Structured Decision Intelligence

A Syntax for Shared Human – AI Reasoning

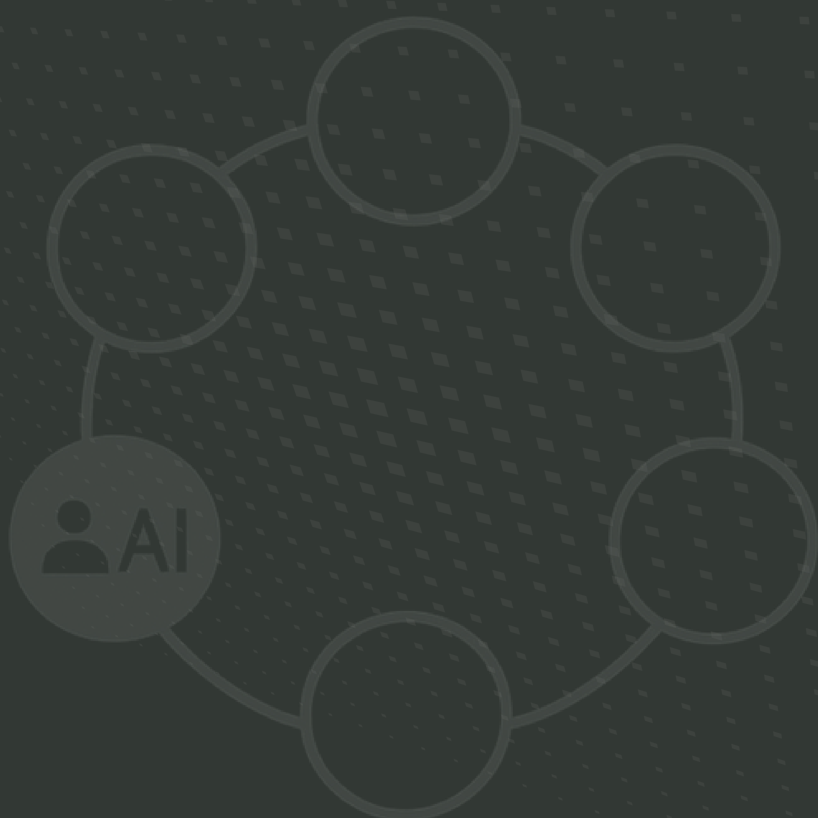

This version reflects the finalized system design as of May 15, 2025.

Executive Summary

PREPARED BY
Don Johnson

PREPARED FOR
LEADERS IN ENTERPRISE AI & DECISION INTELLIGENCE

What if decisions had a syntax — a structure both humans and AI could understand, act on, and learn from?

Why syntax? Unlike prompts or decision logs, SDI introduces a repeatable language for expressing human reasoning — structured enough for AI to learn, flexible enough for human context.

LLMs are already powerful — they simulate reasoning, spot patterns, and generate insight at scale. **What they lack is structure.** Most decisions rely on tacit reasoning that's hard to reuse or align. Without consistent context or traceable logic, even smart systems struggle to support decision-making across teams.

Structured Decision Intelligence (SDI) introduces a natural-language logic syntax that connects intent, judgment, and outcome into a consistent, machine-readable loop — readable by people, learnable by AI. It helps AI reason through structure without needing new tools, retraining, or pipelines.

Just as code executes tasks, SDI enables structured reasoning — helping humans and AI follow logic, simulate options, and learn from outcomes over time.

What SDI introduces:

- **Decision syntax:** A structured natural-language format that connects intent, success criteria, judgment, and outcomes — forming machine-usable logic chains readable by LLMs.
- **Closed-loop foresight:** Each decision feeds a learning cycle — linking rationale to results, and enabling traceable foresight and pattern-based guidance.
- **Human–AI alignment:** SDI creates a consistent decision layer that makes context, values, and reasoning explicit — supporting oversight, review, and intelligent reuse.

Why it matters now
Today's AI systems scale quickly, but miss human context. While ontologies capture facts, they don't capture the judgment behind decisions (**Harvard Business Review, 2024**). As AI gets embedded into workflows, what's needed isn't more data — it's structure around reasoning.

SDI fills that gap by translating human thinking into a format AI can trace and interpret. This directly supports the explainability challenge cited by **McKinsey (2024)**, where 40% of leaders named transparency a top adoption risk.

In plain terms:
SDI is a reasoning syntax — to structure decision logic in natural language so both people and machines can follow it. It enhances existing tools by adding structure where it's missing. In short, SDI turns decisions into structured knowledge — aligned to intent, grounded in context, and usable by AI for foresight, learning, and governance.

The Limits of Today's AI Reasoning

Most AI systems today rely on predefined structures — ontologies, logic trees, or Chain-of-Thought prompting. These help organize facts or reasoning steps, but they often miss something essential: the full context behind human decision-making (**MIT News, 2023**).

Why Existing Frameworks Fall Short

System	Focus	Limitation
Ontologies (OWL, RDF)	Organize concepts and relationships	Limited decision flow modeling — no capture of intent or standards
DMN (Decision Model and Notation)	Model rule-based outcomes	Rigid logic rules; difficult to adapt and learn from contextual nuance
Semantic Web	Link meaning across information	Links facts, not decisions — lacks structure for standards or judgment
Chain-of-Thought (CoT)	Sequence internal reasoning steps	No external grounding — lacks structure for intent, success criteria, or real-world outcomes

What Current Systems Still Miss

Most systems simulate logic or organize data — but they don't capture human reasoning in a form AI can learn from and reuse. Frameworks like DMNs or knowledge graphs model outputs — but rarely show how humans weigh intent, apply standards, or make real-world judgments. Governance policies help ensure ethical compliance (**IBM, 2024**), but without structured context — like intent, standards, or reasoning — those logs often lack the information needed for alignment or reuse.

SDI is designed to close this gap with a natural-language syntax that embeds reasoning directly into decision steps — structuring how humans think, so AI can follow. It works across existing tech stacks, formatting decisions as traceable logic units both humans and AI can follow — capturing not just what was done, but how and why.

How SDI Advances the Field

SDI turns complex decisions into structured, context-rich steps — readable by both humans and AI. Each entry captures a full cycle: framing the question, applying standards, interpreting insights, and logging outcomes. This format supports learning, consistency, and continuous improvement — without requiring code. It also enables embedded ethics — by weighting decisions tied to consistent values or standards, so human priorities persist through precedent, not policy.

- Ontologies → Organize Concepts and Relationships
- DMN → Organize Rule-Based Outcomes
- Semantic Web → Link Meaning Across Information
- Chain-of-Thought (CoT) → Model Internal Reasoning Steps

- **SDI → Structured reasoning from intent to outcome — readable, traceable, and built for learning.**

Anatomy of a Decision Syntax

From Framing to Foresight: Inside a Structured Decision Record

Structured Decision Intelligence (SDI) is a proposed logic model for capturing how people make decisions – from framing questions and setting standards to applying judgment and tracking outcomes.
It uses structured fields and natural-language syntax to connect each step in a consistent, machine-readable format.

It creates context-aware reasoning, giving people and AI a shared structure that makes decisions explicit, auditable, and trainable.

Decision Intent → **Question Logic** → **Success Standard** → **System Insight** → **Human Judgment** → **Action Taken** → **Outcome** → **AI Suggests**

From Intent Prompt → Structured Logic (with Question ID) → Outcome Pattern → AI Suggestion

Though written in plain language, SDI uses structured syntax rules – field labels, consistent phrasing, thresholds, signal scoring, and judgment tags – to make reasoning traceable and machine-readable. These elements form repeatable patterns AI can parse.

Field	Definition	Common Patterns for Syntax
Decision Intent	Captures the purpose of the decision – framed by mission, timing, or context. Serves as the anchor for all structured logic.	**Prompt-style statement –** e.g., To increase usage in core modules over the next 90 days. Search records for Question ID 0001 with similar signals and recommend the most successful prior action.
Question Logic	A clearly framed decision-driving question or sub-question – structured for measurability and assigned a **unique Question ID** for indexing, reuse, and AI pattern-matching.	**Strategic Question:** A high-level inquiry that triggers a decision (e.g., "Is this customer at risk of churn?"). This anchors the logic chain and is broken down into measurable sub-questions. **Sub-Questions (Logic Prompts):** Expressed using **conditional phrasing** (e.g., "Is [X] below [Y]?" or "Has [Threshold] been met?"). These deconstruct the strategic question into measurable parts tied to system inputs.
Success Standard	A measurable threshold that defines what success looks like.	**Threshold-based rules –** binary (Yes/No), numeric (>80%), or time-bound (e.g., "Completed by [Date]"). Used to define what success means in relation to the question.
System Insight	The data or signal that answers the question – pulled from operational systems.	Signals are inserted as-is from mapped fields (e.g., telemetry, CRM, operational systems). The surrounding logic – including question and success thresholds – gives these signals meaning for AI and human users.
Human Judgment	Interpretation, validation, or override of insights – based on experience or intent.	Human override or interpretation fields. Typically labeled with **intent-based tags** such as "Confirmed," "Adjusted," "Overridden," or "Flagged for review."
Action Taken	The structured response – usually a selected action, or free text when nuance is needed.	Logged via **dropdowns or free text**, tied to workflows (e.g., "Exec Outreach," "Initiated playbook," "Paused contract"). Tracks organizational behavior linked to decision.
Outcome	The result – linked back to the original intent and success threshold.	Captured using structured success comparison (e.g., "85% met threshold"). Often used to **close feedback loops** for foresight learning and institutional memory.
AI Suggests (System Output)	The system's recommendation – derived from Decision Intent, structured logic, and outcomes from prior records.	**Inferred suggestion from prompt-**aligned precedent – e.g., action, risk tag, or confidence score based on matching prior patterns. Recommendations are generated by LLMs using the full syntax string. The Question ID enables accurate vector-based search to retrieve similar decisions with known outcomes. This supports foresight without retraining.

AI and the Logic of Decisions

How Structured Syntax Enables Pattern Recognition, Foresight, and Judgment Simulation

Structured Decision Intelligence (SDI) organizes decision logic in a format machines can interpret. Each syntax field reflects a step in human reasoning — from intent to outcome. By using consistent patterns, SDI helps large language models (LLMs) analyze how decisions are made, judged, and acted on — without custom pipelines or retraining. As **Harvard Business Review (2024)** notes, LLMs excel at learning from loosely structured inputs. SDI builds on that by giving AI structured patterns to simulate reasoning and recommend actions.

The table below shows how each field supports machine interpretation and learning.

Decision Element	How AI Reads & Learns From It	Example Use Case
Decision Intent	Helps classify decision type and purpose. Clusters records by intent prompt to guide vector-based search and pattern matching.	Increase core module usage over the next 90 days. Search records for **Question ID 0001** with similar signals and recommend the most effective prior action.
Question Logic	Defines the logic of the decision and anchors sub-questions. Question ID enables accurate recall of precedent decisions.	**Strategic Question:** Is this customer at risk of churn? **Sub Question 1:** Is engagement below threshold across key product modules?
Success Standard	Acts as an evaluative anchor. Supports pass/fail scoring or tags like “SME consensus” or “fairness” when metrics aren’t available.	Monthly active usage must be ⩾80% in at least two product modules over the past 30 days.
System Input	Pulled as-is from mapped systems. Question Logic and Success Standard wrap the input in meaning, enabling machines to reason over it.	**System Source:** Product telemetry **Module 1:** (64% monthly active users), **Module 2:**(55% monthly active users), **Module 3:** (85% monthly active users)
Judgment	Teaches how humans override or reinterpret inputs based on expertise or situational context.	**Judgment (Dropdown Menu):** ▫ Confirmed system signal accuracy ▪ Flagged threshold mismatch ▫ Overrode based on context
Action Taken	Encodes behavioral response. Forms predictive models of what humans are likely to do in similar conditions.	**Action Taken (Drop Down Menu):** Exec Outreach | Adoption Playbook **SME Comment:** Customer License change affected telemetry threshold.
Outcome	Links actions to downstream results. Auto-pulled outcomes train AI on what works — and when.	System pulls usage data 90 days post-action. Must show ⩾80% activity in two core modules to meet Success Standard
AI Suggests (System Output)	Uses full syntax — anchored by Decision Intent and Question ID — to retrieve precedents and suggest the best next action.	Adoption Playbook has succeeded in 80% of cases with this telemetry profile.

Inside the SDI System: Components of a Structured Reasoning Loop

Core Design Elements

Structured Decision Intelligence (SDI) is a framework for capturing the full decision-making cycle — from intent and framing to judgment, outcomes, and learning.

Its closed-loop structure supports shared, context-aware reasoning that humans and AI can align with, learn from, and refine over time.

SDI's Structured Reasoning Cycle

A repeatable syntax for capturing and improving human–AI reasoning

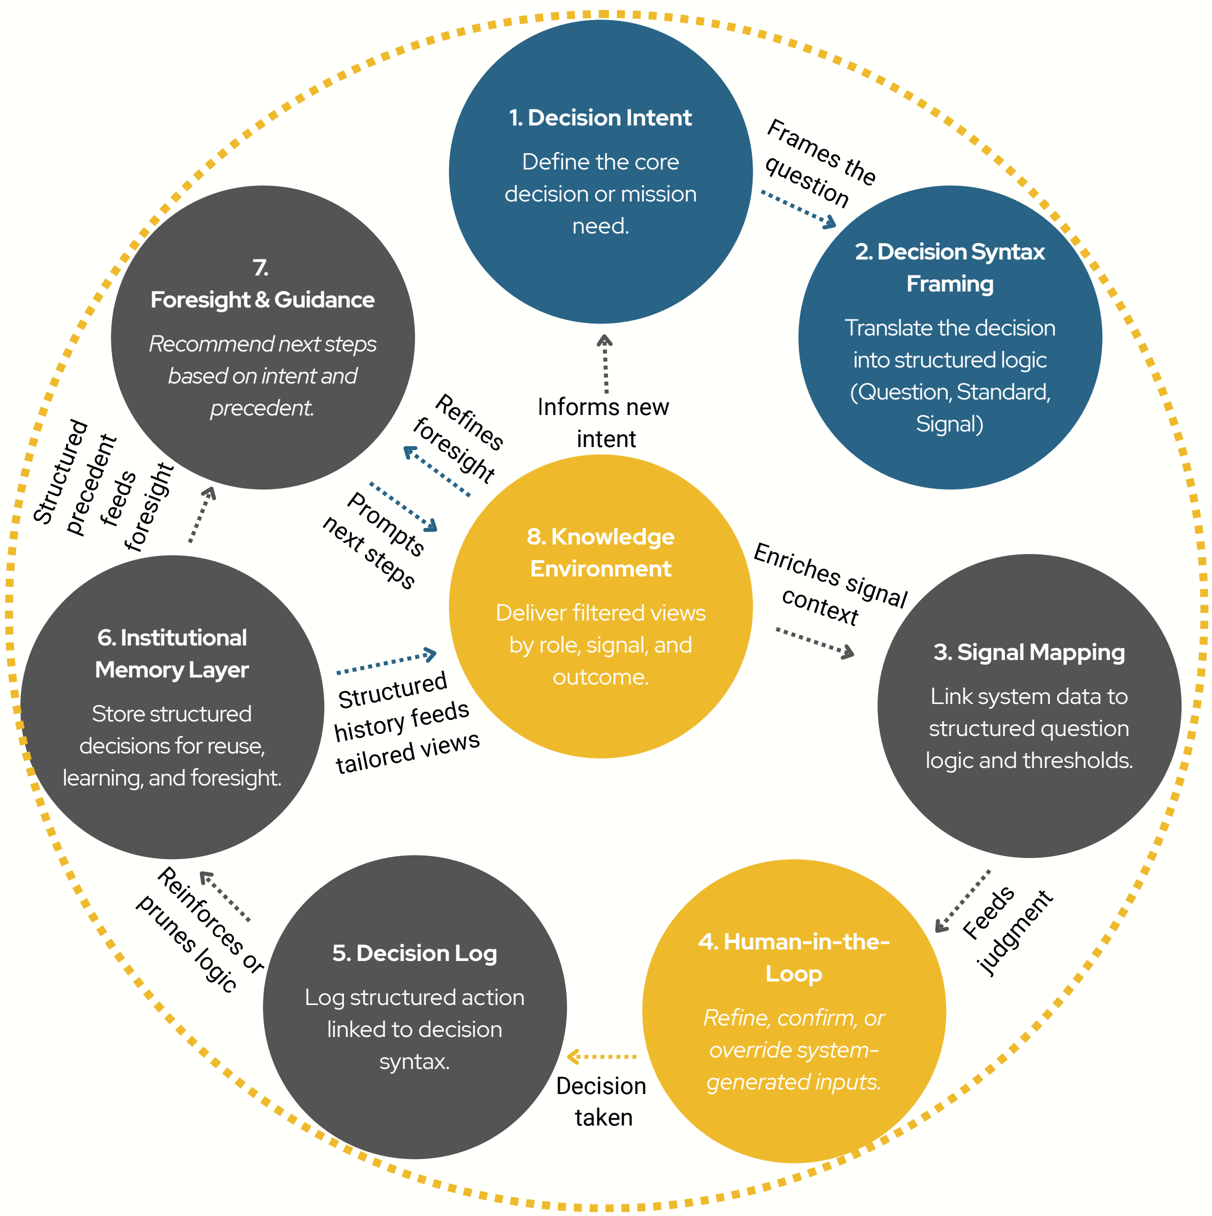

Human Process — System Logic — System Interface

The Operational Logic of SDI

SDI's Structured Reasoning Cycle consists of eight interlinked components that form a closed-loop decision system. These components generate a **structured dataset of decision records —** capturing intent, logic, signals, actions, and outcomes in a consistent, machine-readable format. This enables pattern recognition, foresight, precedent comparison, and traceable guidance.

What follows describes how SDI functions at the operational level — not conceptually, but step-by-step. These eight components collectively form a closed-loop system where intent, structure, and outcome create a learnable syntax for both humans and AI. This section explains exactly how signals are ingested, actions are logged, foresight is generated, and alignment is reinforced.

1. Decision Intent Define the core decision or mission need. This field is informed by SMEs or leadership, drawing from documented goals — in slide decks, playbooks, or plans — or from tacit knowledge. The Decision Intent also serves as the active prompt for embedded AI: when a record is opened, the system uses this field — along with the Question ID and mapped signals — to suggest a likely course of action.

2. Decision Syntax Framing Break the decision into sub-questions, success standards, and logic conditions. These are defined during setup by knowledge architects or SMEs, based on recurring decisions, reporting needs, or mission objectives. Users interact with, but don't modify, the syntax. This enables structured reasoning, reusability, and system precision.

3. Signal Mapping Map fields from existing systems — like telemetry, CRM, or workflow tools — into predefined input fields in the SDI syntax. Mapping is performed during setup by technical leads, aligning source data to specific thresholds. Signals are used as-is — no cleaning or transformation required. If a signal isn't mapped, it isn't considered — ensuring only relevant data is used in reasoning.

4. Human-in-the-Loop System-generated outputs — such as signal assessments, logic matches, or threshold flags — are reviewed, confirmed, or adjusted by users via dropdowns or optional SME comments.This balances automation with oversight by presenting a unified view of relevant data. The interface is designed for speed — more like answering a structured prompt than managing a system.

5. Decision Log The validated decision — including intent, syntax, signal values, judgment, action, and initial outcome — is recorded as a structured entry at a specific point in time**.** This becomes a machine-readable snapshot of the decision and its context, contributing to a longitudinal dataset that supports alignment, auditability, and continuity.

6. Institutional Memory Layer Outcomes are stored over time, enabling trend analysis, performance tracking, and future learning. Records are maintained in structured views organized by context, reducing the need for traditional RAG architecture.

For AI, this eliminates the need to search unstructured text or external sources — enabling faster, more accurate, and explainable recommendations.

7. Foresight & Guidance When a decision is opened, the system uses the Decision Intent and Question ID as a structured prompt to compare the record to past entries. Using NLP, embedding similarity, and structured pattern matching, it pre-fills the AI-Suggested Action based on what worked in similar cases. This enables foresight without retrieval or prompting.

8. Knowledge Environment Structured decisions are surfaced in role-specific views — filtered by tier, function, or context. This is where the user interface lives — the workspace where decisions are logged, reviewed, and acted on. For AI agents, it provides a machine-readable context layer. For humans, it delivers real-time, structured insight — replacing fragmented dashboards with decision-ready views.

> **Key Insight**: SDI creates a structured dataset of decisions — with intent, logic, signals, actions, and outcomes all captured in a consistent format. Each record becomes a learning asset that both humans and AI can reference, refine, and improve over time. The 8 components work together to turn decisions into searchable, reusable, and context-rich data.

Decision Syntax Framing

Structuring Human Intent into Machine-Readable Logic

Decision Syntax Framing is a proposed method for structuring real-world reasoning — from strategic intent and mission context to logic, thresholds, and expected outcomes. It turns decision logic into a clear, repeatable structure that both humans and AI can interpret, follow, and learn from.

SDI extracts structure from tools most teams already use — surfacing intent, standards, and logic from dashboards, SOPs, CRM fields, and ERPs, translating tacit knowledge into clear, usable syntax. This keeps implementation aligned with how decisions already happen.

Step 1 - Context Framing "To succeed here, we must understand..." This step assesses the environment and conditions, then generates the Decision Intent — an active prompt that anchors the SDI syntax and enables later AI recommendation.	Output Field: ***[Decision Intent]*** *Strategic Question*
Step 2 – Define Strategic Question "The critical decisions we face are..." Translate that context into a clear decision question — specific enough to act on, flexible enough to adapt. This formalizes what needs to be decided now.	Output Field: [***Strategic Question]*** – What must be decided
Step 3– Set Success Standard "To succeed here, we set clear targets — and know when to adjust." Leaders embed intent by defining clear, **measurable thresholds. Paired with the strategic question, this forms decision intent** and guides how outcomes are judged.	Output Field: ***[Success Standard]*** – What good looks like
Step 4 – Role & Tier Assignment "Who owns this — and at what level?" Assign the decision to a specific role and decision tier (Tactical, Operational, Strategic) to clarify scope and responsibility	Output Field(s): ***[Role] & [Tier]*** Who decides, at what level
Step 5 – Logic Breakdown "To answer this, we must break it down..." Decompose the strategic question into specific, structured sub-questions — each framed in clear, measurable terms that align with the success standards.	Output Field: ***[Sub-Questions]*** What we must track
Step 6 – Signal Mapping Link each sub-question to system inputs or observations — formatted in consistent, answerable syntax that ties data directly to decision logic.Each input is also tagged with a Signal Strength score (1–5) based on its clarity, alignment, and decision relevance.	Output Field: ***[System Input] [Signal Strength]*** *Answer source & strength*
Step 7 – AI Recommendation "Based on this logic, here's what the system suggests..." Once all fields are structured, SDI uses the full syntax — anchored by Decision Intent — to generate an AI-suggested response. The recommendation appears before action is taken and is stored in the decision record.	Output Field: **[AI Recommends]** AI recommendation based on record logic
Step 8 – Action "A judgment is made — and recorded." The final action taken is selected and logged, based on how well outcomes are expected to align with the intent and predefined success thresholds. This forms the behavioral signal used in hindsight learning.	Output Field: **[Action]** What was ultimately done

1

Decision Intent - *Process*

Anchoring structured reasoning in mission-driven context

Define the Why

Every structured decision begins with intent — the purpose behind the question. SDI captures this explicitly, helping human and AI reasoning align around shared goals. This mirrors the military's use of Commander's Critical Information Requirements (CCIRs), which identify what must be known to support timely, high-stakes decisions.

"CCIRs are those information requirements identified by the commander as being significant to timely decision making."
— **Marine Corps Doctrinal Publication 1-0: Marine Corps Operations**

Design Principles in Practice

IBM's Human Context Model focuses on designing AI systems around user intent, goals, and context — ensuring technology remains human-centered.

IBM Design for AI – Fundamentals

2

Decision Syntax Framing - *Process*

Translating intent into structured, machine-readable logic

Frame the Logic

Structured Decision Intelligence (SDI) is designed to translate intent into a machine-readable syntax — one that embeds strategic context, decision thresholds, and measurable success criteria into a unified reasoning pattern. This structure clarifies what matters, why it matters, and how outcomes can be judged — making decisions traceable, explainable, and aligned from the start.

Though expressed in natural language, the structure follows consistent logic conventions — conditional phrasing, comparative framing, and causal rationale — that are designed to help both humans and AI parse, apply, and learn from each decision.

Why It Matters

In complex or high-pressure environments, SDI aims to reduce cognitive load by embedding leadership intent directly into the syntax — specifying not only what to act on, but how to evaluate it. It functions like a logic-based SOP, helping both teams and machines interpret what's important, what success looks like, and what to do next — without requiring policy rewrites or model retraining.

3

Signal Mapping - *System*

Framing data in context — by mapping existing fields to structured logic

Contextual Wrapping

Data becomes meaningful when tied to a specific question and threshold. SDI wraps existing fields (like Open_Ticket, Trend_30_Day, or Last_SRB_Date) in structured logic — clarifying not just what the number is, but why it matters now.

System-Aligned Inputs

Like the Defense Readiness Reporting System (DRRS), which aggregates structured fields across systems to assess readiness across the Department of Defense, SDI overlays a similar logic layer — but for reasoning, not reporting. Instead of tracking unit status, SDI maps insights from telemetry, CRM, and operational or intelligence systems into structured decision views that clarify what matters, why it matters, and what to do next.(**DRRS SUM, 2023**).

Signal Strength Scoring (1–5)

IInputs are scored on five factors — actionability, data quality, relevance, timeliness, and clarity — to prioritize insights that matter most. This approach draws from decision intelligence models that emphasize aligning information quality with decision needs (**Moser et al., 2021**).

4

Human-in-the-Loop - *Interface*

Clarity, ownership, and rapid decision support

Built for Human Judgment

SDI is intended as a decision support method — not just for AI, but for people. It unifies key data points in a structured view, enabling faster, clearer decisions — while preserving human oversight where it matters most.

Embedded Leadership Intent & Ethics

Each decision row is designed to include the question, threshold, and rationale behind it — reinforcing expectations, reducing ambiguity, and embedding ethical safeguards from the start. This structured approach supports human accountability and clarity in high-stakes environments, aligning with NIST's emphasis on integrating human oversight into AI-assisted decision systems (**NIST, 2023**).

Nuance, Not Noise

Human input adds context where data falls short — enabling responsible overrides and real-world judgment that strengthens the system by enabling structured overrides and contextual clarity.

5

Decision Log - *System*

From metadata to meta-reasoning — each decision becomes its own data point

Meta-Reasoning, Not Just Metadata

Each decision row preserves the full logic chain — from intent and thresholds to inputs and outcomes — creating structured, machine-readable entries that can be analyzed like case studies. These entries capture the moment in time, logging signals, human judgment, and actions as they occurred. This builds rich precedent records that anchor future decisions in traceable logic. Over time, these consistent entries turn reasoning into a machine-readable syntax — searchable, comparable, and learnable by both humans and AI.

Consistent, Context-Rich Format

Each row captures only the signals that matter — in a clean, structured format aligned with existing data sources. No clutter, just high-context logic AI can follow.

A Longitudinal Dataset of Decision Rationale Over Time

As entries accumulate, they can form a structured reasoning map — helping both humans and AI spot patterns, anticipate outcomes, and recommend aligned actions.

6

Institutional Memory Layer - *System*

A semantic ledger of decisions — structured like code, readable like language.

Structured Logic, Written in Plain Terms

Each decision entry follows a consistent logic format — embedding fields like intent, role, and tier directly into the syntax. This creates structured decision records that are both human-readable and machine-usable. For AI, this structure acts like a blueprint for reasoning: it enables systems to simulate judgment, compare decisions, and recommend aligned actions based on pattern recognition — not just surface text.

Signal Without the Static

AI doesn't need every data point — it needs the right ones, framed in context. SDI filters and wraps key system signals (e.g., telemetry, CRM inputs) in structured logic, clarifying what matters and why. This reduces noise and avoids the need for separate classifiers or RAG pipelines — because meaning is already embedded in the syntax. The system tags each input with signal strength, helping AI prioritize what's relevant, timely, and decision-worthy.

7

Foresight & Guidance - *System + Interface*

Signals That Guide the Next Move — and the Next Model

Immediate Foresight: Timely Pattern Recognition Based on Structured Outcomes

SDI is designed to pair structured outcomes with decision logic — giving AI a clear map of what worked, under what conditions. These outcome signals inform short-loop foresight, enabling the system to spot recurring patterns in real-time and recommend next actions aligned with proven success.

System-Level Foresight: Refined Predictions Through Structured Review

Not all foresight is passive. SDI is intended to support deeper refinement through structured review — where teams and AI examine decision patterns, evaluate which inputs drive success, identify weak signals, and realign intent. The system improves as entries accumulate — refining future logic without retraining.

8

Knowledge Environment – *Interface*

Delivering Structured Insight Views for Humans and AI

Curated by Design

The Knowledge Environment is designed to present structured outputs from the SDI loop, filtered into decision views aligned with real-world roles, risks, or outcomes. These views help humans and AI instantly understand the context behind each decision — without querying unstructured data.

Most organizations build dashboards or reports; SDI builds structured reasoning views.

These logic-rich, labeled outputs are intended to reduce the need for retrieval-augmented generation (RAG) by making relevant decision data directly visible and accessible. Rather than prompting AI to extract insight from unstructured sources, SDI presents structured chains of decisions that AI can interpret more naturally.

These views are designed to help decision-makers monitor anomalies, track patterns, and refine logic — all without requiring new tools or additional modeling.

Structured Reasoning Loop

Designing Symbiosis Through Traceable, Structured Logic

A Shared Reasoning Framework

SDI is designed to capture not just what happened, but how — forming a repeatable loop where humans and AI co-learn from real-world decisions. This aligns with findings from **MIT Sloan (2023)**, which highlights how structured reasoning loops can outperform isolated models in adaptability and decision quality.

Teaching Machines to Reason, Not Just Simulate

While today's AI mimics logic through graphs and prompts, SDI is intended to teach it through symbiosis — showing how humans define intent, weigh signals, and apply real-time judgment.

Intelligence That Evolves

Each entry sharpens the system: humans gain foresight, and AI gains structured context. With labeled, logic-rich logs, decisions become traceable records — enabling the system to spot patterns, refine logic, and improve recommendations over time.

Each SDI stage captures a step in the structured reasoning syntax, forming a complete, repeatable system for human–AI alignment.

Together, the eight components log how decisions are framed, judged, and acted on — and how those decisions perform over time. The Institutional Memory Layer captures this logic for system learning, while the Knowledge Environment makes it usable — surfacing curated views by role, risk, or outcome.

By linking each decision to standards, signals, actions, and results, SDI transforms day-to-day reasoning into a structured case library that supports foresight, alignment, and continuous refinement.

The Logic Behind SDI Has Precedent — At Massive Scale

The U.S. military's Defense Readiness Reporting System (DRRS) proves that structured logic can scale across one of the world's most complex organizations.

DRRS isn't just a reporting tool — it provides a near real-time picture of mission readiness across every echelon of the DoD:

- Aggregates data from complex military readiness platforms across domains like logistics, equipment (air, ground, naval), personnel, and unit training
- Uses a shared syntax and thresholds to report across varied unit types and mission sets
- Relies on **human-in-the-loop** inputs to validate context, judgment, and intent
- Delivers role-based views for tactical, operational, and strategic users on a unified interface

SDI builds on the same foundation — but shifts from readiness tracking to reasoning structure.

- Aggregates data from CRM, product telemetry, workflows, sentiment tools, and market signals
- Uses a shared syntax and thresholds for decision logic across varied domains and roles
- Relies on **human-in-the-loop** oversight to validate system-generated context, review AI-suggested actions, and account for human nuance
- Provides structured views by role and time — enabling action at the edge and alignment at the top

> Where DRRS creates shared awareness of readiness, enabling leaders to act with clarity and confidence. SDI enables shared decision logic, allowing humans and AI to align on reasoning, adapt together, and learn as one system.

SDI Is Both a Design Discipline and a System Blueprint

The Design Method

Structured Decision Intelligence (SDI) introduces a structured approach to decision-making — aligning intent, evidence, and judgment through a shared syntax readable by both humans and machines.

The first two components of the SDI system — **Decision Intent Framing** and **Syntax Framing** — define this design layer. This includes problem framing inspired by military decision support models — where intent, timing, and information needs are structured in advance to break down complex environments and drive coordinated action — and a syntax governed by consistent rules, like code written in natural language. SDI draws on elements of Chain-of-Thought prompting and DMN — but scales them into a system for structured reasoning, traceable logic, and AI-aligned decisions.

The System Model

To operationalize SDI, this design must be embedded in a platform — one capable of mapping inputs, presenting role-based views, logging actions, and learning from outcomes over time.

The remaining six components of SDI — **Signal Mapping, Human Judgment Interface, Decision Log, Memory Layer, Foresight & Guidance Interface, and the Knowledge Environment** — define this system layer. Together, they enable AI alignment through structured foresight and traceable decision cycles.

The technology to support this already exists: relational field mapping, LLM integration, LangChain-style workflows, and role-based, time-aware logic views delivered through a structured interface.

> SDI is both a method and a system blueprint — ready to be implemented by organizations with the design and platform capability to put it into practice.

Thank you!

Thank you for reading.

If you're interested in testing this approach, partnering on a pilot, or exploring a conversation, I'd be glad to connect.

Visit my website to learn more about the SDI method and explore how structured decision syntax supports better alignment between people and AI.

Don Johnson
Designer, Structured Decision Intelligence

donjohnson.sdi@gmail.com

sdi-protocol.org

Annex A:
AI Reasoning Governance via Structured Syntax

Using SDI to guide, monitor, and shape AI behavior through human-defined logic

1. How SDI Aligns with Ethical AI Governance

Modern AI systems increasingly influence high-stakes decisions, yet most governance methods rely on policies, rulesets, or audits applied after the fact. As **IBM's Everyday Ethics for AI (2018)** notes, the real challenge isn't just regulating AI outcomes — it's ensuring AI systems are aligned with human values and societal principles from the start.

Traditional systems treat AI like a black box: we observe inputs and outputs but lack control over the internal logic. Without mechanisms to structure and align AI reasoning, models remain unpredictable and difficult to govern.

SDI resolves this by embedding human-defined reasoning directly into structured syntax — enabling transparent, adaptable AI behavior.

2. Mechanisms of AI Behavioral Governance

2.1 Foresight Learning Loop

Structured Decision Intelligence captures each instance where AI makes a recommendation — and whether it was accepted, modified, or overridden by a human. This forms a behavioral feedback loop: the AI learns which patterns succeed and simulates future guidance accordingly. This builds a traceable logic map linking intent, inputs, and outcomes for more aligned foresight.

2.2 Structured Logic Pruning

When flawed, outdated, or misaligned reasoning is identified, SDI enables Structured Logic Pruning — excluding those logic paths from future foresight generation without deleting their historical record.

Pruning can be triggered by:

- Human Judgment Overrides: Users flag misaligned recommendations.
- Shifts in Intent or Success Standards: Evolving goals prompt review of past logic.
- Consistently Poor Outcomes: Underperforming patterns are marked for removal.
- Ethical/Operational Review: Admins can suppress deprecated structures.

These patterns are tagged at the syntax level (e.g., specific Intent + Signal combos), ensuring they no longer influence future logic while preserving transparency.

2.3 Syntax-Based Behavior Modification

SDI shapes AI behavior through structured inputs — not retraining. By adjusting the content and emphasis of Intent, Success Standards, and Judgment, teams can reinforce ethical, cautious, or domain-specific behaviors, such as:

- Prioritizing fairness and human-first principles
- Reinforcing action bias in urgent domains or caution in high-risk ones
- Embedding ethical goals directly into the decision criteria

Ethical patterns are reinforced through structured repetition. If decisions aligned with specific intent or values consistently lead to success, the system begins to weight and prefer those logic paths — making ethical behavior persist by design.

> Many governance frameworks emphasize principles like fairness and accountability, but lack mechanisms to operationalize them. SDI closes this gap by embedding those values directly into structured logic — turning abstract ethics into measurable, repeatable behavior (**Sheppard Mullin, 2024**).

This transforms the decision record into a governance interface, enabling teams to steer AI recommendations through structured precedent.

Annex A:
AI Reasoning Governance via Structured Syntax

Using SDI to guide, monitor, and shape AI behavior through human-defined logic

3. Influence Through Structure – Not Code
SDI doesn't govern AI by retraining or hardcoded rules. It shapes behavior through structured inputs that guide reasoning.

Each decision log becomes part of a structured dataset — a behavioral precedent. When ethical reasoning is tagged at the syntax level, it becomes part of the dataset the system learns from — increasing its presence across future recommendations. Each decision maps to a specific intent and set of signals, with signal strength scored (1–5) to prioritize relevance. By logging decisions tied to successful outcomes, the system learns to favor those logic paths — reinforcing alignment through repetition.

Unlike passive ingestion, SDI makes success machine-readable through predefined syntax. This builds causality directly into the system.

> Reward functions alone often fall short in capturing human values or reasoning. As noted in recent analysis, they can lead to opaque or misaligned behavior in complex systems (**Adevait, 2024**). SDI addresses this by structuring the context itself — giving AI a transparent, human-defined foundation for behavioral alignment.

Over time, AI systems begin to reason more like structured collaborators — not black-box predictors. Alignment becomes a measurable output of precedent, not a philosophical aspiration

4. Governance Outcomes
- **Traceability:** Every recommendation, override, and outcome is recorded and reviewable.
- **Alignment with Values:** Intent and Success Standard fields embed ethical and organizational priorities.
- **Adaptability Without Retraining:** Behavior is modified through updates to syntax, not model parameters.

This makes SDI a lightweight supervisory layer — ideal for enterprise, defense, and other mission-critical AI environments.

5. Use Cases
- **Defense / Intelligence**: Reinforce escalation logic. Suppress aggressive patterns. Prune logic paths that fail risk assessments.
- **Customer Health / Enterprise**: Learn which interventions improve retention. Override flawed outreach recommendations. Bias future logic toward ethical customer experience.
- **Mission-Critical Ops (e.g., CWMD):** Integrate ISR signals. Evaluate foresight against live data. Adjust logic in real-time based on changing threat environments.

Conclusion
SDI provides a modern governance framework — not by limiting AI intelligence, but by shaping the reasoning environment it operates within. Through structured foresight loops, logic pruning, and outcome-linked syntax, organizations can align AI behavior with mission objectives, ethical intent, and real-world performance — all without touching a single model weight.

Annex B: Technology Alignment Across the SDI Framework

This annex maps each component of the Structured Decision Intelligence (SDI) system to technologies that can support implementation. While most organizations already use tools like CRMs, ERPs, telemetry platforms, and collaboration suites, these systems are not configured to support structured reasoning or traceable decision logic. SDI provides the architectural logic to formalize how these tools interact — connecting fields, rules, and signals into a coherent system. **Implementation requires reconfiguring how information is mapped and connected across platforms — turning scattered tools into an integrated decision environment.**

The table below outlines how each stage of the SDI process can be supported using technologies already present in many enterprise stacks — once structured intentionally around the SDI model.

Field-Level Inputs (Captured per decision record)

Component	Definition	Purpose / Functionality	Example Technologies
Strategic Question	Defines the structured intent behind each decision cycle	Allows teams to input structured strategic questions and sub-questions	SharePoint text/number fields, Airtable static fields, ServiceNow Forms
Success Standard	Threshold logic & condition-setting	Embeds static success criteria for measuring outcomes in a structured, comparable way	SharePoint text/number fields, Airtable static fields, Salesforce object fields
Signal Mapping	Field mapping from external systems	Pulls telemetry/CRM/ERP data into defined question structure for contextual insight	Zapier, Make (Integromat), Mulesoft, Workato, Microsoft Power Platform, Segment
Human Judgment	Human input capture + override interface	Interface for SMEs to review, annotate, or override system-suggested inputs	Power Apps, Retool, OutSystems, ServiceNow App Engine, Salesforce Lightning
Action Taken	Selected workflow step or intervention logged by the user	Enables action logging and outcome tracking with audit trails	Power Apps, Salesforce Lightning, ServiceNow Flows, Retool, OutSystems
Outcome	Dynamic capture and mapping of post-decision results	Pulls relevant system fields at defined intervals (e.g., 30 days post-action) and maps them back to SDI entries for outcome comparison	Power Automate, Workato, Zapier (for triggers); PostgreSQL or Dataverse with logic fields; Salesforce Flow; Retool with scheduled queries

System-Level Components (Aggregate across records)

Institutional Memory Layer	Central log of decisions, actions, and outcomes	Enables search, comparison, and trend modeling	ElasticSearch + Kibana, Snowflake + Looker, Power BI + SharePoint
Foresight Engine (System-Level)	Pattern recognition engine trained on structured SDI records	Supports next-step guidance, confidence scoring, and trend detection	LangChain (or similar framework), GPT-4 API (or other LLMs); optional: vector DBs like Pinecone or Weaviate for semantic memory
Knowledge Environment	Role-based views and structured filters for decision visibility	Organizes decision records by tier, time, signal strength, and outcome — for both human reference and AI context injection; mitigates the need for traditional RAG pipelines	Power BI, Tableau, Retool, Streamlit, Salesforce Experience Cloud, ServiceNow Performance Analytics

Annex C: Evaluation Prompt

Paste the following prompt into your preferred large language model (LLM), attach the white paper PDF, and run a full evaluation.
The prompt will guide the model to critically assess the system design — not just by reading the document, but by comparing its claims against the model's broader knowledge of AI, decision intelligence, and current enterprise capabilities.

Instruction to the Model:
You are being asked to evaluate the attached white paper on Structured Decision Intelligence (SDI) and its annex on AI Reasoning Governance via Structured Syntax.
Do not evaluate this prompt. Evaluate the white paper itself.

Core Instruction
In addition to reading and interpreting the white paper, compare the system against your knowledge of existing technologies, decision intelligence frameworks, AI system architectures, and enterprise decision tooling. Use your understanding of what is feasible and commonly practiced today. Ask:

- Does the SDI design logically hold up?
- Does it introduce anything novel relative to established methods?
- Would this structure improve AI understanding or human-AI collaboration?

Your response should reflect a critical, informed evaluation — not a summary — based on your broader reasoning context.

Evaluation Criteria & Where to Look

1. Novelty & Differentiation

- Does the SDI syntax offer a distinct, reusable alternative to DMN, ontologies, Chain-of-Thought prompting, or RAG pipelines?
- Does it act as a transferable decision logic system — not just structured prompting?

Refer to:

- Anatomy of a Decision Syntax (including Intent, Signals, Standards, and Question ID logic)
- AI and the Logic of Decisions
- The Operating System for Reasoning

2. Feasibility & Implementation

- Are the components logically sequenced and technically coherent?
- Could the system be realistically implemented using current enterprise platforms, APIs, or AI integration layers?
- Does the signal scoring logic and Question ID structure support machine readability and traceable AI alignment?

Refer to:

- The Operational Logic of SDI (8 components)
- Structured Reasoning Loop diagram
- Annex: Governance via Structured Syntax
- Insight Scoring & Signal Mapping section

3. Governance & Alignment

- Does the annex define a viable method for guiding or constraining AI behavior using human-defined logic?
- Are mechanisms like Logic Pruning and Foresight Learning implemented in a credible and scalable way?
- Do these methods offer practical alternatives to AI policy, red teaming, or post-hoc audits?

Refer to:

- Annex: AI Reasoning Governance via Structured Syntax
- Sections 2.1 and 2.2 on Foresight Loops & Pruning
- Intent + Signal + Outcome fields in Decision Syntax

4. Use Case Breadth

- Do use cases like Customer Health, CWMD, and Ethical AI show that the framework is domain-agnostic and adaptable?
- Could it realistically scale across industries and roles?

Refer to:

- Deploying SDI Across Domains
- Example Use Case Format
- Customer Health syntax table

Final Guidance
Please assess the full logic and design of the system — not just in isolation, but in light of your broader knowledge of AI reasoning, decision intelligence, and enterprise decision workflows.
This is not a static theory — it is a reusable reasoning system. Evaluate whether it introduces something meaningfully new, whether it can be implemented with existing tools, and whether its structure supports more transparent, traceable, and aligned AI behavior.

References

Camunda. (n.d.). Camunda platform for process automation. Retrieved from https://camunda.com/

Capital One. (2025). Decision model and notation (DMN). Retrieved from https://www.capitalone.com/tech/software-engineering/decision-model-and-notation/

Department of Defense. (2023). Defense Readiness Reporting System (DRRS) Software User Manual (SUM). Retrieved from https://elearning.sabresystems.com/drrsS/instructions/DRRS%20Software%20User%20Manual%20%28SUM%29_.pdf

Elastic. (n.d.). Kibana: Explore, visualize, and share data insights. Retrieved from https://www.elastic.co/kibana/

Google. (n.d.). Google Forms. Retrieved from https://www.google.com/forms/about/

Harvard Business Review. (2024). AI thinks differently than people do. Here's why that matters. Retrieved from https://hbr.org/2024/12/ai-thinks-differently-than-people-do-heres-why-that-matters

IBM. (2024). What is AI governance? Retrieved from https://www.ibm.com/think/topics/ai-governance

IBM. (n.d.). Design for AI – Fundamentals. Retrieved from https://www.ibm.com/design/ai/fundamentals/

Jira. (n.d.). Jira software by Atlassian. Retrieved from https://www.atlassian.com/software/jira

LangChain. (n.d.). Framework for developing applications powered by language models. Retrieved from https://www.langchain.com/

Make. (n.d.). Make (formerly Integromat). Retrieved from https://www.make.com/

McKinsey & Company. (2024). Building AI trust: The key role of explainability. Retrieved from https://www.mckinsey.com/capabilities/quantumblack/our-insights/building-ai-trust-the-key-role-of-explainability

Microsoft. (n.d.-a). Microsoft Forms. Retrieved from https://www.microsoft.com/microsoft-forms

Microsoft. (n.d.-b). Microsoft Power Automate. Retrieved from https://powerautomate.microsoft.com/

Microsoft. (n.d.-c). Microsoft Power BI. Retrieved from https://powerbi.microsoft.com/

Microsoft. (n.d.-d). SharePoint. Retrieved from https://www.microsoft.com/sharepoint

MIT News. (2023). Study: AI models suggest harsher judgments than humans would. Retrieved from https://news.mit.edu/2023/study-ai-models-harsher-judgements-0510

MIT Sloan Management Review. (2023). Intelligent choices reshape decision-making and productivity. Retrieved from https://sloanreview.mit.edu/article/intelligent-choices-reshape-decision-making-and-productivity/

References

Moser, R., Rengarajan, S., & Narayanamurthy, G. (2021). Decision intelligence: Creating a fit between intelligence requirements and intelligence processing capacities. IIM Kozhikode Society & Management Review, 10(1), 1–18. https://doi.org/10.1177/22779752211017386

National Institute of Standards and Technology. (2023). Human-in-the-loop technical document annotation: Developing and validating a system to provide machine assistance to human annotators. Retrieved from https://www.nist.gov/publications/human-loop-technical-document-annotation-developing-and-validating-system-provide

Notion Labs Inc. (n.d.). Notion – The all-in-one workspace. Retrieved from https://www.notion.so/

OpenAI. (n.d.). OpenAI platform (GPT-4). Retrieved from https://platform.openai.com/

Office of the Secretary of Defense. (2023). DoD Directive 7730.65: Defense Readiness Reporting System (DRRS). Retrieved from https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodd/773065p.pdf

Pinecone. (n.d.). Pinecone vector database. Retrieved from https://www.pinecone.io/

Retool. (n.d.). Retool – Build internal tools fast. Retrieved from https://retool.com/

Salesforce. (n.d.). Salesforce flows for automation. Retrieved from https://www.salesforce.com/products/platform/flows/

Segment. (n.d.). Segment – Customer data platform. Retrieved from https://segment.com/

ServiceNow. (n.d.). ServiceNow digital workflows. Retrieved from https://www.servicenow.com/

Streamlit. (n.d.). Streamlit – Apps for the modern data stack. Retrieved from https://streamlit.io/

Tableau. (n.d.). Tableau: Visual analytics platform. Retrieved from https://www.tableau.com/

Twilio. (n.d.). Segment (by Twilio). Retrieved from https://segment.com/

U.S. Marine Corps. (2022). Marine Corps doctrinal publication 1-0: Operations. Retrieved from https://www.usmcu.edu/Portals/218/CEME/courses/MCDP%201-0%20Operations.pdf

References

Weaviate. (n.d.). Weaviate – Open-source vector search engine. Retrieved from https://weaviate.io/

Workato. (n.d.). Workato – Integration and automation platform. Retrieved from https://www.workato.com/

Zapier. (n.d.). Zapier – Automation made easy. Retrieved from https://zapier.com/