

Decision Intelligence Framework: Next-Generation Analytics Beyond Classic BI

1stAjai Gopal Bhartariya
*Research Scholar, Amity Institute
of Information Technology,
Amity University*
Uttar Pradesh, Lucknow, India
ajai.bhartariya@s.amity.edu

2nd Dr. S.K. Singh
*Professor, Amity Institute of
Information Technology
Amity University*
Uttar Pradesh, Lucknow,
India sksingh1@amity.edu

3rd Dr. Ajay Kumar Bharti
*Professor,
Ambalika Institute of Management
and Technology.*
Lucknow, India
ajay_bharti@hotmail.com

ABSTRACT

In modern business organisations a crossroad has been reached in so far as the use of data is concerned. Over the past few decades, Business Intelligence (BI) has become predominant in the organisational decision support using the retrospective reporting and stagnant dashboards. The explosion in data and the increased pace of the market, however, has revealed the constraining nature of this strategy, leading to dashboard fatigue, the creation of bottlenecks in the analysts, and delayed realisation of value. Such issues have propelled the introduction of the new paradigm of Decision Intelligence (DI), a new way of paradigm to re-architect the relationship between information and action by moving past descriptive insights to prescriptive and autonomous decision-making.

Decision Intelligence is a combination of human intelligence and Artificial Intelligence, decision modelling, causal learning, and advanced analytics. This paper identifies the changing DI ecosystem, including architectural changes to headless BI and semantic layers, the democratisation of analytics based on generative and multimodal interfaces, and the use of control-theoretic models, like the OODA loop, to enable autonomous decision making.

The case studies of industries presented by the research based on Unilever, Mastercard, and Diageo show the real economic performance of DI in the fields of optimisation of inventory, fraud detection, and marketing efficacy. Ethical issues are also discussed in the paper, with a focus on how Explainable AI can be used to reduce the problem of algorithmic bias and promote trust.

The results suggest that the organisations of the next generation will not be defined by the amount of data that they have, but the intelligence and promptness built in their decision structures.

Keywords: BI, Augmented Analytics, Decision Intelligence, NLP, Explainable AI, XAI, LLM

I. INTRODUCTION: BI BEYOND DATA PIPELINE

Over decades, the value proposition of data analytics has been dominated by the effectiveness of the data pipe - the Extract, Transform, and Load (ETL) operations that transferred data between the silos of operational activities to central repositories. Although these processes are fundamental, the contemporary enterprise has uncovered the fact that an effective ETL pipeline is not enough to ensure high results. “Business Intelligence (BI) focuses on describing what has happened — using dashboards, KPIs and reports to present historical data. Decision Intelligence (DI) builds on this by connecting those insights to decision-making processes.” [1].

The first is the transformation of data into insights, which is not the easiest one, but the last mile of analytics, which is the conversion of insight into action, is the hardest. Although classic BI can bring out a trend, it hardly dictates a course of action. “DI addresses this by embedding strategic objectives directly into decision-making workflows, ensuring consistency and alignment across teams” [1]. This change is important since, the cost of poor or delayed decision making is high; according to studies, “the average S&P 500 company throws away about 250 million dollars per year through weak decision processes” [2]. Moving towards proactive and independent action is one of the ways forward in entering the information and value gap such that data is not necessarily seen, but acted upon, at the rate of the business.

II. BACKGROUND, RELATED WORK AND GAP ANALYSIS

Modern business analytics can be traced back to the early 1990s which saw the formalisation of data warehousing in addition to the introduction of Online Analytical Processing (OLAP). “The world of data warehouses as we know it can be traced back to 1993 when Dr E.F. Codd conceived the idea of OLAP” [3]. This historic innovation made a separation between transactional processing (OLTP) and

analytical processing which enabled organisations to accumulate huge amounts of historical information to report. The technology stack has changed through the decades since the early on-premises, row-based databases, to the columnar storage engine, “and now we see the market behind to morph into lakehouses” [3]. However, regardless of the infrastructural improvements, the basic Business Intelligence (BI) workflow has been markedly and noticeably unchanged: a business user poses a question or a set of questions, “once those are all ironed out the analyst will then use their BI tool to run the SQL queries against the data warehouse and” [3] the result is presented as a static report or dashboard.

The Retrospective Gap and the Dashboard Fatigue Phenomenon:

Classic BI is diagnostic and descriptive in nature. “Business intelligence focuses on describing what’s happened through dashboards, reports and performance metrics, decision intelligence is focused on what to do next. BI might tell you that sales have dropped in one region” [4]. On the one hand, this retrospective orientation is essential to keep an eye on performance, but it creates a serious decision latency. The information that is given through BI is in essence a snapshot of the past and it is upon the decision maker to close the gap between the historical information and the action to be taken. A dashboard can be seen as a graphic depiction of a drop in sales within the region but it would be quiet on the best intervention that could be done to turn the situation around or the likely outcomes of possible future action.

This dependency on human synthesis has caused a crisis of dashboard fatigue. In a bid to be comprehensive, organisations have flooded users with useless noise by creating a multiplicity of dashboards that visualise all the possible metrics. A case in point is of a SaaS company that had a product-usage dashboard, which had 23 different charts. User testing revealed their customer success team only used a few metrics. The other 15+ visualizations in their BI dashboards were ignored. As the dashboard has been remodeled to show only these three core measures, adoption has taken off by 18 percent to 71 percent. This highlights one of the most important weaknesses of traditional BI, it is a belief that having higher amounts of data leads to improved decisions when in reality this may lead to cognitive overload and paralysis in decision-making.

The Capacity Crisis and the "Tribal Knowledge" Trap:

The human intermediaries are deeply engaged in the canonical business-intelligence (BI) workflow. Business users who do not have the technical skills to access databases directly have to depend on data teams to produce reports. This dependence has caused a capacity crisis whereby an average BI analyst is handling much more than their optimal workload that result into bottlenecks that postpone the accessibility of valuable insights. Analysts thus become in effect factories of reports, “spending their time on

ad-hoc reporting as opposed to doing high-value strategic analysis” [5].

“Decision intelligence is the process of capturing decision points, applying data and human judgment to these decision points, and monitoring the impact of these decisions” [3]. In addition, the process of interpreting BI reports is often subjective and is based on the tribal knowledge of experienced managers instead of formalised decision logic. The intuitive sense of how to read a certain KPI spike is likely to follow an experienced manager when he/she retires or leaves the company. Classic BI does not have a system that can harness and digitise this decision logic. It displays the data, but not the rules or heuristics that are used to do something with it. Such lack of decision memory makes organisations susceptible to inconsistency and subjectivity in the face of a report they can interpret differently in the eyes of disparate managers who will apply their own heuristics when it comes to interpreting a certain report.

Structural Rigidities and Data Fragmentation:

The constraints of the classical BI are structural as well. Traditional relational databases and business information systems often have strict schema constraints, which require a large amount of planning and downtime to add new fields or data types. Within the modern organisation, the data is scattered in a myriad of different systems: the enterprise resource planning (ERP) system of financials, the customer relationship management (CRM) system of customer relations, the warehouse management systems (WMS) system of inventory, and the transportation management systems (TMS) system of logistics.

Classic BI finds it hard to bring together these divergent sources into a single storey at real time. Provided by the snapshot quality of the majority of the BI reporting, it is impossible not to extract data at a certain time (e.g., nightly batches) which means that the insights will always be out-of-date by the time they will reach the decision-maker. This latency makes BI incapable of delivering the real-time, constantly changing intelligence necessary to respond swiftly in high velocity domains, “like detecting fraud, or algorithmic trading, in which the time-span of useful decision-making falls in milliseconds” [6].

Table 1 shows the key differences between classical BI and Decision Intelligence (DI).

III. DECISION ANALYTICS FRAMEWORK: THE RISE OF SEMANTIC LAYER AND HEADLESS BI

The modern data stack is experiencing an architectural change in order to make the gap between raw information and decision-ready intelligence more narrow. Monolithic BI suites of old, where the three data modelling, logic, and visualisation were combined together into one proprietary system are giving way to a more modular architecture based

around the Semantic Layer and the Headless BI. This is not only a technical change, it is a strategic requirement to make Decision Intelligence possible so that all downstream consumers, be they human analysts, AI agents or automated applications may be operating off a Single Source of Truth across the system.

The Architecture of Consistency: The Unified Semantic Layer:

The business logic in the classical setting is usually replicated and spread out through a wide range of reports and tools. A definition of a Gross Margin may be computed differently in a Tableau dashboard utilised by Sales as compared to a Power BI report utilised by Finance and this will lead to a “semantic mismatch and internal narrative conflict” [7]. This is addressed by the unified semantic layer.

This layer then transforms the complex technical database fields and names as normalised and business friendly terms (e.g. total revenue, active customer). Centralisation of such definitions by organisations makes sure that all the stakeholders in the organisation are talking in one language. The semantic layer is not just data labelling, it adds metadata, relations, and context to the data and creates a knowledge graph about the enterprise relations. Such a universal translator capability is essential to Decision Intelligence because it provides the credible base on which “prediction models and autonomous agents may act without fear of absorbing misinterpreted data” [8].

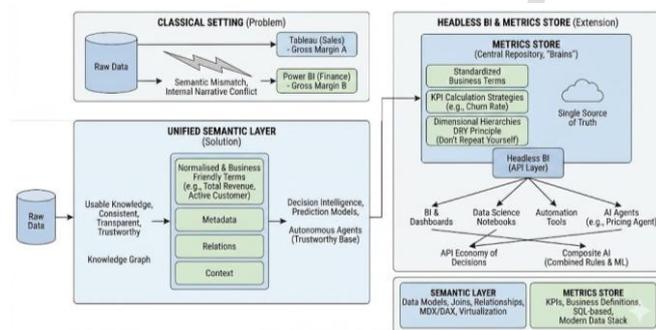


Image 1: Architecture of Consistency

Headless BI and the Metrics Store:

Headless BI extends the semantic layer one step by separating the ability to define metrics unnecessarily with the visualisation interface. “Headless BI is the architectural approach to separating the generation of business metrics from the user interface” [9]. “The team that owns the metrics only needs to define them once with headless BI, and they can use them in many different ways” [10]. In a headless design the metrics are represented as code within a central repository, which is sometimes called a Metrics Store, and which is accessed by any other consuming application.

The Metrics store is the brains of the architecture and it contains enterprise-wide standardised business terms, KPI calculation strategies and dimensional hierarchies. This methodology follows the principle of software engineering of Don't Repeat Yourself (DRY). As opposed to the SQL

query being rewritten every time an analyst needs the Churn Rate dashboard, the definition is contained in the metrics store and is used everywhere. “Metrics stores separate metrics definitions from BI reporting and visualizations... creating a single source of truth” [11].

There is a slight difference between a semantic layer and a metrics store, however:

- **Semantic Layer:** The one concerned with the overall mapping of data models and joins and relationships. It frequently provides support of complex query languages like MDX or DAX and can use virtualization.
- **Metrics Store:** Is specifically concerned with the business definition of key performance indicators (KPIs) and metrics. It is generally SQL-based and is used to provide the modern data stack, such as data science notebooks and automation tools.

Enabling the "API Economy" of Decisions:

The main importance of Headless BI in a Decision Intelligence environment is that it allows supporting the API economy of information. Since metrics are made available through open APIs, they are not only consumed by dashboards, but also by operational workflows and AI agents. A pricing agent (an AI agent) is able to ask the Metrics Store the current inventory level and the price index of its competitors in real-time and make a pricing decision, without requiring knowledge of the underlying database schema.

This architecture can handle the concept of the Composite AI in which specific AI methods are combined (e.g., rules-based logic and machine learning). The semantic layer is a guarantee that these scattered elements have a shared definition of the entities of data and it makes it easy to coordinate complicated decision processes. In the natural sciences, e.g., semantic technologies are becoming essential to synthesise fragmented information within the fields of research, clinical and regulatory, to generate hypotheses more quickly and adopt AI more securely.

Table 2 shows different components of DI.

IV. AI AND AUGMENTED ANALYTICS: DEMOCRATIZING INSIGHT

The introduction of Artificial Intelligence (AI) into the process of analysis has given rise to the vision of Augmented Analytics, which represents a paradigm of utilising Machine Learning (ML) and “Natural Language Processing (NLP) to automatize the labour-intensive part of the data analysis process” [12], i.e., preparation, discovery, and visualisation. It is this evolution that is used to democratize access to advanced insights, allowing the business person to bypass the technical bottlenecks that exist under the traditional Business Intelligence (BI) system

and allows the business person to interact with the data directly.

Generative BI: The Conversational Interface:

Generative Business Intelligence (GenBI) is a combination of Large Language Models (LLMs) and business intelligence. In contrast to the more traditional natural language query (NLQ) systems, which often tend to be brittle and reliant on specific keywords, GenBI makes use of the semantic understanding of LLMs in understanding the free or nuanced natural language queries. Instead of classical dashboards, Gen BI uses AI (LLMs, NLP) and enables users to ask data questions in plain English. Google Looker, which is built-in with Gemini, and Amazon QuickSight Q allow asking questions like, What was the cause of slumping sales in Q3.

The effect of GenBI on the speed of decision is significant. "About 65% of companies report improved decision quality with AI-augmented BI" [13]. GenBI lowers the barrier to entry into analytics by automatically formulating SQL statements and selecting the right visualisation types, thus enabling "non-analysts [to] chat with data in real time" [13].

Nonetheless, GenBI does not replace regular BI systems, but complements them instead. It is also good at answering ad hoc and exploratory questions, which frees data analysts to focus on strategic, high-value projects. It lays a level of support to data queries, "whereby AI serves the daily questions that used to clog the queue of analysts" [14].

Agentic AI: From Analysis to Autonomy

Whereas Generative BI assists human users, the issue of Agentic AI is with respect to taking actions in their behalf. According to Gartner, in 2027, "AI agents will complement or automate half of business decisions" [15]. The difference between these agents and the passive models is their ability to pursue ends, use tools, perform tasks, and create workflows independently.

In a Decision Intelligence system, the AI agent can have a high-level goal, such as minimising holding costs by 10 per cent without service quality declining. The agent would then:

1. **Query the Semantic Layer:** In order to determine existing inventory, sales velocity, and lead time.
2. **Run Predictive Models:** To "predict demand and possible disruptions in the supply chain" [16].
3. **Evaluate Scenarios:** It can use prescriptive logic to find the best reorder points.
4. **Execute Actions:** The purchase orders can be generated in the ERP system or "can be sent to a human planner to be approved" [17].

Such a shift to the so-called Agentic Analytics is made possible by the ability of the agents to link a series of procedural tasks, "such as data cleaning, analysis, reasoning, and execution, into a unified workflow" [18]. As an example, data-quality problems (including missing values) can be automatically identified by the agents, and the remedial policies (including imputation) can be implemented to make sure that informed decisions are made based on sound data.

Automated Insight Discovery:

Also included in the scope of augmented analytics are so-called smart discovery, which are algorithms that systematically search through datasets in order to generate statistically significant patterns, anomalies, or relationships that could be missed by human analysts, "and uncover patterns free from human bias" [19]. These unasked questions are critical to the disclosure of latent risks and opportunities. Indicatively, an augmented analytics engine can help a sales manager to know with early warning that a specific group of customers is at a risk of churning out, before it can affect the top-line revenue due to a subtle change in usage habits, which is not yet reflected.

This feature is changing the paradigm of analytical models that were traditionally pull-based, meaning the user enters in a question, to push-based, meaning the system provides insight automatically. It is expected that by 2025, contextual understanding received actively will be a common feature of advanced analytics systems and will be supported by persistent intelligence pipelines, which will act on real-time data streams.

V. VISUALIZATION AND INTERACTIVE REPRESENTATION: BEYOND THE SCREEN

The traditional method of dashboard (two-dimensional) becomes less effective to display multidimensional data because of its size and limited dimensions. Developers of Decision Intelligence introduce innovative ways of interaction, which take into account the entire range of human senses, such as voice, audio, and multimodal interfaces. The purpose of these advances is to alleviate cognitive load and expansion analytics into the unfeasible screen settings.

Voice-Driven Analytics: The Hands-Free Revolution

Voice based analytics uses advanced Speech-to-Text (STT) and Natural Language Understanding (NLU) systems to provide verbal interface to data systems. The use of this technology is particularly revolutionary to the deskless workers in manufacturing, logistics and healthcare sectors who need hands-free access to information.

The benefits of voice interfaces are quantifiable:

International Conference on Intelligent Technologies & Science - 2025 (ICITS-2025)

- **Speed:** Voice-driven querying deployments “can improve decision-making speed by 78%” [20] compared to the traditional manual lookups.
- **Accessibility:** Voice commands break the hurdle to non-technical users, allowing non-technical staff to make use of data and BI.
- **Industrial Utility:** Specialised voice AI models that have been trained on domain-specific acoustics can be very accurate in noisy or hazardous environments where touchscreens are unsafe or ineffective.

As an example, a factory floor manager can ask a question verbally, e.g. what does Unit 4 have as the current operating temperature? or "Show the statistics of the previous shift of the downtime," immediately getting audio-visual feedback without interrupting the working process.

Multimodal Interfaces and Context Awareness

The potential future development of analytic interfaces involve a multimodal combination of textual, visual, acoustic, video, and sensor-based data to a single, unified analytic stream. It is expected that, by the year 2030, the multimodal artificial intelligence market that was valued at \$1.73 billion in 2024 is projected to increase by 10x.

Multimodal Decision Intelligence system can simultaneously handle diverse inputs to provide contextual support. For example, “a single AI agent might interpret a customer’s voice query, review an image of a product, and generate a helpful video or text response—all in real-time” [21].

This feature is the foundation of the notion of the Ambient Intelligence system which implies that the system smoothly becomes part of the environment and provides the user with small-step improvement of his/her working process without the need to demand attention in any way. “The ambient intelligence pattern integrates suggestions subtly into the workflow, appearing only when context makes them valuable” [22].

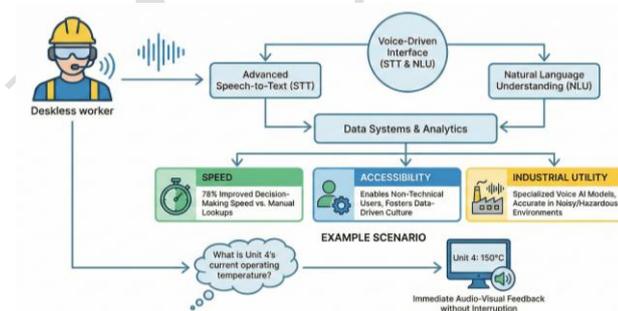


Image 2: Voice-Driven Analytics

Table 3 also shows different modalities and advancements in BI.

VI. VISUALIZATION AND INTERACTIVE REPRESENTATION: BEYOND THE SCREEN

Decision Intelligence is characterised by the fact that it is shifting to prescriptive analytics, as opposed to descriptive and predictive analytics. Unlike predictive models, which predict phenomena (e.g., “Demand will decline by 10%”), “decision intelligence—prescribes (or, more accurately, “recommends”) a set of actions for achieving a defined set of goals” [23] (e.g., “Offer a 5% discount to inventory to maintain revenue neutrality). To achieve this jump, there must be a sophisticated combination of causal thinking, decision modeling that is standardized, and a system of learning that is lifelong.

Causal Inference: Beyond Correlation

One of the issues that have plagued most machine-learning deployments is the confusion of correlation and causation. “You can use causal inference to measure the business impact of deploying a recommendation system or any other AI tool” [24]. For example, “causal frameworks can reduce marketing costs by 30-40% while improving ROI” [25].

Causal Artificial Intelligence applies a scientific approach to analytics, providing organisations with the opportunity to model interventions and assess their real effects. Key concepts include:

- **Counterfactual Reasoning:** The question would be, What would have otherwise happened to this particular customer, had we not given the coupon? “This allows calculating individual treatment effect (ITE) and not just aggregated averages” [24].
- **Uplift Modeling:** “Uplift models directly estimate the **individual treatment effect (ITE)**” [24]. When you are interested in changing the behaviour of a subset of the population, namely the persuadables.
- **Confounder Detection:** Identification of extraneous factors (e.g., a competitors reduction of prices) that can cause bias in the perceived association between an action and its resultant outcome.

With causal models, enterprises are able to optimise decision policies to achieve incremental value. It has been empirically demonstrated that causal schemes have the ability of reducing marketing investments by 3040 percent by removing redundant incentives on those customers who would have converted otherwise in the absence of the intervention.

Engineering Decisions: The DMN Standard

In order to operationalize prescriptive analytics, decisions should be perceived as designed pieces of art as opposed to thought processes. “DMN creates a standardized bridge for the gap between the business decision design and decision implementation” [26]. DMN allows business analysts to

International Conference on Intelligent Technologies & Science - 2025 (ICITS-2025)

graphically represent the rules and inputs and logic that make up a decision, to decouple decision logic and application code.

A DMN model typically consists of:

- **Decision Requirements Diagram (DRD):** This is a high-level schematic diagram that shows the relationship between different decisions and data sources.
- **Decision Tables:** Tabular representations of business rules (e.g., "IF Credit Score > 700 AND Income > \$50k THEN Approve").
- **Boxed Expressions:** Standardized expressions that make calculating and rationally reusing easier, "treating decisions as reusable assets rather than ad-hoc processes" [26].

The result of this standardisation is a set of Decision Services, "stateless, side-effect free elements that may be called to make a consistent decision by any system" [27] (BPMN processes, AI agents, microservices). As a result, the decision logic is the same and transparent regardless of whether the loan application is done via a web interface, via a mobile application, or via a call-centre representative.

Image 3 depicts a decision engineering flow using DMN Standard.

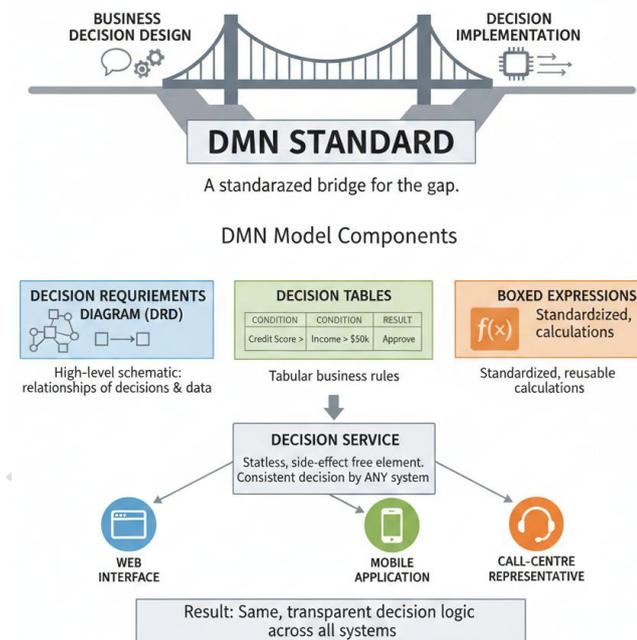


Image 3: Engineering Decisions: The DMN Standard

Strategic Frameworks: OODA and SEAL

On a strategy level, Decision Intelligence systems are often modelled as control loops that guarantee continuous adaptation. The most notable one of them is the OODA Loop (Observe, Orient, Decide, Act), which was initially

Organized by: International Journal of Research and Development in Applied Science and Engineering, India
All Rights Reserved © 2025 IJRDAE (ISSN: 2454-6844)

"articulated by former Air Force fighter pilot John Boyd" [28].

- **Observe:** The system is linked to data streams to identify events (e.g., a disruption in the supply chain).
- **Orient:** The system builds the context based on semantic layers and predictive models (e.g., estimating the effect of the disruption).
- **Decide:** The system uses DMN logic or optimization solvers to choose the best course of action (e.g., re-shipment).
- **Act:** The system would implement the decision through an API and track the result.

Emergent theories like SEAL (Sense, Evaluate, Act, Learn) predict the collaborative aspect of human machine decision-making, which reduces the cognitive load on the human operators but maintains a human-in-the-loop which handles critical checkpoints.

Image 4 below depicts the OODA Loop and SEAL Strategic Decision Frameworks.

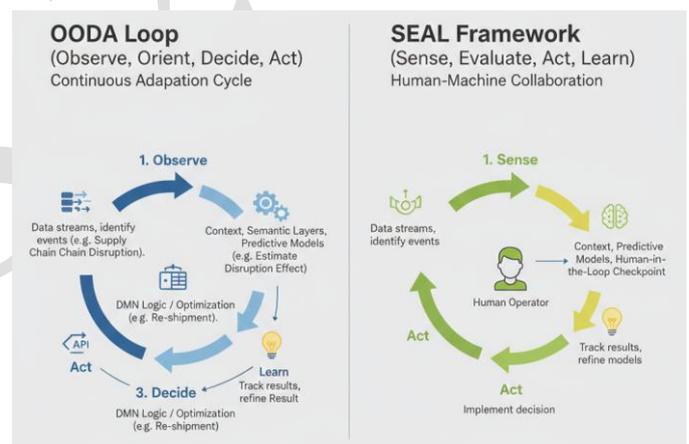


Image 4: OODA Loop and SEAL Framework

The Maturity Model: From Support to Automation

Introduction of Decision Intelligence is a maturity process. There are usually three levels of progress of organisations:

Decision Support: The system offers AI generated insights and dashboards, but the final decision is made by a human (e.g., a dashboard with sales trends).

Decision Augmentation: The system suggests a particular decision and man will confirm it (e.g., "We recommend to reorder 500 units. Approve?").

Decision Automation: It is a system that is capable of making and implementing decisions without close supervision.

The long-term goal is to relocate high-volume, low-risk decisions to the Automation level, thus freeing up human

International Conference on Intelligent Technologies & Science - 2025 (ICITS-2025)

space in the, high-value, complex, and strategic, deliberations.

VII. CHALLENGES: ETHICS, BIAS, AND EXPLAINABILITY

The risks posed by opaque (black box) decision-making processes in the modern organisation have been increased by the transfer of power to independent systems. The risk of algorithmic bias, the complexity of deep-learning models, and the problem of data privacy altogether imply the need to develop a strict ethical code and the systematic use of Explainable Artificial Intelligence (XAI).

Algorithmic Bias and Fairness

The definition of algorithmic bias describes a decision-making system where “decisions made by AI systems may reflect prejudices present in the training data” [29]. Financial market, as an example, the use of automated lending programmes has been discovered to discriminate against some population groups arbitrarily in favour of others, resulting in what is popularly known as automated discrimination.

In order to address these risks, fairness-conscious machine-learning methods should be implemented on decision-intelligence (DI) platforms. This involves:

- **Pre-processing:** Considering balanced training sets so that there is a diversity of representations.
- **In-processing:** Directly embedding fairness constraints within the objective of the model (e.g. having the model pay fines due to disparate impact).
- **Post-processing:** To make model outputs equitable among the groups of protection. Auditing these systems should be done by the ethical governance bodies or by the so-called AI Review Boards, which include a variety of stakeholders (legal, compliance, sociology, technology).

The Trust Gap and Explainable AI (XAI)

The currency of AI adoption is trust. Unless the business leaders understand the logic behind such decisions, they would hardly grant an AI agent the power to make multi-million-dollar supply-chain decisions. XAI acts as the trust layer between the user and platform.

The design patterns to effective XAI to the business stakeholders depend on the following design patterns:

- **Narrative Generation:** The system comes up with a natural-language description, instead of showing direct feature-importance scores (e.g., FICO Score: 0.4 weight): “this customer's high-risk score is driven by 40% recent missed payments” [30].

- **Counterfactual Explanations:** These offer alternative action based on the ability to show what would have changed the outcome. Had you cut down on the usage of your credit by 10 percent, your application would have been approved. This is essential to conform to the regulations in areas like banking (GDPR/ECOA).
- **Confidence Indicators:** It is explicitly stating the degree of confidence in the model (e.g., 87% confidence match) and helps users to adjust their trust to the automated suggestion.

Data Privacy and Governance

Emerging vectors of data leakage are presented by the application of generative AI into DI platforms. It is also possible that any discussion of proprietary financial data using public large language models (LLMs) can reveal sensitive data unless adequately controlled. The architectures used by organisations must also use private and local LLM or powerful data-masking protocols. In addition, the use of synthetic data to train models should be handled with care; synthetic data failed to generate successful results can trigger a model collapse or regulatory failure. Automated lineage tracking that is recording the exact sources of data and logic versions used to make a particular decision is an essential condition of auditable DI.

VIII. CONCLUSION

The shift from Business Intelligence to Decision Intelligence marks a definitive maturation in the enterprise data capability. For thirty years, the industry has focused on the "how" of data—how to store it, how to query it, and how to visualize it. Decision Intelligence shifts the focus to the "why" and the "what next." By bridging the "retrospective gap" of classic BI with prescriptive logic, and by solving the "capacity crisis" through AI-driven automation, DI offers a path out of dashboard fatigue and into the era of the autonomous enterprise.

This transformation rests on three pillars:

- **Architectural Modernization:** The adoption of Headless BI and Unified Semantic Layers to create a programmable, consistent fabric of business logic.
- **Cognitive Augmentation:** The use of GenBI, Voice, and Multimodal interfaces to make data accessible, intuitive, and omnipresent.
- **Decision Engineering:** The formalization of decision logic through standards like DMN and causal AI, moving decision-making from an art form to a rigorous engineering discipline.

The evidence from early adopters like Unilever and Mastercard confirms that this is not merely a theoretical upgrade but a driver of massive economic value—reducing

International Conference on Intelligent Technologies & Science - 2025 (ICITS-2025)

waste, stopping fraud, and optimizing growth. However, the path forward is not solely technological. It requires a commitment to ethical AI, the cultivation of trust through explainability, and the courage to redesign organizational processes around the speed of algorithms rather than the

speed of meetings. As we approach 2027, the divide between market leaders and laggards will be defined by their "Decision Velocity"—the speed and accuracy with which they can translate the noise of data into the signal of action.

Table 1: Traditional BI VS Decision Intelligence

Feature	Classic Business Intelligence (BI)	Decision Intelligence (DI)	Implications for the Enterprise
Primary Orientation	Descriptive & Diagnostic ("What happened?")	Predictive & Prescriptive ("What to do?")	Shifts focus from observation to action; reduces decision latency.
Cognitive Load	High: Human must interpret and synthesize.	Low: System recommends or automates.	Reduces "dashboard fatigue" and analysis paralysis.
Workflow	Linear: Data → Analyst → Report.	Closed-Loop: Data → Decision → Action → Feedback.	Enables continuous learning and optimization of decision logic.
Logic Storage	Mental: Logic exists in users' heads (Tribal Knowledge).	Digital: Logic encoded in models/rules (Metrics Store).	institutionalizes expertise; reduces key-person risk.
Latency	Batch-based (Snapshot).	Real-time / Event-driven.	Allows for immediate response to emerging threats or opportunities.

Table 2: Components of DI

Component	Function	Role in Decision Intelligence
Data Warehouse / Lake	Raw Storage	The "Memory": Stores the historical and real-time data.
Semantic Layer	Context & Mapping	The "Dictionary": Translates technical data to business concepts.
Metrics Store	Logic & Calculation	The "Brain": Centralizes the rules for calculating KPIs.
Headless API	Distribution	The "Nervous System": Delivers metrics to humans and machines.

International Conference on Intelligent Technologies & Science - 2025 (ICITS-2025)

Downstream Apps	Consumption	The "Limbs": Dashboards, AI Agents, and Automation scripts.
------------------------	-------------	---

Table 3. Different Modalities and Advancements in BI

Interface Modality	Key Technology	Primary Use Case	Advantage
Generative Chat	LLM / NLP	Ad-hoc exploration / Executive reporting	Instant SQL generation; explains "why"
Voice Command	NLU / Speech-to-Text	Field operations / Logistics / Healthcare	Hands-free; 30-78% faster decisions
Sonification	Audio Mapping	Security monitoring / High-dim analysis	Reduces visual cognitive load; peripheral awareness
Multimodal	Sensor Fusion / Vision	Claims processing / Predictive maintenance	Holistically analyzes unstructured & structured data

REFERENCES

- [1] Ultimate Guide to Decision Intelligence (DI) - Qualtrics, <https://www.qualtrics.com/articles/strategy-research/decision-intelligence/>
- [2] How a Decision Intelligence Framework Boosts Enterprise Efficiency - SAP BW Consulting, <https://www.sapbwconsulting.com/blog/decision-intelligence-framework>
- [3] A Data Warehouse Alone Won't Cut It for Modern Analytics - DZone, <https://dzone.com/articles/a-data-warehouse-alone-wont-cut-it-for-modern-anal>
- [4] Evaluating decision intelligence tools vs BI solutions - ThoughtSpot, <https://www.thoughtspot.com/data-trends/business-intelligence/decision-intelligence-tools>
- [5] David, Yigal & Harison, Elad. (2022). Paying Lip Service?. International Journal of Enterprise Information Systems. 18. 10.4018/IJEIS.313049. https://www.researchgate.net/publication/364947110_Paying_Lip_Service
- [6] Decision intelligence vs business intelligence: Which to use? - Linkurious, <https://linkurious.com/decision-intelligence-vs-business-intelligence/>
- [7] How a Decision Intelligence Framework Boosts Enterprise Efficiency - SAP BW Consulting, <https://www.sapbwconsulting.com/blog/decision-intelligence-framework>
- [8] Team O. Gartner: semantic technologies take center stage in 2025 powering AI, metadata, and decision intelligence [Internet]. Gartner. 2025. Available from: <https://www.ontoforce.com/blog/gartner-semantic-technologies-take-center-stage-in-2025>

International Conference on Intelligent Technologies & Science - 2025 (ICITS-2025)

- [9] Aslett M. Headless BI and Data Democratization Enables AI and Data-Driven Enterprises [Internet]. ISG. 2025. Available from: <https://research.isg-one.com/analyst-perspectives/headless-bi-and-data-democratization-enables-ai-and-data-driven-enterprises>
- [10] Kyligence. What is Headless BI and Why It Matters? [Internet]. Kyligence. 2023. Available from: <https://kyligence.io/blog/what-is-headless-bi-and-why-it-matters/>
- [11] Kyligence. Metrics Store vs Semantics Layer [Internet]. Kyligence. 2023. Available from: <https://kyligence.io/blog/metrics-store-vs-semantics-layer/>
- [12] Augmented analytics for democratizing data insights. IARJSET [Internet]. 2025 Apr 26;12(4). Available from: <https://doi.org/10.17148/iarjset.2025.12451>
- [13] Ab. The rise of Generative BI Platforms: What it means for data Teams [Internet]. Autonmis. 2025. Available from: <https://autonmis.com/blogs/the-rise-of-generative-bi-platforms>
- [14] Lewinson E. Metrics Layer: a single source of truth for all KPI definitions [Internet]. Towards Data Science. 2025. Available from: <https://towardsdatascience.com/metrics-layer-a-single-source-of-truth-for-all-kpi-definitions-671c9aa04a91/>
- [15] Gartner announces the top data & analytics predictions [Internet]. 2025 [cited 2026 Feb 5]. Available from: <https://www.gartner.com/en/newsroom/press-releases/2025-06-17-gartner-announces-top-data-and-analytics-predictions>
- [16] Aera Technology. AeRa Technology [Internet]. Aera Technology. Available from: <https://www.aeratechnology.com/decision-intelligence-faq>
- [17] Durak N. How Technology Makes Continuous Innovation Possible: A Case Study with Unilever [Internet]. Logistics Viewpoints. 2023. Available from: <https://logisticsviewpoints.com/2023/10/05/how-technology-makes-continuous-innovation-possible-a-case-study-with-unilever/>
- [18] Tencent. Next-Gen Data Intelligence Platform Construction Guide: Data+AI Enabled [Internet]. 2025. Available from: <https://staticintl.cloudcachetci.com/yehe/backend-news/06.26-%E8%8B%B1%E6%96%87%E7%89%88-Data%2BA1%EF%BC%8C%E4%B8%8B%E4%B8%80%E4%BB%A3%E6%95%B0%E6%99%BA%E5%B9%B3%E5%8F%B0%E5%BB%BA%E8%AE%BE%E6%8C%87%E5%8D%97%E6%89%8B%E5%86%8C-%E5%8D%95%E9%A1%B5%E7%89%88.pdf>
- [19] Siraj F. Ansari, Srujan K. Gunta. Augmented analytics for democratizing data insights. IARJSET [Internet]. 2025 Apr 26;12(4). <https://doi.org/10.17148/iarjset.2025.12451>
- [20] How AI improves campaign tracking accuracy [Internet]. Available from: <https://www.enreach.ai/insights/how-ai-improves-campaign-tracking-accuracy>
- [21] Team C. Multimodal AI in 2025: Transforming communication and the road ahead for platforms like ChatMaxima [Internet]. Available from: <https://chatmaxima.com/blog/multimodal-ai-in-2025-transforming-communication-and-the-road-ahead-for-platforms-like-chatmaxima/>
- [22] Guide to AI integration patterns for enterprise-scale systems | Transcenda [Internet]. Available from: <https://www.transcenda.com/insights/guide-to-ai-integration-patterns-for-enterprise-scale-systems>
- [23] Predictive Analytics vs. Prescriptive Analytics - Gurobi Optimization [Internet]. Gurobi Optimization. 2025. Available from: <https://www.gurobi.com/resources/predictive-analytics-vs-prescriptive-analytics/>
- [24] I by IMD. How 'causal' AI can improve your decision-making - I by IMD [Internet]. IMD Business School for Management and Leadership Courses. 2025. Available from: <https://www.imd.org/ibyimd/artificial-intelligence/how-causal-ai-can-improve-your-decision-making/>
- [25] Finarb. Finarb - AI & Data Solutions | Transform Your Business with Advanced Analytics [Internet]. Finarb. Available from: <https://finarb.ai/blog/causal-inference-business-decisioning-correlation-causation>
- [26] FlexRule. Decision Model and Notation - Decision Intelligence Platform [Internet]. Open Decision Intelligence Platform. 2026. Available from: <https://www.flexrule.com/decision-model-and-notation/>
- [27] Wikipedia contributors. Decision model and notation [Internet]. Wikipedia. 2026. Available from: https://en.wikipedia.org/wiki/Decision_Model_and_Notation
- [28] OODA Loop. Decision Intelligence and Establishing an intelligent Enterprise — OODALoop [Internet]. OODALoop. 2024. Available from: <https://oodaloop.com/corporate-sensemaking-establishing-an-intelligent-enterprise/>
- [29] Rinf.Tech. Ethical Considerations in AI Decision-Making [Internet]. rinf.tech. 2024. Available from: <https://www.rinf.tech/ethical-considerations-in-ai-decision-making/>
- [30] Sinha H. AI bridges the trust gap: How explainable AI aligns modelers and business leaders. Forbes [Internet]. 2025 Feb 19; Available from: <https://www.forbes.com/councils/forbestechcouncil/2025/02/19/ai-bridges-the-trust-gap-how-explainable-ai-aligns-modelers-and-business-leaders/>