White Paper

Unlock the Potential of AI/ML with Decision Intelligence

LUCAS GLASS, Vice President, Analytics Center of Excellence, IQVIA
SHARATH BENNUR, Director, AI Solutions Delivery, IQVIA

Executive summary

Despite the years of sustained investment in data science teams, many enterprises are struggling to show ROI from their investments in Artificial Intelligence and Machine Learning (AI/ML). This is highlighted by a survey of organizations in the life sciences industry where 30% of organizations said their top challenge was difficulty identifying AI/ML use cases, and 18% are being challenged to prove business value of their AI/ML solutions. Additionally, in a survey of IQVIA clients, only 1 to 2% of clinical trial professionals considered AI/ML generated insights as a top priority. The common theme - an inability to demonstrate the value of insights generated by AI/ML and to reliably scale the generation of new AI/ML use cases across the business.

Fortunately, Decision Intelligence – a nascent field that combines behavioral research into decision-making with a thoughtful application of AI/ML – can help fill these gaps. By enabling structured discovery of decision points and their decision contexts, decision intelligence can help identify where enhanced data, AI/ML, or other statistical and rule-based techniques can enhance the user's decisions. It also facilitates an understanding of user preferences as it relates to trusting and consuming AI/ML. These insights help drive smarter decisions and business value by integrating AI/ML into a user's workflow seamlessly, in a manner that increases adoption of the insights and builds trust in technology products.

CHALLENGES FOR AI/ML DEVELOPMENT

Artificial Intelligence and Machine Learning (AI/ML) have the potential to radically transform all parts of the molecule research and development (R&D) process. Applications include optimizing molecules, improving administrative efficiency, optimized clinical trial planning and design, virtual assistants, trial finance predictions, trial risk predictions and patient-trial matching. While AI/ML has a lot of potential to positively effect these processes, organizations face challenges and bottlenecks in leveraging it successfully across the organization. In a recent broad study on the status of AI adoption[1], when asked about the main bottlenecks holding back AI adoption, the top two responses were:

- Company culture not recognizing the need for AI (22%)
- Difficulties in identifying appropriate business use cases (20%)

In a different study on scaling of AI in life science companies by Deloitte[2], 30% of organizations had difficulty in identifying use cases with the greatest business value, 28% had operational challenges with integrating AI into the organization and 18% of organizations were challenged to prove business value of the AI/ML being built. Similarly, when IQVIA asked its clients about the role of AI/ML generated insights in clinical research, only 1-2% of clinical trial professionals saw AI/ML generated insights as a top priority[3].

What is holding back organizations from delivering on the potential of AI/ML? While the ability to identify and deliver AI/ML products is not unique or limited, the capacity of the organization to scale these abilities is. Limiting factors include the availability of capable data science teams and supporting processes, and of teams that can efficiently translate business needs within the organization into actionable AI/ML insights. Bottlenecks also extend to identifying when and where users would benefit from the application of AI/ML and then measuring the value added to users. These challenges are also accentuated by existing workflows and processes, which, when combined with heuristics and biases that users might have, make the adoption of AI/ML recommendations difficult at both the individual and organizational level. Data science and product management teams are typically not setup to examine these questions holistically to drive such insights.

So how have organizations been addressing these challenges? Recently, there has been a surge of interest in decision intelligence led by work out of Google[4]. We believe that decision intelligence is key to focusing decision support resources – technology, data, or AI/ML – on those decision points within an organization that will have the most impact. A core element of decision intelligence involves analyzing business processes and workflows to identify how and where employees make business-critical decisions, and then use the decision intelligence framework to analyze the context of the decision. This can provide a critical bridge to understanding where the technology, data or AI/ML can help improve the decisions themselves and the efficiency of surrounding processes.

Humans are not ideal decision-makers; in fact, our decision-making is influenced by many biases that can impair our judgement. Additionally, we tend to use mental shortcuts (heuristics) that allow us to bypass the high cost of processing a lot of information.

Let's use a simple example to illustrate this. One of the authors of this white paper uses Google Maps to navigate to work every day. Google has a complex set of graph neural networks at the backend of the application making recommendations for the best routes to take[5]. Unfortunately, the author often uses a heuristic (the route he takes most often) to choose his route – so, on many days the recommended shorter route from Google Maps is overruled by this heuristic. Similarly, on other days, the author prefers a route with less stop-go highway traffic (a bias) that causes him to choose a slower route. This is an example of how humans are highly skilled at circumventing slower, rational decision-making in favor of fast unconscious decisions – which are not always optimal or accurate.

The less-than-ideal aspects of human decision-making are what make decision intelligence useful, as it enables the discovery of user's decision contexts, their biases and heuristics. This in turn can guide the choice of decision support systems in a technology product to support better decision-making.

OPTIMIZING AI/ML AND DATA APPLICATIONS THROUGH DECISION INTELLIGENCE DISCOVERY

How can decision intelligence be used to optimize the AI/ML discovery process? Our first step is to map user workflows and business processes to understand where business-critical decisions are being made. The next step is to examine the decision context of those decisions (Figure 1).

FIGURE 1: ILLUSTRATION OF A DECISION AND ITS CONTEXT

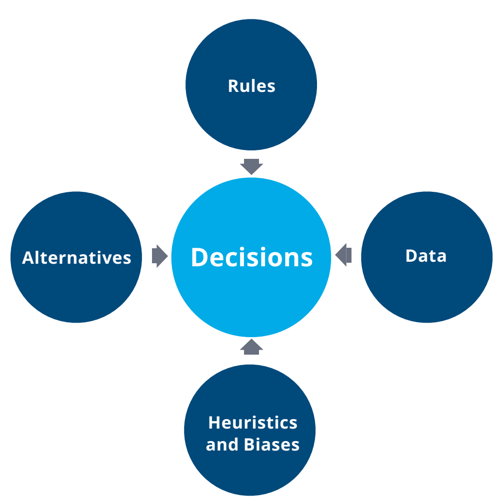

TABLE 1: THE ROLE OF DATA IN DECISIONS

DATA-DRIVEN	Data has to meet an exact threshold to drive a decision. E.g., If the demand is greater than X, we will invest in a new plant
DATA-INFORMED	Data is used to understand past and future performance and form new strategies E.g., If the demand is positive and our sales teams predict positive demand for next year, we will invest in a new plant
DATA-INSPIRED	Data is used for trendspotting and observe general behaviors. E.g., Our competitor opened a new plant in this country, and we should as well

The decision context includes different aspects of the decision – who the decision-maker is, what the default decision is in the absence of data, what alternatives a user might choose, and what supporting data could be used to validate an alternative choice. Note that the use of data does not indicate that decisions are driven by data – decisions can be data-driven, data-informed or data-inspired, with each level being less informed by the data itself respectively (Table 1).

Once the decision context for a decision is understood, our next step is to assess how we can better support those decisions. Decision intelligence is not prescriptive regarding which decision support tool is appropriate, as the cost-benefit tradeoff between alternatives can vary greatly among individual cases. Instead, we suggest focusing on the decision and improving support for the decision to make it more data-informed or data-driven. Decision support can be improved through: finding new or augmenting current data sources, analytics dashboards or data journalism, improved statistics or analysis of the current data, and sometimes through rules-driven automation or the use of AI/ML.

Another aspect of this discovery process is highlighting any heuristics or biases that influence decisions. The important principle here is to recognize heuristics and biases and determine their redressal. In our view, not all heuristics (e.g., lack of certain data) may be addressable. We suggest triaging the discovered heuristics and biases with stakeholders and business leaders to plan, prioritize and execute remediation as appropriate. A careful evaluation of the tradeoff between the cost of addressing the heuristic (e.g., acquiring new data) vs continuing to use the heuristic is needed.

Readers familiar with behavioral analysis, behavioral economics and psychology would be familiar with some heuristics and biases including the availability heuristic, outcome bias, hindsight bias, framing effects, loss aversion, sunk cost effects, confirmation bias, conjunction fallacies, and others. While this white paper does not attempt to describe the heuristics/biases in detail, we urge readers to become familiar with them in order to better identify them and plan remediations as a part of their decision intelligence discovery and exploration activities.

The final step in our discovery process is to connect decision-making support with AI/ML. A useful rubric for identifying AI/ML use cases is to assess the complexity or uncertainty of the data. High complexity/uncertainty can highlight where AI/ML could be most impactful. AI/ML can be used to automate the decision itself or enhance and predict the data used in decision-making. Heuristics or biases sometimes indicate the lack of reliable data for a decision, suggesting that data enhancements or AI/ML can address this.

DRIVING AI/ML DISCOVERY

Key elements to find AI/ML use cases that can drive business value:

- Understand the decision-context and data used for decision-making through user research
- Assess the uncertainty and complexity to determine if AI/ML are appropriate
- Analyze dimensions of the AI/ML that can drive adoption
- Define KPIs to evaluate and measure business value and user-adoption

DELIVERING TRUSTED AI/ML THAT USERS ADOPT

Users are critical to informing the entire developmental cycle for decision support and decision intelligence. This includes conducting user research and identifying personas to discover decision contexts, which are key inputs for determining the requirements for improved decision support.

When AI/ML is the decision support being considered, it is critical to understand the importance of various AI/ML dimensions such as explainability, diversity, fairness and actionability (amongst others) for users who consume the recommendations. These are important inputs for the data science team as they assist in defining requirements and deliverables for AI/ML projects, and help determine how these behaviorally driven insights are delivered to users while building trust in the AI/ML recommendations.

Yet, building trustworthy AI/ML that users want is not enough. The user experience (UX) of AI/ML generated insights lag behind the UX for traditional tech. Thus, the adoption of AI/ML can be greatly enhanced by thoughtfully integrating AI/ML output in a manner that users can consume easily.

For example, some user personas may prefer to look at a dashboard to dig into the data and understand it before arriving at a decision, whereas other user personas may prefer the delivery of a data point in a human-readable format through a chatbot. Analyzing personas and weighing the relative cost of building either a dashboard or a chatbot, for example, can help teams prioritize between AI/ML integration and delivery options. Ideally, AI/ML integration and delivery is seamless, such that it fits naturally into the user's workflow.

TRACKING THE SUCCESS OF AI/ML AND DECISION INTELLIGENCE

Once business-critical decision points have been discovered, their decision contexts determined, and options for improved decision support considered, how do we know if we are succeeding?

To be truly successful, AI/ML must be complemented by a more precise definition of the success criteria – not just from the machine learning perspective but also from the perspective of the business and users. In order to achieve this, we suggest complementing machine learning success criteria (e.g., R-squared, AUC scores, F1 scores, etc.) with business and user success KPIs. This requires identifying what success looks like from a business perspective upfront (e.g., X% increase in adoption of AI/ML recommended countries), and what success looks like from a user perspective (e.g., Y% reduction in time taken for evaluating countries to include in a trial).

These criteria can help define and focus a data science team's goals and provide a business justification for continued investment in a particular project or team. Feedback about the utility of the AI/ML can be collected either implicitly within the product (e.g., measure how often the country recommendations from a AI/ML system are being adopted/ignored by users) or through a more explicit direct survey of user preferences (e.g., user surveys asking how they perceive the augmentation of their decision-making through AI/ML, etc.).

Proving business value and the lack of executive commitment to AI/ML were cited as major challenges by life science companies adopting AI/ML in Deloitte Global 2020[2] and in IQVIA's own market research. Decision intelligence provides a critical bridge to understanding how AI/ML is targeted to focus on providing decision support where it's most impactful.

Additionally, consider a diverse set of decision support approaches (e.g., enhancing data, statistics business rules, and AI/ML). This in turn addresses another area of concern for users: that they are not always comfortable relying on the judgement of an AI/ML system to make business-critical decisions.

Case study: Country selection for clinical trials

As a case in point drawn from our work at IQVIA, let us consider a step in the clinical trial planning process where countries are included or excluded from a particular clinical trial. The choice of countries for a trial is critical, as it determines the speed and success of the trial. Let's start by examining the decision context – recall that the decision context includes the decision makers, the data being used and any heuristics or biases that influence the decision.

In this case, the decision-maker for the trial's strategy is the Therapeutic Strategy Lead, a leadership role that determines the trial's overall strategy. What is not obvious at first glance is that there are additional decision-makers – specifically, regional teams whose recommendations are critical for determining if a country is appropriate for a particular trial, and analytics teams who provide further nuance to the data used for this decision.

Next, let's consider the data used to determine if a trial could be successful in a particular country. These include the following for a particular country: the number of potential patients for an indication, historical recruitment rates for that type of trial, the number of sites IQVIA supports in the country, and so on. In order for a country to be considered viable for a clinical trial, the supporting data for a country on many of these data points has to exceed that of other countries being considered. Together, these make up the decision context for this decision to include/exclude a country in a particular trial.

As mentioned, understanding heuristics, biases and existing process are also key to improving decision support. In some countries, there is minimal or no historical data for clinical trials at a specific site or in a therapeutic area. In these situations, 'estimates' are used to fill in the gaps. While these heuristics may be reasonable given the absence of data, it also suggests that better decision support is needed – which in this case could be accomplished by acquiring/augmenting source data or using AI/ML to predict missing data.

Similarly, the number of potential patients with an indication in a country can be very hard to determine, as many patients are often undiagnosed. Recall that complexity of the data is a key indicator for using AI/ML. Here, the complexity of the underlying EHR data plays a key role in suggesting that AI/ML would be useful in predicting the number of patients with the rare disease. Similarly, when historical trials are used for predicting recruitment rates for a new trial, there is often uncertainty in knowing which trials are appropriate to use for comparison. Here again, uncertainty is a key driver for recommending the use of AI/ML to assess which therapeutic areas or trials are best used in this comparison.

Case study: Matching patients to clinical trials

Let's consider a technology solution being built internally in IQVIA to recruit patients for clinical trials by matching them with the appropriate trials.

Initial versions of the solution screened patients for a primary trial of their choice and potentially a secondary trial that they might be interested in. Here, the patient has to decide whether to accept or reject any additional clinical trials that might be suggested. The decision context would include the default decision to reject any additional clinical trials being offered outside of their original therapeutic area of interest. The alternative choices could be to accept or reject the suggested trial but be open to other trials in the future.

Another aspect of the decision intelligence framework is considering the data used for the decision. What factors might a patient consider in deciding to accept an alternate clinical trial? These include other co-morbidities that occur alongside the patient's current indications; the patient's current healthcare burden, and the potential burden of participating in a trial. These factors vary by trial.

These insights about the decision context and some additional heuristics discovered led us to re-examine what data was being gathered from patients in order to better match them, and how our matching algorithm was set up. For example, understanding any co-morbidities that a patient might have could help match them to other trials, and could help estimate their current healthcare burden more accurately. The relationship between co-morbidities is complex, which in turn suggests that this is potentially an AI/ML use case. While this work is on-going, we believe that these insights into patient decisions are critical for building better decision support into our technology products.

HOW IS DECISION INTELLIGENCE BEING USED TODAY?

At IQVIA, the challenges we've faced internally in scaling AI/ML as well as in increasing user adoption and trust in AI/ML solutions are similar to those in other organizations. This provided the impetus to build the decision intelligence strategy and practices that are discussed here, and to deploy them within our organization, as discussed in the case studies provided.

A deliberate deep dive using decision intelligence frameworks into our workflows, processes and products has provided numerous benefits. Once the business-critical decisions are identified, we are able to focus our decision support – irrespective of whether they are technology, data or AI/ML solutions – on those decisions. These tools also help us discover heuristics and biases that might be sources of friction in the organization.

Where AI/ML is the appropriate decision support solution, we also have a better definition for what user expectations are for using/adopting AI/ML and for seamless delivery.

We are using this approach to discover and integrate improved decision support and AI/ML in products for clinical trial planning and execution, for managing and tracking the costs of clinical trials, for patient recruitment, and much more. In our view, the principles of decision intelligence being used in the Analytics Center of Excellence (ACOE) at IQVIA are applicable to any organization that is trying to build AI/ML at scale and wants to further enable the success of their AI/ML initiatives.

While some of these insights can be discovered using traditional UI-UX research frameworks, our focus on decision-making and AI/ML enables the discovery of AI/ML use cases in a novel manner not addressed by current UI-UX research techniques. What follow are some of our learnings and best practices.

- **Incorporate behavioral research tools and expertise.** Use a combination of behavioral researchers and tools in a scalable manner across product teams to generate insights (Figure 2). For organizations attempting to build their own decision intelligence practice, ideal teams or individuals combine product management, behavioral research and data science skills. This new approach of combining behavioral research with AI/ML is still in its infancy and there is a dearth of experienced professionals in this space. In addition to finding the right people, this research also requires the right tools. In the ACOE, we have built both new tools and re-purposed elements of UI-UX research frameworks to drive our decision intelligence research.

- **Track usage.** The technology stack in our products have to support the requirements described in earlier sections. Life science and healthcare organizations have unique challenges when it comes to collecting user data, but collecting data is nonetheless essential to understand the success of initiatives. It's critical to be able to track the full lifecycle of the decision and any decision support tools built into the product.

- **Collect user feedback.** We are building the ability to collect user feedback into our products, both implicitly within the product and explicitly from users. Traditionally, the emphasis has been on tracking and monitoring the models themselves and the data used to train/generate the models. However, this does not address user adoption. We suggest tracking the full lifecycle of all decision support tools, including AI/ML, to understand whether it was utilized for decision-making. In this way, the learnings can be used to influence future development.

FIGURE 2: REQUIREMENTS FOR IMPLEMENTING DECISION INTELLIGENCE

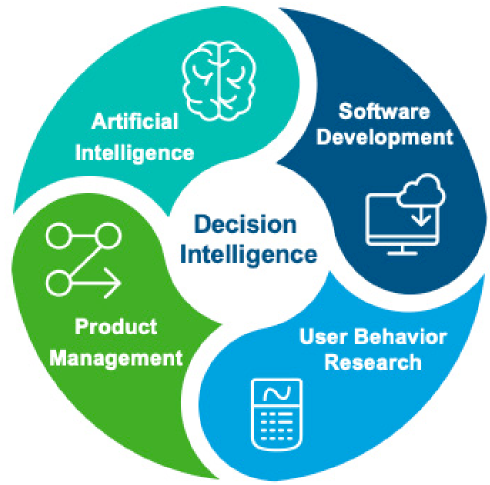

- **Invest in education.** Another important aspect of using decision intelligence in the organization is education - not only for stakeholders, but also for product management, data science and developer teams. The long-term goal is to ensure that all team members and roles involved in building the decision support tools – irrespective of whether they take the form of data, statistics, rules-based systems or AI/ML – are familiar and comfortable with incorporating decision intelligence into the product research and development process.

Ultimately, we want our team members to understand not only the principles of decision intelligence, but also which questions need to be addressed from a user's perspective. For example, do data science teams and developers understand why an AI/ML solution is being built, who its users are, how the users will interact and consume the AI/ML solution, where the solution fits into a user's daily life or workflow, and how the solution is helping users? We are we gathering feedback from the business and stakeholders and measuring the added value realized from adopting our decision support solutions.

WHAT DOES THE FUTURE HOLD?

It's still early in our journey of building and integrating a decision intelligence strategy and practice within the ACOE. Using decision intelligence internally has provided a number of additional insights and led to the discovery of heuristics in different internal processes, but this is just the beginning.

Within the ACOE, we are considering two future areas of focus. The first is to understand how we can reliably measure changes in AI/ML adoption and additional decision support.

For example, do we know where AI/ML is used? Is it easy to understand and consume, is it trustworthy, is it going to help users with their jobs? Table 2 lists examples of the types of questions users might be asked when an AI/ML enhanced application that recommends clinical trial sites is being rolled out. This table is loosely based on work from the field of Information Sciences for measuring technology adoption[6].

TABLE 2: AI ACCEPTANCE SURVEY QUESTIONS

Concept	Description	User survey question
PERCEIVED USEFULNESS	Does the user think an AI/ML solution will make them better at their job?	Do you think the site recommendations help you do your job better?
PERCEIVED EASE OF USE	Does the user think the AI/ML solution is easy to use?	Are the site recommendations easy to understand?
INTRINSIC MOTIVATION	Does the user want to use the AI/ML solution?	Do you want to use the site recommendations?
EFFECT TOWARDS USE	How does the user feel about using the AI/ML solution?	Are you happy or unhappy to use the site recommendations?
RESULTS DEMONSTRABILITY	Does the user understand how the AI/ML solution makes their work easier or better?	Do you understand how the site recommendations make your work easier and/or better?
AI-ML INFERENCE QUALITY	Does the user think the output of the AI/ML solution is of good quality?	Do you understand why these specific site recommendations were made?
IMPROVED DECISION-MAKING	Does the user think the AI/ML solution helps them make better decisions?	Do the site recommendations help choose better sites for a trial?

Our second area of focus is to use insights from adjacent fields like behavioral economics and psychology to further the effectiveness of decision support in our products and programs. There has been a significant body of work from researchers looking at the efficacy of behavioral interventions in order to drive optimal or desired behaviors from users. In the future, we want to adopt principles such as setting default actions, adjusting the choice architecture for decisions, providing planning prompts, and leveraging social norms in our products to promote both adoption of AI/ML solutions and to fully leverage the lift from the AI/ML solutions. The primary challenge to achieving these admittedly ambitious goals is both having the necessary technologies built into our products, and having the ability to measure whether interventions are adding value to both users and the business.

CONCLUSION

The key to unlocking the potential of AI/ML and building great products that empower users to make better decisions, is to understand the context and heuristics around how users make decisions in their work and daily lives. Users want to make better decisions, but they don't always know how. This is where they can greatly benefit from better decision support broadly and, when appropriate, from AI/ML solutions.

What is unique about decision intelligence is not just the focus on business-critical decisions. Through unearthing insights into a user's heuristics and biases, we have a much better chance at addressing the entire decision context holistically within our products.

In the Google Maps example provided at the beginning of this white paper, a solution built using insights from decision intelligence could take into account the author's heuristics and biases by showing the preferred routes for the author first – provided those routes are not significantly longer than the results from the graph neural networks. When there is a significant difference between the author's heuristic and reality, this author would prefer a nudge from Google Maps encouraging a better decision. The future looks exciting as we look to extend these insights into product-driven interventions that can drive further adoption and value.

References

1. O'Reilly Media. 2020. AI adoption in the enterprise 2020. March 18. https://www.oreilly.com/radar/ai-adoption-in-the-enterprise-2020/

2. Deloitte Global. 2020. Scaling up AI across the life sciences value chain. November 04. https://www2.deloitte.com/us/en/insights/industry/life-sciences/ai-and-pharma.html

3. IQVIA's Voice of the Customer Survey

4. Kozyrkov, Cassie. 2019. Introduction to Decision Intelligence. August 02. https://towardsdatascience.com/introduction-to-decision-intelligence-5d147ddab767

5. DeepMind. 2020. Traffic prediction with advanced Graph Neural Networks. September 03. https://deepmind.com/blog/article/traffic-prediction-with-advanced-graph-neural-networks

6. Sohn, Kwonsang , and Ohbyung Kwon. 2020. "Technology acceptance theories and factors influencing artificial Intelligence-based intelligent products." Telematics and Informatics 101324
   Venkatesh, Viswanath , Michael G Morris, Gordon B Davis, and D Fred Davis. 2003. "User Acceptance of Information Technology: Toward A Unified View." MIS

CONTACT US

iqvia.com