrate how the framework can be used; they are not presented as independent evaluations, completed client engagements, or evidence of measured impact. Claims should therefore be cited with that distinction intact.

Rights and contact

ABSTRACT

Abstract

Organizations have unprecedented access to data, analytics, and artificial intelligence, yet consequential decisions often remain slow, inconsistent, opaque, or weakly connected to action. This paper argues that the missing layer is architectural: the deliberate design of how a decision is framed, informed, contested, made, enacted, and improved. It defines Decision Architecture as the intentional design of the information, context, choices, workflows, safeguards, and guidance that shape the journey from uncertainty to responsible action.

The paper introduces the Felentra Decision Architecture Framework: Purpose ☐ Context ☐ Uncertainty ☐ Options ☐ Judgment ☐ Action ☐ Learning. Each stage addresses a distinct failure mode. Purpose prevents optimization of the wrong outcome. Context keeps recommendations situated. Uncertainty separates evidence from assumption. Options make the choice set visible. Judgment makes values and trade-offs explicit. Action connects conclusions to executable next steps. Learning closes the loop through outcomes and unintended effects.

Decision Architecture is positioned as a bridge discipline. It does not replace decision science, user-experience design, organizational governance, professional expertise, or AI assurance. It connects them around the unit that ultimately matters: a consequential decision in context. The paper traces the intellectual lineage of the field, explains why generative and predictive AI increase the need for explicit decision design, describes methods and artifacts for practice, applies the framework to three illustrative cases, and proposes principles, ethical obligations, evaluation criteria, and future research directions.

CORE THESIS

Better information does not automatically produce better decisions. Better decisions require deliberately designed conditions for understanding, judgment, action, accountability, and learning.

Keywords: decision architecture; decision support; human judgment; responsible AI; uncertainty; choice architecture; systems design; human oversight; organizational learning.

CONTENTS

Contents

EXECUTIVE SUMMARY

The decision is the unit of design

Every policy, product, workflow, and intelligent system contains an implied theory of how people should move from what is known to what should be done.

The problem

Modern organizations are rich in signals and poor in synthesis. Dashboards reveal more, models predict more, and generative systems can produce plausible answers on demand. Yet the user still confronts a practical question: What should we do next? The gap between an answer and a responsible action is filled by context, uncertainty, values, authority, timing, and consequences. When those elements remain implicit, the system has an architecture—but it is accidental.

Accidental architectures hide assumptions in defaults, interfaces, escalation rules, data pipelines, incentives, and organizational habits. They often optimize what is easy to measure rather than what matters. They may provide an output without a rationale, a recommendation without a responsible owner, or an action without feedback. Decision Architecture makes these structures visible and redesignable.

The proposition

A decision should be designed as an end-to-end journey. The quality of a recommendation is only one component. A high-quality architecture also defines the purpose, represents relevant context, exposes uncertainty, constructs legitimate options, supports accountable judgment, enables proportionate action, and learns from outcomes.

NORTH-STAR QUESTION	Given what we know, what remains uncertain, and what matters most—what is the most responsible next action?

EXECUTIVE SUMMARY

What the framework contributes

A common language

The framework gives cross-functional teams a shared sequence for examining a decision. Researchers can ask what evidence informs each stage. Designers can identify the information and interaction needed. Domain experts can define legitimate options and thresholds. Engineers can implement provenance, confidence, escalation, and feedback. Leaders can assign authority and accountability.

A practical discipline

Decision Architecture works at three levels. At the decision level, it improves a particular choice. At the system level, it aligns data, interfaces, policies, incentives, and people around that choice. At the institutional level, it creates repeatable standards for governing portfolios of decisions. The method begins with one consequential decision because broad transformation becomes tractable when anchored in an observable unit of work.

A research agenda

The framework creates testable questions: Does explicit uncertainty calibration improve reliance? Do richer option sets reduce premature closure? Does recording rationale improve organizational learning? Which forms of explanation support judgment without increasing automation bias? How should decision quality be measured when outcomes are delayed or partly outside the decision maker's control?

BOUNDARY	Decision Architecture supports human judgment; it does not silently transfer moral or legal responsibility to a model, interface, or workflow.

1 / WHY DECISION ARCHITECTURE?

Information is not action

The central problem is no longer access alone. It is the conversion of abundant, uneven, and sometimes conflicting information into justified action.

The last-mile failure

Many systems are built to retrieve, display, rank, or predict. They succeed on those terms and still fail the decision maker. A risk score may not show which intervention is feasible. A dashboard may not distinguish signal from noise. A generated answer may not know who has authority to act. The missing work occurs after information appears and before consequences unfold.

Architecture already exists

Every decision environment has defaults, categories, timing, permissions, narratives, and paths of least resistance. These structures direct attention and shape behavior whether or not anyone designed them deliberately. Decision Architecture begins by treating those structures as consequential design material.

The design object

The object is not an isolated click or recommendation. It is the complete journey through which a person or group understands the situation, develops a choice set, exercises judgment, commits to action, and learns. This broader object prevents local interface improvements from masking systemic failure.

1 / WHY DECISION ARCHITECTURE?

The decision gap

The distance between “what the system says” and “what we should do” is where responsibility lives.

Four gaps

An epistemic gap separates available information from warranted belief. A context gap separates general evidence from this situation. A normative gap separates predicted outcomes from what ought to matter. An execution gap separates a choice from coordinated action. No model output closes all four by itself.

A bridge discipline

Decision Architecture connects analytical, behavioral, technical, organizational, and ethical work around one operational question. It borrows without collapsing distinctions. Decision analysis helps compare alternatives; human factors examines interaction with automation; service and product design shape experiences; governance assigns rights and duties. Architecture integrates these contributions into a coherent decision pathway.

Why now

As intelligent systems become more capable and more embedded, organizations can automate fragments faster than they can understand the whole. The risk is not merely a wrong prediction. It is a system that makes the wrong question easy, the responsible option invisible, or accountability impossible to locate.

1 / WHY DECISION ARCHITECTURE?

What counts as a decision?

A decision is a commitment—explicit or implicit—that changes what actions remain available.

Decision versus output

A forecast estimates what may happen. A classification assigns a category. A recommendation proposes a course. A decision authorizes or withholds action under a set of responsibilities. Conflating these objects encourages organizations to treat analytical performance as decision quality.

Consequentiality

Architecture is most valuable when a decision materially affects rights, safety, resources, opportunity, trust, or future options; when uncertainty is meaningful; when multiple actors must coordinate; or when effects accumulate across repeated choices. The appropriate rigor should scale with impact and reversibility.

Hidden decisions

Many consequential decisions are embedded in routine operations: which cases are reviewed, what evidence is requested, when a person can appeal, which alternatives are shown, or when a recommendation becomes the default. Mapping hidden decisions often reveals more leverage than optimizing the visible final approval.

1 / WHY DECISION ARCHITECTURE?

Decision quality is multidimensional

A good outcome can follow a poor process, and a poor outcome can follow a defensible process under genuine uncertainty.

Beyond outcome accuracy

Evaluation should distinguish outcome quality, process quality, epistemic quality, ethical quality, and learning quality. Outcome quality concerns what happened. Process quality concerns whether relevant evidence and perspectives were considered. Epistemic quality concerns calibration and provenance. Ethical quality concerns rights, fairness, agency, and accountability. Learning quality concerns whether the system improves.

Avoiding hindsight bias

If evaluation rewards only realized outcomes, decision makers may be punished for well-reasoned choices and rewarded for luck. A decision record should capture what was known, what was uncertain, the alternatives considered, the rationale, the responsible owner, and the review point.

A practical test

An architecture is stronger when a reasonable reviewer can reconstruct why the decision was made, determine whether the process was proportionate to its stakes, observe whether commitments became action, and see how outcomes changed the next cycle.

1 / WHY DECISION ARCHITECTURE?

From feature thinking to decision thinking

Products become more useful when teams stop asking only what functionality to add and ask which decision deserves support.

A different brief

Feature briefs specify screens and capabilities. Decision briefs specify the consequential choice, the people affected, the purpose, the constraints, the unknowns, the legitimate options, the authority to decide, and the evidence of success. Technology follows from that analysis.

A sharper metric

Usage is not equivalent to value. More recommendations, longer sessions, or faster throughput can coexist with worse judgment. Decision-centered measures include time to warranted action, comprehension, calibration, option quality, reversibility, contestability, follow-through, and learning.

Begin with one decision

The most useful entry point is the decision that, if consistently improved, would create the greatest positive impact. Starting with one decision narrows the system boundary without pretending the surrounding organization does not matter.

2 / THE HISTORY OF DECISIONS

From philosophy to probability

Decision Architecture inherits a long argument about reason, evidence, value, and action.

Ancient and practical reason

Classical traditions distinguished theoretical knowledge from practical wisdom: knowing general truths is different from judging what to do in a particular situation. This distinction remains foundational. A system may possess analytical power while lacking the situated judgment required for responsible action.

Formal choice

The development of probability and expected utility made uncertainty calculable. Bernoulli's account of utility, later axiomatized by von Neumann and Morgenstern, provided a normative language for comparing risky prospects. Bayesian reasoning offered a coherent way to update belief with evidence.

The contribution and limit

Formal models made assumptions explicit and enabled disciplined comparison. But their usefulness depends on how outcomes, probabilities, preferences, and constraints are represented. Architecture attends to the work before and after calculation: framing the problem, constructing options, assigning authority, implementing action, and learning.

2 / THE HISTORY OF DECISIONS

Bounded rationality and organizations

Real decisions are made with limited attention, incomplete knowledge, conflicting goals, and finite time.

Satisficing

Herbert Simon challenged the image of exhaustive optimization. Under bounded rationality, people search selectively and often accept an option that meets an aspiration level. The implication for architecture is not that people are deficient; it is that environments must be designed around realistic cognitive and organizational limits.

Organizations as decision systems

Rules, routines, reporting lines, incentives, and information channels determine which problems receive attention and which alternatives can be considered. A decision is therefore rarely the work of a single mind. It is produced by a distributed system whose structure affects the result.

Design implication

Good architecture manages attention, establishes stopping rules, makes dependencies visible, and creates escalation paths. It does not demand impossible comprehensiveness. It provides enough structure to support a defensible judgment at the moment action is required.

2 / THE HISTORY OF DECISIONS

Heuristics, biases, and ecological fit

Human judgment uses shortcuts that can be powerful, predictable, and context-sensitive.

Behavioral findings

Research on heuristics and biases showed systematic departures from simple rational-choice models: framing, anchoring, availability, loss aversion, overconfidence, and other effects can shape judgment. These findings made the design of choice environments an explicit concern.

Two cautions

A bias catalogue is not a complete theory of decision quality, and "debiasing" can become paternalistic if designers assume their own objective is neutral. Heuristics can also be adaptive in environments with limited time and reliable cues. The relevant question is which strategy fits which environment.

Architectural response

Make important frames visible, invite alternatives, calibrate confidence, separate reversible from irreversible choices, and allow contestation. The goal is not to eliminate human intuition. It is to create conditions in which intuition, analysis, and expertise can check one another.

2 / THE HISTORY OF DECISIONS

Choice architecture and behavior

The arrangement of options influences what people notice and choose.

Defaults and presentation

Choice architecture demonstrated that defaults, ordering, simplification, salience, and feedback can alter behavior without removing options. The insight is durable: there is no neutral presentation. Even an unordered list has an order.

Decision Architecture is broader

Choice architecture often focuses on the immediate environment of selection. Decision Architecture includes that environment but expands the object to purpose, evidence, uncertainty, option formation, accountability, execution, and learning. It asks not only how a choice is presented, but whether the right decision has been framed and whether consequences improve the system.

Ethical implication

Influence should be disclosed, proportionate, and aligned with legitimate purposes. Meaningful choice requires more than the formal presence of alternatives; it requires intelligibility, feasibility, and a genuine path to decline or appeal.

2 / THE HISTORY OF DECISIONS

Decision support, human factors, and design

Computing shifted the question from how people decide to how people and systems decide together.

Decision support systems

Early decision support combined data and models with managerial judgment rather than imagining that computation would replace management. Human-computer interaction and information visualization later emphasized usability, mental models, feedback, and the costs of cognitive overload.

Automation and reliance

Human-factors research showed that automation can be used, misused, disused, or abused. Reliability alone does not determine appropriate reliance; the allocation of functions, transparency, workload, expertise, and organizational incentives also matter.

Design thinking

Human-centered and participatory design brought observation, reframing, prototyping, and iteration into complex problem solving. Decision Architecture shares this empirical, iterative stance but anchors it in a specific object: the conditions through which consequential judgments become accountable action.

3 / WHY AI CHANGES EVERYTHING

From tools to participants

AI increasingly generates the language, rankings, predictions, and options through which a situation is understood.

A qualitative shift

Traditional software usually applies rules specified in advance. Contemporary machine-learning and generative systems infer patterns and produce probabilistic outputs that can be fluent, adaptive, and difficult to inspect. Their outputs may enter the decision journey before users recognize that a decision is being shaped.

Scale and speed

AI can compress research, synthesis, and production cycles. This expands access and creates extraordinary leverage. It also scales errors, embeds recommendations across many contexts, and shortens the time available for reflection. Architecture must therefore govern not only model behavior but the pathway from output to consequence.

The central distinction

AI can contribute intelligence: patterns, predictions, summaries, simulations, and candidate options. It does not thereby inherit legitimate authority to determine what matters, accept moral responsibility, or represent all affected interests.

3 / WHY AI CHANGES EVERYTHING

Fluency is not warrant

The form of an answer can exceed the strength of its evidence.

Epistemic risk

Generative systems can produce coherent explanations for uncertain, incomplete, or false claims. Predictive systems can output precise scores that conceal dataset shift, measurement error, subgroup variation, or contested labels. Interfaces often intensify this risk by presenting a single result without provenance or alternatives.

Calibration by design

A responsible architecture distinguishes observed facts, user-provided information, model inference, policy assumptions, and unresolved unknowns. Confidence should be meaningful to the task, not ornamental. When uncertainty cannot be quantified reliably, the system should communicate its source, direction, and decision consequence.

Evidence posture

Users need to know what would change the recommendation, where the evidence came from, and when expert review is required. Explanations should support a decision—not merely rationalize an output after the fact.

3 / WHY AI CHANGES EVERYTHING

Automation bias and responsibility gaps

A human in the loop is not automatically a human in control.

Nominal oversight

Oversight fails when reviewers lack time, information, authority, or a realistic ability to disagree. Requiring a click from a human may transfer blame without transferring meaningful control. High-volume workflows can turn review into ritual confirmation.

Appropriate reliance

Systems should help users decide when to rely, verify, defer, escalate, or reject. That requires clear task boundaries, known limitations, uncertainty cues, accessible evidence, and institutional protection for disagreement. Expertise should shape both the model's role and the interface around it.

Responsibility must remain locatable

Every consequential action needs a named accountable role. Responsibility cannot be assigned to "the algorithm." Suppliers, deployers, organizational leaders, professional users, and operators may hold different duties, but those duties must be explicit.

3 / WHY AI CHANGES EVERYTHING

AI expands the option space

Generative capability changes not only evaluation but the creation of alternatives.

Opportunity

AI can surface overlooked scenarios, generate counterarguments, simulate consequences, translate across expertise, and reduce the cost of exploring alternatives. This can counter premature closure and make high-quality analytical support more widely available.

New failure modes

Generated options can be infeasible, unsafe, homogenized, or subtly constrained by the prompt and training distribution. An abundance of possibilities can create a new overload problem. Ranking may prematurely collapse a plural choice set into a single apparent answer.

Architectural requirement

Option generation and option authorization should be separated. Candidate alternatives should be tested against constraints, stakeholder rights, reversibility, and evidence. The system should make omitted categories visible and preserve the ability to introduce a human-generated option.

3 / WHY AI CHANGES EVERYTHING

Governance must meet the workflow

Principles become real only when translated into operating decisions.

From policy to controls

Risk frameworks emphasize governance, mapping context, measurement, and management. These activities must appear in product requirements, data practices, access controls, interface states, escalation paths, incident response, and post-deployment monitoring.

Risk is contextual

The same model can create different risks in entertainment, conservation, finance, employment, or health. Architecture begins with the decision and its stakes, not a generic capability rating. It asks who can be harmed, what rights are implicated, how reversible the action is, and what recourse exists.

Lifecycle discipline

Responsible deployment is not a one-time model review. Contexts drift, users adapt, data changes, and organizations discover unintended effects. Learning must include monitoring, complaints, overrides, near misses, and qualitative evidence—not only aggregate performance.

3 / WHY AI CHANGES EVERYTHING

The new design brief

Design the collaboration among people, models, evidence, rules, and institutions.

Role clarity

Specify what the AI may generate or recommend, what it must not decide, when it must ask for more information, and when it must hand control to a qualified person. Tie each boundary to the stakes and reversibility of the decision.

Interaction quality

Show sources where practicable, distinguish inference from fact, offer relevant alternatives, expose material uncertainty, and provide a visible route to challenge. Avoid persuasive anthropomorphism when it could inflate trust or obscure institutional responsibility.

Organizational readiness

A well-designed interface cannot compensate for missing policy, inadequate staffing, misaligned incentives, or absent accountability. Decision Architecture treats the technical and institutional layers as one system because users experience their combined effect.

4 / THE FELENTRA FRAMEWORK

The seven-stage architecture

The sequence creates discipline; the context determines the work.

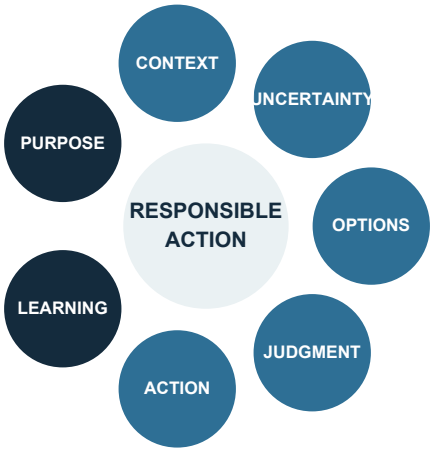

Purpose

Define the decision and the outcome it should serve.

Context

Represent the people, environment, constraints, objectives, evidence, and dependencies.

Uncertainty

Identify what is unknown, changing, inferred, or disputed.

Options

Make legitimate choices and meaningful alternatives visible.

Judgment

Evaluate trade-offs, consequences, rights, and values under clear authority.

Action

Convert the decision into a feasible, proportionate, and accountable next step.

Learning

Use outcomes, overrides, feedback, and unintended effects to improve the next decision.

4 / THE FELENTRA FRAMEWORK

1. Purpose

A system can be efficient and still optimize the wrong thing.

Define the decision

Write the decision as a verb with an owner and a time horizon: allocate, approve, intervene, defer, report, escalate, or stop. Separate the decision from the analytical questions that inform it.

Specify value

Name the intended outcome, affected parties, constraints, and unacceptable harms. Distinguish proxy metrics from the purpose they approximate. If multiple purposes conflict, do not hide the conflict inside a composite score.

Tests

Can the team state the decision in one sentence? Is the accountable owner clear? Are success and non-negotiable boundaries defined? Would affected people recognize the stated purpose as legitimate?

4 / THE FELENTRA FRAMEWORK

2. Context

Context changes the answer because the same evidence can justify different actions in different situations.

Map the environment

Identify decision makers, subjects, beneficiaries, reviewers, and implementers. Map legal duties, policy constraints, resources, timing, cultural conditions, prior commitments, and downstream dependencies.

Minimum sufficient context

Collect what materially changes the choice; avoid collecting data merely because it is available. Context has privacy and burden costs. The architecture should explain why sensitive fields are needed and what happens when they are absent.

Tests

Which variables change the recommended action? Whose perspective is missing? What local knowledge is required? Where might general evidence fail to transfer? What context is inferred rather than observed?

4 / THE FELENTRA FRAMEWORK

3. Uncertainty

Uncertainty is not a defect to hide; it is information for judgment.

Create an uncertainty register

Record unknown facts, model uncertainty, ambiguity about values, volatility in the environment, and disagreement among experts. For each item, note its source, importance, reducibility, and the action it should trigger.

Use decision-relevant expression

A probability is useful only if users understand what it refers to and how it should affect action. Combine calibrated quantities with plain-language consequence statements, scenarios, sensitivity analysis, and evidence thresholds.

Tests

What would reverse the recommendation? Which assumptions carry the result? What is unknowable before action? Is uncertainty communicated consistently, including when it weakens the preferred option?

4 / THE FELENTRA FRAMEWORK

4. Options

The quality of a decision is constrained by the quality of its choice set.

Construct before ranking

Include status quo, delay, staged action, reversible experiment, escalation, and hybrid alternatives when legitimate. Separate idea generation from evaluation so early preferences do not narrow the field prematurely.

Make feasibility visible

An option is not meaningful if the decision maker lacks resources, authority, comprehension, or a real opportunity to select it. Show dependencies, costs, time, reversibility, and affected parties.

Tests

Which category of alternative is missing? Who benefits from the current framing? Is there a safe-to-try option? Is “do nothing” explicit, along with its consequences? Can affected people propose or contest an option?

4 / THE FELENTRA FRAMEWORK

5. Judgment

Judgment converts evidence and values into a warranted choice.

Make trade-offs explicit

Name the criteria, their rationale, and any thresholds. Do not let a ranking formula silently decide contested values. Where criteria are incommensurable, preserve the deliberation rather than forcing false precision.

Combine modes of knowing

Analytical models, professional expertise, lived experience, and institutional rules each contribute. Architecture specifies when each is relevant and how disagreement is handled. Structured dissent and pre-mortems can reveal assumptions that consensus suppresses.

Tests

Who has authority and why? What evidence is decisive? Which value judgments are embedded? Can a reviewer understand the rationale? Is there a conflict-of-interest process and a route to challenge?

4 / THE FELENTRA FRAMEWORK

6. Action

A decision has not created value until it changes what happens next.

Close the execution gap

Specify the next action, owner, timing, dependencies, resources, permissions, communication, and confirmation. Match the action to confidence and stakes: uncertainty may justify a reversible pilot rather than paralysis or full commitment.

Design reversibility and escalation

Define stop conditions, review points, exception handling, and escalation paths. Make it easy to correct course. For high-impact actions, require independent checks or dual authorization where appropriate.

Tests

Can the responsible person act now? Does everyone know who does what by when? What happens if the action fails? Who is informed? Can the decision subject seek explanation, correction, or appeal?

4 / THE FELENTRA FRAMEWORK

7. Learning

Every outcome should improve the next decision without turning people into unprotected data sources.

Close the loop

Capture outcome measures, process signals, overrides, appeals, near misses, and unintended effects. Compare results with the expectations and assumptions recorded at the time of decision.

Learn at multiple levels

Update the individual case, the decision rule, the product or workflow, and—when evidence supports it—the governing policy. Distinguish local adaptation from generalizable learning.

Tests

When will the decision be reviewed? What evidence counts as success or harm? Are feedback channels accessible? Can the system learn from disagreement rather than treating override as error? Who authorizes changes produced by learning?

4 / THE FELENTRA FRAMEWORK

The framework is a loop, not a funnel

Outcomes change context, reveal uncertainty, and sometimes require purpose to be reconsidered.

Iteration

The seven stages are presented sequentially for clarity, but practice is recursive. New context can reframe purpose; options can reveal missing evidence; implementation can expose constraints; learning can invalidate the original model.

Traceability

Each stage should leave a proportionate record: decision statement, context map, uncertainty register, option set, rationale, action plan, and learning review. Together these artifacts create a decision trace without requiring maximal documentation for every low-stakes choice.

Scaling rigor

Use lighter architecture for reversible, low-impact decisions and deeper review for high-impact, irreversible, or rights-affecting decisions. The sequence stays stable while evidence, participation, assurance, and documentation scale with risk.

5 / FROM FRAMEWORK TO PRACTICE

The Decision Architecture Review

Begin by observing how one consequential decision actually happens.

Discover

Interview participants, examine policies and interfaces, follow cases, and identify the formal and informal decision path. Look for rework, delay, hidden judgment, missing context, duplicated checks, and points where responsibility becomes ambiguous.

Map

Create a current-state decision map showing actors, information, transitions, systems, approvals, and feedback. Overlay uncertainty, failure modes, rights, and incentives. The map is a hypothesis to test with the people who live the workflow.

Redesign

Define the target architecture, prototype critical interactions, test with representative users, and specify governance controls. Prioritize changes by impact, dependency, effort, and reversibility. Implementation should include measurement and review from the start.

5 / FROM FRAMEWORK TO PRACTICE

Core artifacts

Artifacts make assumptions discussable and decisions inspectable.

Decision brief

A one-page statement of purpose, owner, scope, affected parties, stakes, time horizon, and success criteria.

Context map

Actors, evidence, constraints, incentives, systems, and dependencies that materially shape the decision.

Uncertainty register

Unknowns and assumptions, their consequence, how they might be reduced, and what action they trigger.

Option canvas

Alternatives compared on relevant criteria, feasibility, reversibility, rights, and downstream effects.

Decision record

The selected option, rationale, dissent, evidence, responsible owner, commitments, and review date.

Learning review

Expected versus observed outcomes, unintended effects, overrides, appeals, and authorized changes.

5 / FROM FRAMEWORK TO PRACTICE

Measures that matter

Measurement should improve the architecture, not reward superficial activity.

Process measures

Time to warranted action, completion of required evidence, diversity of options, escalation appropriateness, override patterns, and implementation follow-through.

Human measures

Comprehension, confidence calibration, perceived agency, workload, trustworthiness, accessibility, and the ability to contest or seek help.

Outcome and learning measures

Goal attainment, harm distribution, reversibility cost, recurrence, model and policy drift, feedback latency, and the proportion of lessons that lead to authorized changes. Metrics should be interpreted together; speed without quality or satisfaction without fairness is incomplete.

5 / FROM FRAMEWORK TO PRACTICE

Team and governance

Decision Architecture is cross-functional work with explicit ownership.

Roles

A decision owner holds accountability for the consequential choice. A domain lead defines legitimate evidence and options. Design and research leads represent the lived workflow. Data and engineering leads implement reliable systems. Risk, legal, privacy, security, and ethics functions define safeguards. Affected stakeholders contribute experience and challenge.

Decision rights

Use a rights map to specify who proposes, advises, decides, executes, reviews, and can stop or appeal. Avoid committees in which everyone participates but no one owns the result.

Cadence

Review assumptions before launch, monitor early use, examine exceptions and harms, and schedule periodic reassessment. High-impact changes require an explicit approval path and a rollback plan.

6 / ILLUSTRATIVE CASE STUDIES

How to read the cases

These examples demonstrate use of the framework; they do not claim externally validated outcomes.

Analytical purpose

The cases test transfer across domains: public safety and conservation, personal finance, and consumer subscriptions. Their differences expose why context is essential while their common sequence shows the framework's portability.

Evidence status

Each case is derived from Felentra's stated product or research concepts. The discussion identifies design requirements and hypotheses for evaluation. Any future efficacy claim should report study design, population, baseline, outcome measures, limitations, and conflicts of interest.

A common question

In each domain, information is available but incomplete. The design challenge is to help a person select a responsible next step while preserving uncertainty, agency, and accountability.

6 / CASE 1

Human–wildlife conflict intelligence

The decision is not merely “What species is this?” but “How should I respond safely and responsibly?”

Architecture

Purpose: protect people while supporting conservation. Context: species, location, behavior, season, user vulnerability, local capacity, regulation, and privacy. Uncertainty: identification confidence, immediate risk, location accuracy, and incomplete reporting. Options: withdraw, observe, seek guidance, report, contact authorities, or withhold precise location. Judgment: balance safety, conservation, urgency, privacy, and law. Action: deliver a context-sensitive response and escalation route. Learning: use verified reports and outcomes to improve guidance.

Design implications

Identification should not be presented as certainty. High-risk cues should trigger conservative guidance and human escalation. Sensitive species locations require access controls and data minimization. Offline or low-connectivity use may matter. Guidance must reflect local authority and avoid encouraging unsafe observation.

Evaluation hypotheses

Measure correct safety action under scenarios, calibration of identification confidence, escalation appropriateness, reporting completion, privacy comprehension, response time, and unintended behaviors. Field validation should involve wildlife professionals and affected communities.

6 / CASE 2

Debt to Legacy

Financial information becomes useful when it supports a next action without manufacturing certainty.

Architecture

Purpose: improve stability and long-term resilience. Context: income, debt, essential expenses, savings, goals, dependents, benefits, obligations, and time horizon. Uncertainty: income volatility, emergencies, rates, and shifting priorities. Options: pay debt, save, invest, defer, renegotiate, reduce spending, or increase income. Judgment: balance liquidity, cost, risk, wellbeing, and long-term goals. Action: offer a rationale and reversible next step. Learning: compare projected and actual progress.

Design implications

Recommendations must distinguish education from regulated advice and should not imply a universal ordering of goals. Sensitive financial data requires strong privacy boundaries. The system should explain assumptions, show trade-offs, permit user priorities, and provide hardship or professional-help pathways.

Evaluation hypotheses

Test whether users understand the rationale, choose feasible actions, maintain an emergency buffer, avoid harmful overpayment, and sustain progress. Segment results carefully; aggregate improvement can hide harm to people with volatile income or constrained choices.

6 / CASE 3

StreamPilot

The useful question is not “What can I watch?” but “Which subscriptions should I keep?”

Architecture

Purpose: reduce unnecessary spending without losing valued entertainment. Context: subscriptions, costs, watchlists, usage, preferences, household members, and billing cycles. Uncertainty: price changes, catalog movement, new releases, and future viewing. Options: keep, rotate, pause, cancel, downgrade, or bundle. Judgment: balance cost, access, convenience, and household preference. Action: present a strategy with rationale and reversible steps. Learning: compare later viewing and spending.

Design implications

The system should not optimize cancellations at the cost of user value. Recommendations should disclose affiliate incentives, respect shared-household preferences, and avoid dark patterns. Calendar-aware reminders and easy reversal can reduce the cost of experimentation.

Evaluation hypotheses

Measure net savings, valued-content access, regret, cancellation completion, re-subscription, recommendation acceptance by rationale, and household conflict. The key result is not engagement with the product but a better-aligned subscription portfolio.

6 / CASE SYNTHESIS

Domain-sensitive, not domain-bound

The sequence transfers; the content and safeguards do not.

Common structure

All three cases begin by reframing an informational task as a consequential decision. Each requires relevant context, explicit uncertainty, a real option set, value-sensitive judgment, an executable next step, and feedback.

Critical differences

Wildlife decisions may involve immediate safety and protected location data. Finance implicates vulnerability, regulation, and long time horizons. Entertainment decisions are lower stakes and highly reversible. The architecture should therefore vary in escalation, assurance, data collection, and documentation.

Lesson

Portability does not justify universality claims. A framework is useful when it prompts the right domain questions and makes boundaries explicit—not when it supplies the same answer everywhere.

7 / PRINCIPLES

Eight principles for practice

Principles guide judgment when procedures cannot anticipate every situation.

1. AI is not the decision

AI may surface patterns and possibilities; legitimate responsibility remains human and institutional.

2. Context changes the answer

Recommendations must be situated in people, environments, constraints, objectives, and consequences.

3. Trust must be designed

Transparency, privacy, consistency, competence, recourse, and accountability are structural.

4. Clarity is a form of intelligence

Reduce unnecessary complexity without hiding material nuance.

5. Guidance must lead toward action

Connect information to an appropriate, feasible next step.

6. Uncertainty must remain visible

Distinguish fact, assumption, estimate, inference, disagreement, and unknown.

7. Decisions should create learning

Use outcomes and unintended effects to improve the system.

8. Good decisions compound

Design not only one choice but the pattern of choices that follows.

7 / PRINCIPLES

Principles require tensions, not slogans

Responsible practice is visible in how competing goods are handled.

Clarity versus completeness

Simplification helps action, but omitted nuance can mislead. Use progressive disclosure: show the decision-relevant core and make evidence, assumptions, and detail accessible.

Personalization versus privacy

More context can improve relevance while increasing surveillance and exclusion risk. Collect the minimum sufficient context, prefer user control, and state how missing data affects guidance.

Speed versus deliberation

Urgency can justify streamlined paths, yet irreversible decisions need friction. Design deliberate pauses, second review, or staged commitment when stakes warrant them.

Consistency versus discretion

Standards can reduce arbitrary variation but may fail exceptional cases. Define legitimate discretion, document departures, and learn from patterns of override.

7 / PRINCIPLES

A practitioner’s commitment

The discipline earns trust through observable practice.

We will

Define purpose before selecting technology; make material uncertainty visible; preserve meaningful choice; identify accountable owners; protect sensitive information; design for access and inclusion; test with affected people; monitor outcomes and unintended effects; and revise the architecture when evidence demands it.

We will not

Recommend AI because it is fashionable; present generated confidence as evidence; hide value judgments in technical scores; use nominal human review as liability theater; optimize engagement against user welfare; or claim that a framework removes the need for domain expertise.

The standard

A Decision Architecture practitioner should be able to explain whom the system serves, which decision it shapes, what it assumes, how it can fail, who can challenge it, who is responsible, and how the next version will learn.

8 / ETHICS AND GOVERNANCE

Influence creates responsibility

Decision systems shape attention, perceived possibility, and behavior even when they do not formally decide.

Human dignity and agency

Treat people as participants with reasons and rights, not merely as targets for optimization. Preserve meaningful choice, informed refusal, and appropriate control over personal data.

Non-maleficence and proportionality

Match intervention, automation, and data collection to the stakes. Consider severity, likelihood, scale, reversibility, and distribution of harm. A small average benefit does not excuse concentrated injury.

Justice and inclusion

Ask who receives benefit, who bears error, who is absent from the data, and who can access recourse. Accessibility and language are not peripheral usability concerns; they determine who can exercise agency.

8 / ETHICS AND GOVERNANCE

Transparency, explanation, and contestability

A person affected by a decision needs more than a statement that a model was used.

Layered transparency

Disclose the system's role, purpose, material inputs, limitations, responsible organization, and routes to help. Provide deeper technical and governance information to auditors and qualified reviewers.

Decision-relevant explanation

Explain the basis of guidance, the factors that materially influenced it, the uncertainty that matters, and what could change the outcome. Avoid explanations that are plausible narratives without faithful connection to system behavior.

Contestability and recourse

People should be able to correct relevant data, introduce missing context, challenge an outcome, obtain human review where appropriate, and receive a timely response. Appeals are also learning channels: patterns can reveal architectural failure.

8 / ETHICS AND GOVERNANCE

Privacy, security, and data stewardship

Context can improve decisions, but indiscriminate context collection creates its own harms.

Purpose limitation

Collect and use data for explicit, legitimate decision purposes. Minimize scope, retention, and access. Separate data needed for immediate guidance from data proposed for learning or research.

Sensitive inference

AI may infer attributes that users did not disclose. Treat consequential inferred data with caution, test its necessity and validity, and provide controls appropriate to the risk.

Security as decision integrity

Compromised data, prompts, models, or workflows can change outcomes. Threat modeling should include manipulation of recommendations, unauthorized access, provenance loss, and attacks on feedback channels—not only confidentiality.

8 / ETHICS AND GOVERNANCE

Accountability and assurance

Governance should specify decisions, evidence, roles, and triggers—not merely publish principles.

Accountability chain

Identify who approves the purpose, accepts residual risk, validates domain performance, manages incidents, handles appeals, authorizes updates, and can suspend the system. Supplier contracts should not obscure deployment responsibility.

Assurance case

For consequential systems, maintain a structured argument linking claims of safety or trustworthiness to evidence and controls. Include known limitations, unresolved risks, monitoring, and the conditions under which the claim no longer holds.

Independent challenge

Use red teams, audits, affected-stakeholder review, and external evaluation proportionate to impact. Independence matters most when incentives favor launch, scale, or positive findings.

8 / ETHICS AND GOVERNANCE

An ethical gate for deployment

Before a system shapes consequential action, answer eight questions.

Gate

Is the purpose legitimate and specific? Is human authority meaningful? Are affected parties and rights mapped? Is material uncertainty visible? Are options genuine and non-coercive? Are privacy and security proportionate? Is there accessible contestation and remedy? Is monitoring tied to a named power to intervene?

Stop conditions

Do not deploy—or pause deployment—when the decision cannot be responsibly framed, evidence is inadequate for the stakes, oversight is ceremonial, affected people lack recourse, critical harms cannot be monitored, or no accountable party can act on findings.

Documentation

Record the decision to deploy with the same discipline as other consequential decisions: alternatives, rationale, dissent, controls, responsible owner, residual risk, and review date.

9 / FUTURE DIRECTIONS

A discipline to be built in public

Decision Architecture should develop through application, critique, measurement, and interdisciplinary exchange.

Theory

Clarify the relationship among decision architecture, choice architecture, decision intelligence, decision analysis, human-centered design, systems engineering, and governance. Develop a vocabulary that is precise enough for research without becoming inaccessible to practice.

Methods

Create validated instruments for decision mapping, uncertainty communication, option-set quality, meaningful oversight, contestability, and learning. Compare lightweight and intensive review methods across levels of risk.

Evidence

Run field studies that measure comprehension, calibration, action, fairness, and long-term outcomes. Publish negative results and boundary conditions. Build cross-domain case repositories with consistent reporting standards.

9 / FUTURE DIRECTIONS

AI-mediated institutions

The next frontier is not a better chatbot but a better institution that uses AI.

Collective decisions

AI can support agendas, synthesize public input, map disagreement, and explore scenarios. Architecture must prevent procedural scale from becoming democratic shallowness. Representation, legitimacy, and minority rights remain institutional questions.

Agentic systems

Systems that plan and act across tools compress the distance between suggestion and consequence. They require explicit authority boundaries, transaction limits, confirmation points, provenance, rollback, monitoring, and emergency stop mechanisms.

Adaptive architectures

Personalized and learning systems can alter the decision environment over time. Adaptation should be governed as an ongoing series of decisions: what may change, on whose evidence, within which bounds, with what notice, and under whose authority.

9 / FUTURE DIRECTIONS

Education and professional practice

A mature discipline needs shared competencies and accountable standards.

Curriculum

Practitioners need foundations in decision science, systems thinking, research, design, data and AI literacy, ethics, facilitation, governance, and domain inquiry. They also need the humility to recognize when specialist expertise is required.

Practice standards

Develop codes of conduct, case reporting templates, peer review, continuing education, and methods for disclosing conflicts and evidence status. Certification should follow a credible body of knowledge and demonstrated practice, not precede it.

Institutional home

The field will likely remain plural: universities, professional associations, public institutions, companies, and independent practices will contribute. Its strength should come from an open core of methods and evidence, with room for domain-specific standards.

9 / FUTURE DIRECTIONS

Research questions

The framework is valuable only if it generates better inquiry and survives serious testing.

Human judgment

Which uncertainty representations improve calibration across expertise levels? When does explanation increase appropriate reliance, and when does it rationalize automation? How should dissent be designed into group decisions?

Systems and organizations

How do incentives and decision rights affect adoption? Which artifacts create traceability without excessive burden? How does architecture change when outcomes are delayed, distributed, or politically contested?

Ethics and society

What constitutes meaningful choice in AI-mediated environments? How should cumulative effects across repeated low-stakes decisions be assessed? Which forms of participation genuinely redistribute influence rather than merely collect feedback?

CONCLUSION

Better decisions create better outcomes

The promise is not certainty. It is a more responsible way to move through uncertainty.

The essential move

Decision Architecture shifts attention from the isolated output to the conditions that make action warranted. It asks teams to define the purpose, understand the context, expose uncertainty, construct options, support judgment, enable action, and learn.

The role of AI

AI makes this work more urgent and more possible. It can widen access to analysis, generate alternatives, and accelerate learning. It can also obscure uncertainty, narrow agency, and scale weak architectures. Technology should strengthen judgment—not quietly replace it.

An invitation

The field is deliberately unfinished. Its concepts should be tested, its boundaries debated, its methods compared, and its claims revised. The practical starting point remains simple: identify one consequential decision and make its architecture visible.

Final question

Given what we know, what remains uncertain, and what matters most—what is the most responsible next action?

APPENDIX A

Decision Architecture Canvas

A compact working surface for a first review. Complete it with the people who make, execute, and experience the decision.

PURPOSE	What decision is being made? Who owns it? What outcome and boundaries define success?
CONTEXT	Who is affected? What evidence, constraints, incentives, resources, and dependencies matter?
UNCERTAINTY	What is unknown, inferred, changing, contested, or likely to reverse the recommendation?
OPTIONS	What alternatives—including delay, staged action, escalation, and status quo—are legitimate and feasible?
JUDGMENT	Which criteria and values govern the trade-offs? Who advises, decides, and can challenge?
ACTION	What happens next, by whom, by when, with what confirmation, safeguard, and stop condition?
LEARNING	What outcomes, overrides, appeals, and unintended effects will be reviewed, when, and by whom?

APPENDIX B

Decision record template

Decision statement

What commitment is being made, by whom, and for what time horizon? _____
_____

Evidence and context

What evidence and local context were material? What provenance and limitations apply? _____
_____
_____

Uncertainty and assumptions

What remained unknown? Which assumptions carried the result? What would change the choice? _____
_____
_____

Alternatives considered

Which options were considered, rejected, deferred, or left open—and why? _____
_____

Rationale and dissent

How were criteria and trade-offs applied? What disagreement or conflict of interest was recorded? _____
_____
_____

Action and accountability

Who will act, by when, with which safeguards, confirmation, and escalation path? _____
_____
____

Review and learning

When will outcomes be reviewed? What measures, feedback, appeals, and stop conditions apply? _____
_____
_____

APPENDIX C

Evaluation scorecard

Use ratings as prompts for evidence, not as a substitute for judgment.

Dimension	Key question	Evidence
Purpose	Is the decision and its legitimate outcome explicit?	Brief; stakeholder validation
Context	Are material people, constraints, and dependencies represented?	Map; research; data rationale
Uncertainty	Are unknowns and assumptions decision-relevant and visible?	Register; calibration; sensitivity
Options	Are alternatives genuine, feasible, and sufficiently diverse?	Option canvas; omission review
Judgment	Are criteria, values, authority, and rationale explicit?	Decision record; dissent
Action	Is the next step executable, reversible where possible, and owned?	Plan; controls; confirmation
Learning	Do outcomes and unintended effects change future decisions?	Monitoring; reviews; changes
Ethics	Are agency, fairness, privacy, security, and recourse protected?	Impact assessment; audit

Interpretation

Rate evidence as absent, emerging, adequate, or strong. Any critical ethical gate can override an aggregate score. Record the rationale and reviewer, then compare results over time rather than treating the scorecard as certification.

REFERENCES

References

Arrow, K. J. (1951). Social Choice and Individual Values. Wiley.

Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk. Econometrica, 22(1), 23–36. (Original work published 1738). https://doi.org/10.2307/1909829

Beyer, H., & Holtzblatt, K. (1998). Contextual Design. Morgan Kaufmann.

Buchanan, R. (1992). Wicked problems in design thinking. Design Issues, 8(2), 5–21. https://doi.org/10.2307/1511637

Dewey, J. (1938). Logic: The Theory of Inquiry. Henry Holt.

Dorst, K. (2011). The core of “design thinking” and its application. Design Studies, 32(6), 521–532. https://doi.org/10.1016/j.destud.2011.07.006

European Parliament and Council. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. Official Journal of the European Union. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. Annual Review of Psychology, 62, 451–482. https://doi.org/10.1146/annurev-psych-120709-145346

Hollnagel, E., & Woods, D. D. (2005). Joint Cognitive Systems. CRC Press.

ISO. (2019). ISO 9241-210:2019: Ergonomics of human-system interaction—Human-centred design for interactive systems.

REFERENCES

References / continued

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica, 47(2), 263–291. https://doi.org/10.2307/1914185

Klein, G. (1998). Sources of Power: How People Make Decisions. MIT Press.

March, J. G., & Simon, H. A. (1958). Organizations. Wiley.

National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). https://doi.org/10.6028/NIST.AI.100-1

Norman, D. A. (2013). The Design of Everyday Things (rev. ed.). Basic Books.

OECD. (2019, amended 2024). Recommendation of the Council on Artificial Intelligence (OECD/LEGAL/0449). https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. Human Factors, 39(2), 230–253. https://doi.org/10.1518/001872097778543886

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. IEEE Transactions on Systems, Man, and Cybernetics—Part A, 30(3), 286–297. https://doi.org/10.1109/3468.844354

Rittel, H. W. J., & Webber, M. M. (1973). Dilemmas in a general theory of planning. Policy Sciences, 4, 155–169. https://doi.org/10.1007/BF01405730

Schön, D. A. (1983). The Reflective Practitioner. Basic Books.

REFERENCES

References / continued

Simon, H. A. (1947). Administrative Behavior. Macmillan.

Simon, H. A. (1955). A behavioral model of rational choice. Quarterly Journal of Economics, 69(1), 99–118. https://doi.org/10.2307/1884852

Simon, H. A. (1969). The Sciences of the Artificial. MIT Press.

Thaler, R. H., & Sunstein, C. R. (2008). Nudge: Improving Decisions About Health, Wealth, and Happiness. Yale University Press.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. Science, 185(4157), 1124–1131. https://doi.org/10.1126/science.185.4157.1124

United Nations Educational, Scientific and Cultural Organization. (2021). Recommendation on the Ethics of Artificial Intelligence. https://unesdoc.unesco.org/ark:/48223/pf0000381137

von Neumann, J., & Morgenstern, O. (1944). Theory of Games and Economic Behavior. Princeton University Press.

Wiener, N. (1950). The Human Use of Human Beings. Houghton Mifflin.

Woods, D. D. (2016). The risks of autonomy: Doyle's catch. Journal of Cognitive Engineering and Decision Making, 10(2), 131–133. https://doi.org/10.1177/1555343416653562

ABOUT

About Felentra

Felentra is an emerging Decision Architecture practice developing frameworks, products, research, education, and selected advisory engagements.

The practice

Felentra works at the intersection of decision design, technology, research, and responsible action. It begins with the decision—not with software—and examines how context, uncertainty, options, judgment, workflows, safeguards, and feedback shape outcomes.

The founder

Felentra was founded by Felix Olusola Abayomi. His work spans technology, conservation, systems thinking, digital product development, and doctoral research into decision support for complex human–wildlife interactions. The recurring question behind the practice is: How can we design systems that help people make better decisions under uncertainty?

Begin with one decision

What is the one decision that, if consistently improved, would create the greatest positive impact in your organization?

CONTACT	hello@felentra.com · United States / Remote · Decision Architecture for responsible action.