

The Emergence of Decision AI

Smarter Action with Artificial Intelligence

KEVIN PETRIE

SEPTEMBER 2022

RESEARCH SPONSORED BY



THIS PUBLICATION MAY NOT BE REPRODUCED OR DISTRIBUTED
WITHOUT ECKERSON GROUP'S PRIOR PERMISSION.

About the Author



For 25 years, **Kevin Petrie** has deciphered what technology means to practitioners, as an industry analyst, writer, instructor, marketer, and services leader. Kevin launched, built, and led a profitable data services team for EMC Pivotal in the Americas and EMEA, and ran field training at the data integration software provider Attunity (now part of Qlik). A frequent public speaker and author of two books on data streaming, Kevin also is a data management instructor at eLearningCurve.

About Eckerson Group

Eckerson Group is a global research and consulting firm that helps organizations get more value from data. Our experts think critically, write clearly, and present persuasively about data analytics. They specialize in data strategy, data architecture, self-service analytics, master data management, data governance, and data science. Organizations rely on us to demystify data and analytics and develop business-driven strategies that harness the power of data. [Learn what Eckerson Group can do for you!](#)



About This Report

Research for this report comes primarily from numerous briefings with software vendors. This report is sponsored by Diwo, who has exclusive permission to syndicate its content.

Table of Contents

Executive Summary	4
A New Chapter for Decision-Making.....	5
Drilling In	10
Getting Started	14
About Eckerson Group	16
About the Sponsor	17

Executive Summary

The emerging discipline of decision AI creates the opportunity for enterprises to make smarter, faster, and more adaptable decisions than traditional business intelligence methods. Decision AI uses artificial intelligence (AI) to automate the decision lifecycle and overcome human limitations. It comprises both decision intelligence and contextual intelligence.

- > **Decision intelligence** detects what happened and analyzes why it happened. Then it predicts what will happen, prescribes what to do, and automates the resulting decision and action. Decision intelligence uses visualization, augmented analytics, and machine learning (ML) to help make decisions, then workflow orchestration to automate the resulting action.*
- > **Contextual intelligence** helps make those decisions actionable and specific, and more so over time because it learns. It assesses outcomes, compares them to expectations, then refines algorithms to improve future decisions. Contextual intelligence uses machine learning to perform these tasks and visualization to present evidence along the way for human review.*

A decision AI initiative should start with operational use cases such as finance, sales, supply chain management, marketing, and pricing. If enterprises achieve results with operational use cases, they next can consider tackling strategic use cases such as projects related to market analysis and product design.

Decision AI initiatives contend with challenges related to system design, data management, and AI model training. To improve the odds of success, enterprises should seek to build or buy a decision AI platform that is comprehensive, rigorous, creative, rational, fair, transparent, adaptive, fast, elastic, and open. They also should start with simple, practical use cases, compare commercial and homegrown options, and assemble a cross-functional team for their decision AI initiative.

A New Chapter for Decision-Making

As long as we've stood upright, humans have followed a similar thought process to make decisions. A cave dweller that detected something 10,000 years ago—say, the track of a wildebeest—would have analyzed where the wildebeest went, then predicted where it would go, and thereby decided where to hunt. If they made a timely decision and assessed the results afterwards, they got better at finding food over time. If they had too many tracks to follow in different directions, they got confused...and hungry.

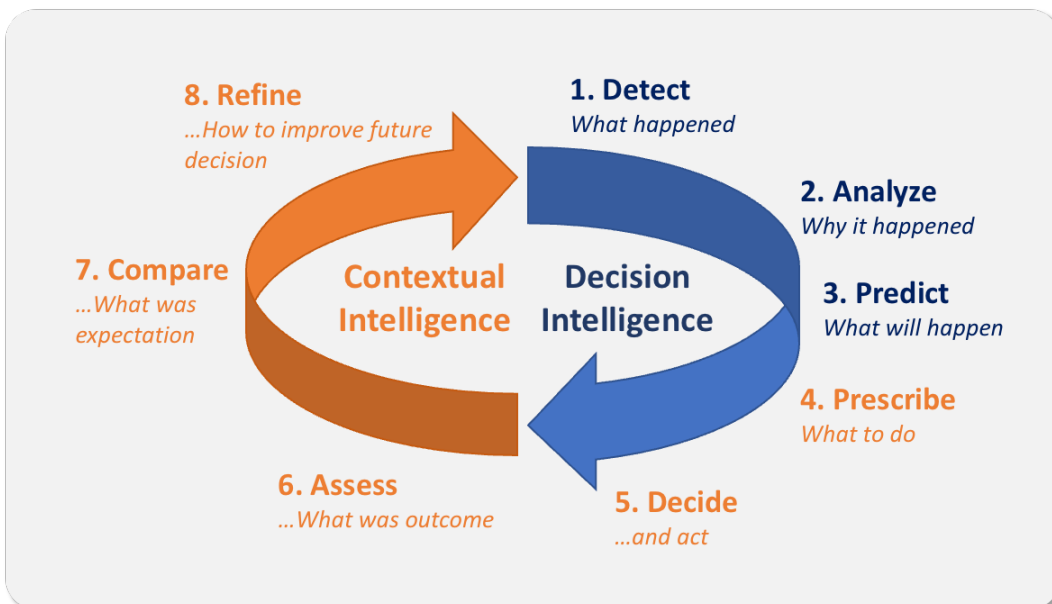
Humans face similar promise, and similar peril, when they make decisions in a modern business environment. Analysts and business managers detect trends or events, analyze them, predict the next event, and decide what to do about it. They use business intelligence (BI) tools to understand what happened, data science tools to predict what will happen, and both types of tools to understand why. But sometimes they have too many charts to synthesize, too many data points to correlate, and too little context to make the right decisions. Confused, they default to intuition and gut instinct.

This aborted process can result in bad decisions, late decisions, or no decisions. Business performance suffers.

Enter Decision AI

There might be hope. The emerging discipline of decision AI gives enterprises a new analytical option to make timely, accurate choices. Decision AI uses artificial intelligence (AI) to automate the decision lifecycle and overcome human limitations. Like other technologies, decision AI prepares data for analytics and helps users visualize analytical outputs. But its real innovation lies in the combination of decision intelligence and contextual intelligence (see figure 1.)

Figure 1. Decision AI Combines Decision and Contextual Intelligence



Decision Intelligence

Decision intelligence detects what happened and analyzes why it happened. Then it predicts what will happen, prescribes what to do, and automates the resulting decision and action. Decision intelligence uses visualization, augmented analytics, and machine learning (ML) to help make decisions, then workflow orchestration to automate the resulting action. Vendors such as Diwo, Noodle, and Aera provide tools that support decision intelligence, including recommendations.

Contextual Intelligence

Contextual intelligence helps make those decisions actionable and specific, and more so over time because it learns. It assesses outcomes, compares them to expectations, then refines algorithms to improve future decisions. Contextual intelligence uses machine learning to perform these tasks and visualizations to present evidence along the way for human review. It overlaps with decision intelligence because it also helps prescribe what to do and decide how to act. Diwo, for example, combines contextual intelligence and decision intelligence in this way.

Decision AI seeks to help humans achieve smarter, faster, and more adaptable business thinking than traditional BI reports and dashboards. Designed and implemented well, it has the potential to overcome human limitations in the following ways. Decision AI can be:

- > **Smarter.** Decision AI examines many data points and combinations, then isolates the most relevant ones. This helps replace human shortcuts and hunches with a comprehensive, evidence-based conclusion about what factors really matter.
- > **Faster.** Decision AI automates the lifecycle of analysis, decision, and action. This helps replace human hesitation and second guesses with rapid processing of the facts.
- > **More adaptable.** Decision AI automatically catalogs decisions and their outcomes, then performs a post mortem to improve future decisions. This helps replace excuses and bad habits with logical ideas to avoid repeat mistakes.

Designed and implemented well, decision AI has the potential to drive smarter, faster, and more adaptable decisions than traditional BI.

Case Study

So how might the vision play out in reality? Suppose a logistics company needs to streamline its delivery routes to support a surge in holiday shipping in New York City. Their supply chain manager reviews shipment orders and delivery schedules across eight warehouses. The decision AI platform guides them through a series of logical steps.

Decision intelligence. Alerted by the platform, this manager *detects* that one warehouse shifted its deliveries to nighttime in the borough of Queens. Prompted by popup windows, they *analyze* why this happened and identify the root cause: roadwork on a major highway that is scheduled to last one week. Based on historical disruptions, the decision AI platform *predicts* that the roadwork will take longer and snarl traffic until New Year's Day, forcing drivers to work nights through the holidays.

Now, the platform prescribes a decision: re-route shipments through a different warehouse that can reach the same neighborhoods without hitting roadwork. The manager validates this thinking with the supply chain analyst, then decides to take the recommended action. This kicks off an automated workflow that dispatches drivers to the other warehouse.

Contextual intelligence. Now things get interesting. The decision AI platform *assesses* the outcome of that decision through the holidays and *compares* it to expectations, noting that drivers deliver 90% of their packages on time without working past 7pm. It logs this success within a catalog of all relevant decisions and outcomes for review by the supply chain manager. In addition, the decision AI platform suggests ways to *refine* future decisions like these. In this case, it prompts the manager to add real-time alerting of roadwork announcements from the city. The real-time alerts give them advance warning of traffic disruptions, enabling faster and more proactive responses.

Challenges

That's the vision of decision AI. As with any new technology, this vision collides with hard challenges. This starts with system design. Business and data leaders can struggle to insert machine logic in the right places, then incorporate human review and approval at the right moments. Data teams, meanwhile, have the familiar struggles of consolidating and preparing high volumes of diverse datasets, then changing data sources as needs change. They also must train their decision AI's models to explore the most telling key performance indicators (KPIs) and charts, while minimizing hunches and dead ends.

Many of these challenges also apply to traditional BI. Data teams that overcome these challenges with decision AI have the potential to yield greater results than with BI. That is, they can make more decisions, and more accurate decisions, for a given amount of time and effort.

Decision AI faces challenges in the areas of design, data management, and ML model training.

Benefits

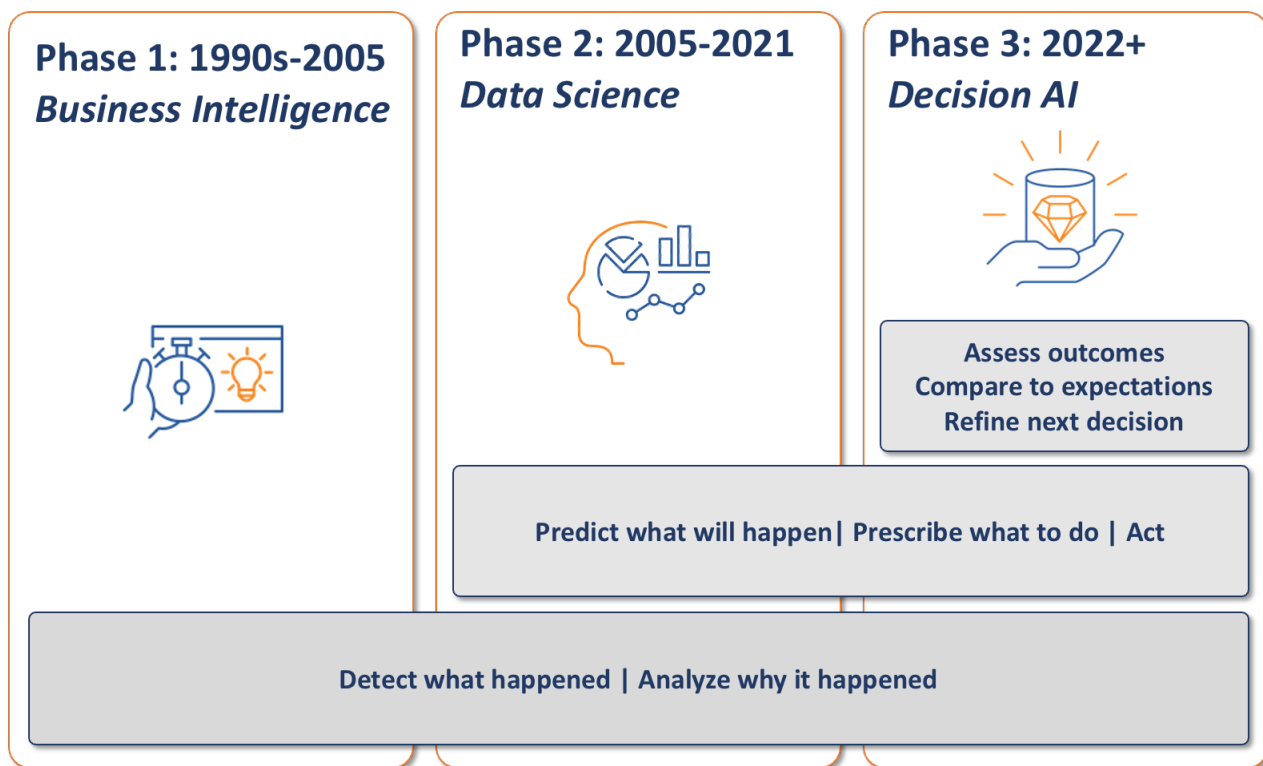
Enterprises that address these challenges well can help ensure their decision AI initiative justifies the effort and risk. They can improve efficiency as they streamline processes, shorten decision time,

and reduce mistakes. They can increase the productivity of their data consumers—including data analysts, business analysts, and business managers—as well as their data engineers. They also can foster innovation by arming decision makers with new insights and freeing up their time to address big questions. In these ways, decision AI can help enterprises reduce risk, reduce cost, and increase revenue. It can help them navigate the hyper competitive markets, geopolitical instability, and economic uncertainty of the 2020s.

Market Evolution

Decision AI represents the third phase of an analytics market that evolved over decades. The three phases include business intelligence, data science, and decision AI. Each cumulative phase complements its predecessor rather than replacing it. (See figure 2.)

Figure 2. Market Evolution



Phase 1: Business Intelligence

The discipline of BI gained traction in the 1990s to help decision makers analyze company performance and business trends. Vendors such as IBM, Microsoft, and Oracle offered BI tools that generate periodic reports and dashboards based on structured data within traditional data warehouses.

Business analysts and data analysts use BI reports and dashboards to detect and visualize what happened based on KPIs, trends, anomalies, and alerts. They also use BI tools to analyze why things happened by looking at root causes, correlations, and patterns. Recent innovations such as augmented analytics make this detection and analysis easier by using AI/ML to prepare data and generate insights.

Phase 2: Data Science

The discipline of data science gained traction in the early 2000s to help decision makers analyze future events and trends that shape business outcomes. An ecosystem of commercial and open-source data science tools arose, including libraries such as PyTorch or TensorFlow and notebooks such as Jupyter. They layer on top of new data lakes such as Delta Lake, which process high volumes and varieties of multi-structured data to support real-time or periodic use cases.

Data scientists use such tools to predict what will happen and evaluate the opportunity or risk. They apply AI/ML techniques such as regression, classification, and clustering to study multi-structured data and estimate the likelihood of successful future outcomes. They train AI/ML models to optimize predictions, simulate outcomes, and then prescribe what to do. In addition, developers build applications and workflows that operationalize the resulting action. As part of this trend, decision intelligence arose in recent years to improve operational decisions by automating the lifecycle of BI and data science tasks.

Phase 3: Decision AI

The discipline of decision AI is emerging now to help decision-makers improve the lifecycle on an automated basis. Decision AI tools such as Diwo—or homegrown alternatives— complement BI tools such as Qlik or Tableau and layer onto cloud data platforms such as Snowflake or Databricks. Decision AI processes high volumes of structured and semi-structured data to drive operational decisions in real time (minutes) or on a periodic basis (hours to days).

Business analysts and data analysts use decision AI tools to detect, analyze, predict, prescribe, and act. Then they build contextual intelligence by automating work that analysts used to perform manually: assessing outcomes, comparing them to expectations, and refining how they handle the next decision. While humans still contribute and oversee the process, they have the potential to achieve new levels of intelligence and automation.

Decision AI represents the third phase of a market evolution, building on the earlier phases of business intelligence and data science.

Drilling In

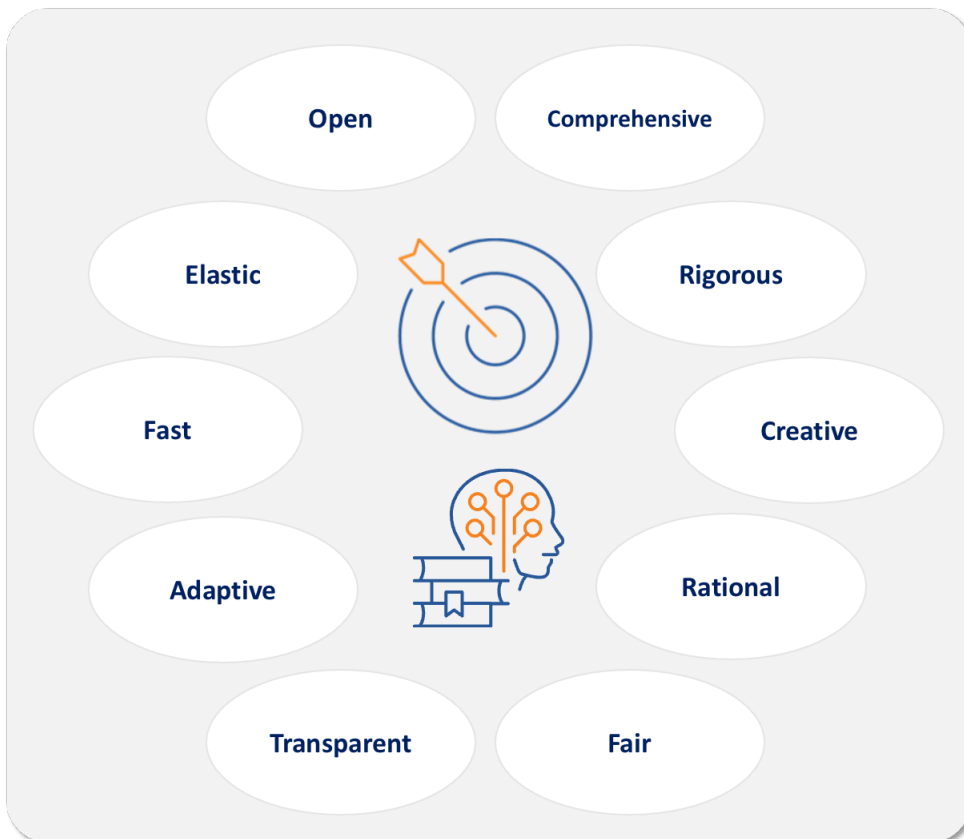
So how can enterprises ensure decision AI delivers on its promise? As with other types of AI, we can view this software as a smart child. It has the cognitive power to achieve great things, but needs parental oversight and periodic course corrections. With this in mind, let's consider the must-have characteristics and use cases for decision AI.

Must-Have Characteristics

Before committing to a decision AI initiative, enterprises must ensure they have the right tool for the job. A decision AI platform must be comprehensive, rigorous, creative, rational, fair, transparent, adaptive, fast, elastic, and open. These must-have characteristics have to strike a balance, for example between creativity and rationality, and between rigor and speed. (See figure 3.)

A decision AI platform must be comprehensive, rigorous, creative, rational, fair, transparent, adaptive, fast, elastic, and open.

Figure 3. Must-Have Characteristics of Decision AI



Comprehensive. The decision AI platform must cast a wide net by synthesizing, correlating, and analyzing all the data at its disposal. This generates more comprehensive and accurate business views, which in turn helps predict outcomes and prescribe actions with higher confidence. In addition, the decision AI platform should catalog decisions and actions, creating a comprehensive historical view that supports contextual intelligence and makes future decisions more accurate.

Rigorous. At each point in the lifecycle, decision AI must ask: what are we missing? For example, it might apply alternative ML models to the same set of facts and see how results vary. Or it might change key facts to check the sensitivity of predictions to different events. Humans also can make decision AI more rigorous because, as with parenting a child, they provide oversight. They should be able to review explanations of models and approve or reject their findings. They also should be able to help assess and compare old decisions.

Creative. As we learned in economics 101, information must change decisions and actions to have any value. Accordingly, the decision AI platform needs creative ML models that surprise decision makers by spotting anomalies, finding new correlations, and generating new insights. This yields novel predictions and recommendations to capture opportunity or reduce risk. Decision AI needs such creativity to give enterprises upside that justifies the cost.

Rational. Like all software, decision AI helps remove human emotion and its baggage—hesitation, habits, hunches, and so on—from the decision lifecycle. It studies facts and how they relate to one another, then draws rational conclusions about the logical next step. But to outperform humans, decision AI must go further and apply rational thinking to humans' tribal knowledge. It must ask probing questions of analysts and managers, then codify their responses to help shape future recommendations.

Fair. If not designed well, AI/ML models can perpetuate bias. Data scientists might unwittingly build their own biased assumptions into their ML models or train models on incomplete datasets that under-represent certain groups. On the other hand, ML models that overcome these design flaws can reduce traditional bias by replacing human assumptions with a rational look at the facts. To reduce bias, decision AI must gauge the sensitivity of model predictions to gender, race, or other group characteristics. It also must alert decision makers of unrepresentative training data and signs of their own biased decisions.

Transparent. To ensure fairness and drive positive outcomes, analysts and managers must understand how their machine thinks. Whenever the decision AI platform predicts an outcome, prescribes an action, or implements a decision, it must explain its logic in clear terms. What evidence, causality, and assumptions drive the analysis? What is the risk of bias? Are predicted outcomes likely to happen? Analysts and managers must understand these factors before approving or rejecting the decisions that the platform recommends.

Adaptive. As with all AI/ML software, the value of decision AI depends on its ability to learn and adapt. This enables both decision intelligence and contextual intelligence in our lifecycle. ML models enable decision intelligence by learning patterns in data that become the basis for detection, analysis, and prediction—as well as the recommendations and decisions that follow. ML models enable contextual intelligence, meanwhile, by adapting to the lessons of past decisions and outcomes. The decision AI platform can refine models by re-weighting certain factors, incorporating new data points, or making other changes.

Fast. Many operational decisions have a narrow window of time—perhaps minutes, perhaps hours—to capture the value or reduce the risk associated with a business event. A competitor discount, customer purchase, or supplier delay all might demand fast thinking. Within that window of time, the decision AI platform must synthesize high volumes of varied data to provide the decision-maker with a prompt and well-grounded recommendation. It can meet these low-latency, high-throughput requirements by relying on cloud compute resources.

Elastic. Decision AI platforms must consume compute resources elastically to support workload changes. Cyber Monday sales, quarter-end reports, or supply chain disruptions might trigger bursts of data processing that consume extra compute on short notice. Here again cloud infrastructure can help, although the decision AI platform should give users advance notice of the cost implications of scaling up and help them scale back down after the workload burst.

Open. Enterprises cannot afford to lock themselves into a monolithic system. Rather, their decision AI platform must integrate with an ecosystem that includes BI tools, AI/ML libraries, notebooks, cloud data platforms, catalogs, and data pipeline tools. It also should accommodate open file formats and application programming interfaces (APIs) that help ensure both tool interoperability and data portability. By maintaining open access to various commercial and open-source elements, the decision AI platform can help enterprises keep their options open for future innovation.

Use Cases

Decision AI best supports use cases that involve repetitive and time-consuming decisions based on volumes of data that would overwhelm humans using traditional BI tools. The use cases fall into the categories of operations and strategy. Let's consider examples for each.

Operational Use Cases

Enterprises should start with operational use cases, in particular those that involve tactical decisions based on identifiable patterns in structured and semi-structured data. Operational use cases span functions such as finance, sales, supply chain management, marketing, and pricing.

- > **Finance.** Every day, finance managers and analysts decide how to handle company funds. For example, a commercial bank decides whether to offer loans—or write off risky loans—by analyzing individuals’ credit histories, demographic trends, and the latest economic indicators. Decision AI can crunch a lot of data to learn patterns, predict outcomes, and recommend more comprehensive, rigorous, and rational decisions than humans might make otherwise.
- > **Sales.** Every week, sales leaders and SalesOps managers analyze their pipeline. They predict which deals will close and whether they’ll achieve their forecast this quarter. They assess lost opportunities by comparing them to expectations, and refine sales tactics based on lessons learned. Decision AI can add rationality and context to this process by pointing out numerical patterns that challenge sales leaders’ blind spots.
- > **Supply chain management.** Logistics managers and analysts juggle many risks as they manage supply chains in a volatile world. They must detect changes in inventories, shipment schedules, or component availability, analyze the implications, then predict outcomes and decide how to respond. Decision AI can add rigor and rationality to this process by comparing outcomes for different decisions and analyzing their sensitivity to different factors.
- > **Marketing.** Each quarter, marketing managers decide how to spend budget on lead-generation programs. They tabulate historical results, analyze basic KPIs such as cost per lead and revenue contribution, then decide which future programs will best meet their lead goals. In the end, their decisions derive mostly from habit and hunches. Decision AI can offer more comprehensive analysis of program elements—such as content topics, program vendors, prospect lists, and target personas—and then recommend more creative programs.
- > **Pricing.** Product owners weigh many factors as they set prices. They analyze product features, historical sales, customer preferences, and competitor prices, then predict revenue and unit sales based on different price alternatives. Again, habit and hunches tend to prevail. Decision AI can help make a rational price decision, then adapt that decision based on a clear-eyed assessment of the outcome.

Operational use cases span functions such as finance, sales, supply chain management, marketing, and pricing.

Strategic Use Cases

If enterprises achieve results with operational use cases, they next can consider tackling strategic use cases. These use cases follow a looser schedule and have more variable data sets than operational

use cases, which makes the patterns more difficult to identify. Still, decision AI has the potential to drive faster, smarter action in the strategic domain. Strategic use cases include projects such as market analysis and product design.

- > **Market analysis.** Product owners and marketing leaders periodically analyze the market for their offerings, for example to support a product launch. They analyze market demographics, survey findings, and historical data about company and product performance. They predict customer demand to decide on a course of action. Decision AI can add rigor and rationality to the analysis by comparing the outcomes of different predictions.
- > **Product design.** Product owners also analyze customer requirements to help guide product design. They assess survey data, historical sales, and competitor offerings, then predict what their target customers want to buy. As with market analysis, decision AI can add rigor and rationality to the analysis and resulting design decisions.

Strategic use cases span projects such as market analysis and product design.

Getting Started

Decision AI creates attractive opportunities for the modern enterprise. By combining decision intelligence and contextual intelligence in one automated lifecycle, it can overcome human limitations to drive smarter, faster, and more adaptable action. To capture this opportunity, enterprises need a decision AI platform that is comprehensive, rigorous, creative, rational, fair, transparent, adaptive, fast, elastic, and open. Consider these additional guiding principles as you scope the decision AI opportunity for your enterprise:

Start with low-hanging fruit. As with other AI initiatives, decision AI offers the most achievable results with simple, tactical use cases. Ask the managers and analysts within your business functions where they feel the most pain. Where is data overload preventing them from making timely and accurate decisions? These obvious points of pain often represent the most promising use cases for starting a decision AI initiative. You also can start your initiative by focusing on decision intelligence, then branching into contextual intelligence.

Compare commercial and homegrown options. Multiple vendors offer decision intelligence tools, and new vendors such as Diwo combine decision intelligence with contextual intelligence to automate the learning process over time. These commercial tools reduce your team's manual workload, but might

limit how you address custom requirements specific to your business. Compare both approaches before committing to a commercial tool. For either option, be sure to scope the work required to train AI/ML models on your data.

Assemble a cross-functional team. Many stakeholders with varying expertise and skillsets shape the success of a decision AI initiative. The team starts with business owners that set strategic objectives and oversee the initiative. It also includes business and data analysts that design and manage the platform, data scientists that train its AI/ML models, data engineers that integrate the data, and developers or DevOps engineers that build the workflows for taking action.

About Eckerson Group



Wayne Eckerson, a globally-known author, speaker, and consultant, formed **Eckerson Group** to help organizations get more value from data and analytics. His goal is to provide organizations with expert guidance during every step of their data and analytics journey.

Eckerson Group helps organizations in three ways:

- > **Our thought leaders** publish practical, compelling content that keeps data analytics leaders abreast of the latest trends, techniques, and tools in the field.
- > **Our consultants** listen carefully, think deeply, and craft tailored solutions that translate business requirements into compelling strategies and solutions.
- > **Our advisors** provide one-on-one coaching and mentoring to data leaders and help software vendors develop go-to-market strategies.

Eckerson Group is a global research and consulting firm that focuses solely on data and analytics. Our experts specialize in data governance, self-service analytics, data architecture, data science, data management, and business intelligence.

Our clients say we are hard-working, insightful, and humble. It all stems from our love of data and our desire to help organizations turn insights into action. We are a family of continuous learners, interpreting the world of data and analytics for you.

Get more value from your data. Put an expert on your side. [Learn what Eckerson Group can do for you!](#)



About the Sponsor

Diwo's Decision Intelligence (DI) platform solves the "last mile of analytics challenge" by continuously scanning your data environment to discover why metrics change and anomalies arise. Diwo then surfaces contextual insights and actionable recommendations that a user can act upon in minutes, giving them a faster time-to-decision compared to traditional BI.



What makes Diwo different from traditional business intelligence tools is our focus on decisions, not dashboards. Diwo's DI platform enables the automation and augmentation of decision-making to optimize business, financial, and operational performance. Diwo increases the speed and quality of decisions across the organization by delivering AI-powered recommendations with a clear explanation of their impact (revenue, cost, operational) directly to a business user via a customized UI. Diwo combines human intelligence with patented contextual intelligence, AI, and a unique decisioning experience to help users respond to changes in their business by making faster, more confident decisions that drive measurable value.

Organizations across industries, including retail, consumer products, financial, and manufacturing, are adopting the Diwo DI platform to eliminate BI bottlenecks and drastically reduce the time and effort necessary to make strategic, tactical, and operational decisions.

To learn more, visit www.diwo.ai.