

An AI mediated framework for improving human judgement through Socratic Reflection

Decision Intelligence Engine

Many important decisions are evaluated only after their outcomes are sown, even though research proves that metacognition and long transfer are best achieved through journaling during and right after the learning process, or in this case. The decision making process.

DECISION INTELLIGENCE ENGINE

Library Insights Reviews

Decision Intelligence Engine

What decision deserves your attention today?

Describe the decision you're facing...

Begin Reflection →

A structured space for clearer thinking. No answers — only better questions.

As a result, individuals often confuse the quality of a decision with the quality of its outcome, making it difficult to distinguish good reasoning from good luck.

Abstract

Important decisions are often evaluated only after their outcomes become known. However, research in metacognition, self-regulated learning, reflective practice, and long-term knowledge transfer suggests that cognitive growth is most effectively supported when individuals actively reflect on their thinking during and immediately after the reasoning process itself. The Decision Intelligence Engine was designed around this principle.

Rather than functioning as an AI assistant that provides answers, the system acts as an epistemic partner that helps users think more rigorously. Through structured Socratic dialogue, the engine challenges assumptions, explores alternative perspectives, surfaces cognitive biases, and encourages deeper reflection before a decision is made.

Throughout the interaction, the system continuously constructs a structured representation of the user's reasoning while simultaneously collecting longitudinal data about the decision-making process. Over time, these interactions become a personal dataset that enables users to better understand how they think, how their confidence changes, which biases recur, and how their decision quality evolves.

The objective is not to automate decision making, but to improve the quality of human judgment.

Problem

Decision making is one of the most fundamental cognitive processes humans perform, yet very few people systematically examine how they arrive at their conclusions.

Instead, individuals typically evaluate only the outcome of a decision. Once the outcome becomes known, hindsight bias alters memories of the original reasoning process. Assumptions become invisible, uncertainty is forgotten, confidence is reconstructed, and alternatives that once appeared plausible are often dismissed.

This creates two important problems.

- Well-reasoned decisions may appear unsuccessful simply because unfavorable outcomes occurred.

- Poor reasoning may become reinforced when fortunate outcomes create the illusion of good judgment.

As a result, opportunities for cognitive growth are lost because the reasoning process itself is rarely preserved.

System Design

The Decision Intelligence Engine transforms decision making from a conversation into a structured process of reflection.

Rather than opening into a traditional chatbot interface, users enter a dedicated thinking environment designed specifically for reasoning. The system presents one carefully selected question at a time, encouraging deliberate reflection instead of rapid interaction.

The AI intentionally avoids providing immediate recommendations. Instead, each response is designed to deepen the user's understanding of their own thinking through successive questions.

Throughout the conversation, the system dynamically builds a structured model of the decision, including:

- goals
- constraints
- supporting evidence
- assumptions
- uncertainties
- alternatives
- tradeoffs
- confidence
- expected outcomes

This evolving structure allows users to visualize their reasoning as it develops while simultaneously creating a permanent record of the decision process before hindsight can influence recollection.

AI-Mediated Socratic Reflection

The core contribution of the Decision Intelligence Engine is its conversational architecture.

Rather than functioning as an answer-generation system, the AI serves as a Socratic facilitator.

Its objective is to improve the user's reasoning by asking progressively better questions.

Examples include:

- What assumption are you making that may not be true?
- What evidence would change your mind?
- Which alternative deserves more consideration?
- What information are you currently missing?
- How might someone who disagrees with you interpret this situation?
- If this decision proves unsuccessful, what will most likely have been overlooked?

These questions encourage metacognitive reflection while helping users recognize cognitive biases, incomplete evidence, and hidden assumptions before making important decisions.

The AI therefore becomes a cognitive partner rather than a decision maker.

Measuring the Decision Process

Unlike conventional AI systems that focus primarily on producing outputs, the Decision Intelligence Engine continuously measures characteristics of the reasoning process itself.

Examples include:

- confidence calibration
- evolution of confidence throughout the conversation
- number of alternatives explored
- assumptions identified and revised
- counterarguments considered
- unresolved uncertainties
- evidence quality
- decision complexity
- reflection depth
- recurring cognitive biases

These variables create a multidimensional representation of how an individual thinks rather than simply what they decide.

Longitudinal Learning

Each completed interaction generates a structured Decision Portrait that summarizes both the decision and the reasoning process that produced it.

Over time, these Decision Portraits accumulate into a longitudinal dataset capable of revealing patterns across many decisions.

The system can begin identifying:

- recurring cognitive biases
- systematic overconfidence or underconfidence

- strengths across different decision domains
- recurring sources of uncertainty
- improvements in reasoning complexity
- changes in evidence utilization
- growth in reflective thinking

Rather than evaluating isolated decisions, the engine measures the evolution of human judgment itself.

Applications

Although initially designed as an individual cognitive support system, the underlying framework has broader applications across multiple domains.

Potential applications include:

- executive decision making
- operations management
- strategic planning
- entrepreneurship
- financial analysis
- organizational learning
- higher education
- leadership development
- AI-assisted learning research

Because every interaction produces structured cognitive data, the platform also offers opportunities to study metacognition, confidence calibration, decision quality, and human-AI collaboration within educational and organizational settings.

Current Status

The Decision Intelligence Engine has been developed as a functional prototype using Base44.

The current implementation includes a dedicated thinking interface, AI-mediated Socratic dialogue, structured reasoning capture, dynamic decision mapping, decision health metrics, and longitudinal decision logging.

Future work will focus on expanding the analytics framework, refining measures of cognitive growth, and validating whether repeated interaction with the system measurably improves decision quality over time.

Conclusion

The Decision Intelligence Engine represents a shift in the role of artificial intelligence within decision support.

Rather than replacing human judgment, it is designed to strengthen it.

By combining Socratic dialogue, structured reasoning, metacognitive reflection, and longitudinal analysis, the system encourages users to think more deliberately, question their assumptions more rigorously, and continuously improve the quality of their decision making.

This project reflects the broader mission of the Decision Systems Lab: designing intelligent systems that do not merely produce better answers, but help people become better thinkers.