IN THIS CHAPTER

» Becoming familiar with the decision intelligence approach

» Comprehending the method, principles, and priorities of decision intelligence

» Working your way from design to reality

» Seeing the difference an inverted V makes

» Implementing for the win

Chapter 1

Short Takes on Decision Intelligence

Do you find yourself looking at a spreadsheet or viewing charts or gazing glassy-eyed at a fancy visualization that some bit of artificial intelligence magic has produced for you and wondering what you should do next? You're not alone. Millions of other business and finance people are doing the same thing. So are legions of leaders and decision-makers in other industries.

While you're trying to puzzle out which parts of those "actionable insights" being handed to you *are* in fact actionable and, if so, what action would apply, you've likely wished for something a bit more cut-and dried when it comes to determining what your organization would implement — and you certainly wouldn't mind being considerably more certain about what's going to happen post-implementation.

Would your best bet in such a situation involve letting the miracle of artificial intelligence (AI, for short) make your decisions for you? Well, it turns out that AI isn't so miraculous. In fact, an estimated 80 percent of all AI projects fail, where *failure* here is defined in terms of failing to deliver a measurable business value. That means most AI projects end up on the trash heap for leaning too heavily on the experimental side and being useless on the applied side.

It is painfully (and expensively) obvious that this strategy isn't quite working out the way everyone hoped. An alternative approach is needed to make data more helpful and better aligned with consistently delivering business value. One such approach flips the model from data driven processes to decision driven processes. Known as decision intelligence, human and machine decision-making skills are combined with decision theory, decision sciences, and data sciences in a customizable mix that pins decisions to a precise and expected business value.

The concept isn't entirely new — one of its oldest published mentions cropped up in 2002 in Uwe Hanning's scholarly paper "Knowledge Management + Business Intelligence = Decision Intelligence" — but it has evolved over time, incorporating long-accepted scientific formulas from several well-established sciences. This means its inner workings are well known and tested. Switching over to a decision intelligence approach is therefore no gamble — it's simply a supremely logical way for you to achieve the business outcomes you desire. Decision intelligence leaves little to chance, in either its own construct or the value it consistently delivers.

What differentiates one decision intelligence project from another is the talent and acumen of the decision makers at the helm. They make the recipe that cooks the business value into the process. And they decide when and whether to invite data and machines to the planning table.

Decision intelligence is highly agile and versatile. Decision makers can use it to make decisions either on the back of a napkin or with the help of the most sophisticated AI on the planet.

The Tale of Two Decision Trails

The business world has long been madly in love with the notion of being a data-driven enterprise, but it's also beginning to feel the pain of being in a bad relationship. Few actually want to break off their relationship with data entirely, mainly because most are loathe to ditch their significant investments in data, analytics, and related technologies. Add to this the fact that, for many, it would feel like a colossal failure and a huge embarrassment to fall short of becoming the data driven enterprise that all investors and stockholders expect these days.

Looking for a way forward, many started to ask themselves this question: "What can we do with all the data investments we've already made and already own in order to make better decisions?" In other words, folks realized that a *rethink* was needed rather than a *redo*. And many of those same folks decided that re-strategizing and restructuring how these same investments are used and aligning them with specific business impacts was the answer to the questions that had been plaguing them.

REMEMBER

A decision intelligence approach doesn't mean that there's no place for more traditional data mining tactics. Most organizations are using a combination of both, and it's already proving to be a winning play for many of them.

Pointing out the way

AI and data analytics no doubt deliver real business value in some use cases. They're helpful when it comes to recognizing patterns in massive amounts of data and spitting out equations, scores, predictions, and estimates. The point is that such facts *point to* possible decisions but suggest none. (That's why I refer to the outputs from such tools as *pointers* in this section.)

These tools are also capable of automating certain decisions based on business rules that are determined and set by you or your organization. At its core, AI is automated decisions at scale. Traditional analytics must be integrated with automation software to cause an action to occur.

But before the various software — analytics, AI, and various forms of automation — begin their work in producing insights and automating your decisions, someone has to either program the analytics and automation software, and/ or train the AI. This group of data professionals often provide the interpretations of the outputs as well (usually as visualizations and/or automated AI-generated narratives).

In other words, people in specific job roles who do these tasks typically determine which insights — pointers, in other words — are accessible to other people in the organization who either use the software in a much more limited way or only view the results on dashboards to consume the information. Given the high degree of data illiteracy throughout organizations and across countries and industries, this process is both logical and necessary.

The downside here is that it is also limiting what information end-users can access when it comes to their own decision-making processes and what prompts the direction their thinking takes. This is why data democratization and AI democratization — decentralization so that more people in the organization can use the tools — is so critical to businesses. By making these tools far more user-friendly, professionals in other disciplines and employees at all levels of the business can make better use of these resources.

However, both data and AI democratization still require data professionals to develop more intuitive and highly automated software to remove barriers before non-data professionals can use the tools in ways that bring their own talents and skill sets to bear. Think of this as very similar to the path other software has taken. For example, Microsoft Office enables people to create documents, notes, spreadsheets, and PowerPoint presentations without knowing how to write code, what keyboard commands to give, or anything at all about how the software works. This is the path analytics and AI software are headed down now.

So, who are these data professionals who are making and/or using analytics and AI to provide you with the pointers you're currently getting from various analytics software?

Typical job roles in data mining and analytics are data scientist, business analyst, data mining specialist, and data mining engineer (and variations of the same) to reflect a specific industry such as healthcare data analyst and risk-mining data scientist.

In AI, job roles include AI scientist, AI researcher, business intelligence developer, robotic scientist, software architect, data scientist, and data engineer, among others.

All these jobs will continue to be important positions in many organizations, and the demand for people with these skills will remain high for the foreseeable future.

However, much of their work is also being automated as part of the data and AI democratization movements.

As to specific examples of the work that these professionals collectively and individually produce for use in several business areas, below are some of the more common use cases for traditional data analytics and/or AI automated decision making:

- **Anomaly detection,** also known as outlier analysis, is a step in data mining (which can be aided by AI/ML or not) that finds deviations in the data from the norm, such as events (purchases on a charge or debit card in another country from where the cardholder is known to live or be, for example), and data point changes (attempts to change a social media account's password via a device or browser that the true account holder is not known to use before, for example).

- **Pattern recognition is** the automated recognition of patterns discovered in the regularities in data. One example would be finding earlier signs of cancer in patient data than doctors and diagnosticians previously knew existed.
- **Predictive modeling,** also known as predictive analytics, analyzes historical patterns in the data using a mathematical process to predict future events or outcomes. One example would involve predicting when a machine part will need repair or replacement based upon its past usage compared to how long identical parts lasted under the same conditions.
- **Recommendation engines** analyze data to make recommendations or suggestions based on users' past behaviors. Examples include analyzing your purchasing patterns in order to offer you a coupon for a grocery item you should be ready to buy again soon, or to recommend a movie based on movies you watched and rated earlier.
- **Personalization systems** use data analysis to customize a service, product, or automated communication. Examples include marketing emails sent to large numbers of customers, each personalized with the customer's name and a custom discount offer for a favorite product or service.
- **Classification and categorization systems** automate the organization of vast amounts of data. Examples include sorting data files and data sets according to importance, topic, secrecy level, or other identifier; legal requirements governing the handling of specific data points (think of laws like General Data Protection Regulation (EU GDPR) which limit where personal data can be stored); and the nature of the data (such as structured machine data or unstructured Twitter posts). Data must be correctly classified and categorized for analytics or AI to work correctly. Automation is the ticket here because there's so much data that it's impossible to do it manually.
- **Sentiment and behavioral analysis** is contextual data mining to discover and analyze the subjective expressed responses (sentiments or feelings) about a brand, product, service, idea, political candidate, and so on in online conversations or customer channels (conversations and customer ratings found in texts, on websites and blogs, in voice recordings or streams during phone calls, and app rating systems. Did you rate that Door Dash driver's service in the app? Yeah, that sort of thing!) Behavioral analysis can extend beyond sentiment analysis to include things like how long you spent reading a news article on your phone and how many times you return to a website, to what time of day and what device you normally use to post on Facebook.
- **Chatbots and conversational systems** frequently appear as a popup sales or customer service chat box on websites where you can ask questions about a product or service or your account and get an automated answer from the resident AI-powered chatbot. Some of these are so good it's hard to tell they

aren't human customer service agents. Data on the user and on the stated problem is collected and analyzed to rapidly respond with answers the user needs. Examples of other conversational systems include every digital assistant you've ever heard of: Alexa, Siri, Google Assistant, Bixby, and Cortana. Each is a data king, with Alexa and Google Assistant reigning over two of the largest kingdoms in terms of technical and market prowess.

- **Autonomous systems** are actually a network or a collection of networks that are all managed by a single entity or organization. Data is live streamed and typically analyzed at the sensor or gateway level, although some data is often sent to a data center for additional analyses later. Think the Internet of Things, such as self-driving cars, robotic systems in manufacturing, and smart cities that use information and communication technologies (ICT) to increase operational efficiency, share information with other systems (such as self-driving cars), and promote sustainable development.

There's no question that the above list is populated with wonderous achievements that would not be possible (or at least not at such huge scales and fast speeds) without data, analytics, and AI. Nevertheless, the promised "actionable insights" produced by analytics and presented in many of today's fancy visualizations and dashboards to business users are often merely pointers. They point to something you might want to use as a key factor in your decision, but they aren't in a position to make that decision for you. You have to conjure some mad data interpretation skills and do some creative problem-solving on your part to figure that one out on your own.

Making a decision

Pointers (also known as *actionable insights*) are typically useful in so far as they go. The trouble is that they point to possible decisions but don't suggest any. Users are often unsure about what action to take, or which option would produce greater value for the business. By contrast, the decision is the be-all-and-end-all of decision intelligence, and everything else in the process supports that decision.

REMEMBER

Whether data driven or decision driven, in both cases humans are the decision makers in this context. It's just that they decide at the tail end of the process in traditional analytics, whereas they decide in the lead position of the decision intelligence process. The starting point for the decision maker matters in terms of the level of control a person has over the impact and value. It's hard to exert much control from the rear.

A history lesson

Disgruntlement with the limitations of traditional data mining is growing. Increasing frustration often leads to both the business side and the AI and IT sides starting to wonder aloud: "What's the point?" But as geeks and businesspeople are wont to do, they realized that there is a point — they just hadn't arrived at it yet.

Eventually, the data driven model was flipped into a decision driven model as people experimented with making the point first and working their way back to the start from there. Decision intelligence is the name of the game where the gamers can all be winners. Now every move made has a point — a point that has value. That's because the point is based on a decision aimed at creating a specific business impact.

The current turn to decision intelligence

Several leading AI luminaries and tech giants have been pioneers in, and first adopters of, decision intelligence. They've already added the title of chief decision scientist to their leadership ranks. One example is Google's eminent chief decision scientist, Cassie Kozyrkov, who spends her days at Google democratizing decision intelligence and developing a more reliable AI approach. She also teaches it to others via conference speeches, YouTube videos, and writings in many online publications.

Kozyrkov appears to embody decision intelligence, partly because of her formal training in economics, mathematical statistics, psychology, and neuroscience. Decision intelligence incorporates all these disciplines — and then some. Although not all who share her title possess the same skill mix, they nevertheless do share strong critical thinking skills as well as a thorough understanding of creative problem-solving strategies, decision theory, and decision science approaches. (For those not familiar with the term, *decision science* focuses on decisions as the unit of analysis; it is the interdisciplinary application of business, math, technology, design thinking, and behavioral sciences to the decision-making process.)

Every day, more leaders are stepping forward to endorse the decision intelligence framework and explain its workings. Many of them work in AI, but others hail from disciplines collectively known as the decision sciences. Business leaders outside the technical domain are also catching on and reveling in their official return to the helm, as opposed to following data's lead (which most never did anyway), and armed with a better strategy. They're also happy about being able to keep their traditional analytics and tools. You don't win battles by limiting your options or abandoning your investments.

Deputizing AI as Your Faithful Sidekick

At its essence, AI automates decisions that are executed rapidly in an exceedingly large number of instances, often simultaneously. You train it by having it work with task-related data sets so that it can recognize what it's looking at in other data sets and learn from the patterns it finds there. Then it makes decisions based on well-defined business rules. (The reality is a good bit more complicated than that, but that's pretty much the gist of it.)

For example, banking institutions use AI to automatically decide which loan applications to approve and which to reject. This is how you can get an answer on your loan application within seconds, no matter how many other people are applying for loans at the same time you are. AI makes these decisions based on the rules it has been given, such as a range of acceptable credit scores, length and types of employment history, items of public record, and other such risk weighting values. AI is able to make such decisions on each individual application, yet at enormous scale and all of it within seconds or minutes. Therefore, borrowers can receive immediate responses to their applications, and lenders can secure more loan deals in minutes than they previously were able to secure over a period of months and at the larger payroll cost of many manhours.

AI is set to continue to serve in this and other automated roles for the foreseeable future. As a technology, it will continue to improve as all technologies do, but placing it within a decision intelligence framework means that its performance will improve exponentially because it is given not only rules to follow but also a target to aim at. Its tasks will be set upon a path of specific actions necessary for creating a specific business impact, and it will faithfully pound away at these tasks until its model decays or someone makes a new model to create another path leading to another targeted impact.

REMEMBER

Other technologies, such as robotic process automation (RPA) and application programming interfaces (APIs), integrate processes. (RPAs are now called virtual workers because they mimic how human workers work, including interacting with user interfaces in the same way.) As RPAs continue to automate processes that were previously difficult to automate, AI can be added to make some automated decisions affecting these processes as well. In other words, the whole of technology engaged in decision making is getting smarter and better and more able to work together.

All this might sound like a setup for a dystopian science fiction movie, but in reality, these developments are nothing to fear. In decision intelligence, whichever technologies you end up using are chosen specifically to augment human soft skills, like creative problem solving, critical thinking, empathy, emotional intelligence, creative design, creative disruption, intuitive intelligence, and intuitive decision making — skills considered nearly impossible to mimic and automate. Even gut instinct can be considered a soft skill, and it too is well out of AI's reach. What ends up happening in decision intelligence is that all human strengths attributable to good judgment and smart decisions are by necessity added to the mix.

REMEMBER

AI is better cast in the role of sidekick, where it augments human decisions rather than dictates or directs them. The same is true of analytics tied to other automated processes as well.

Much of the decision intelligence revolution is happening out of the end user's line of sight, but there's one place where anyone can see the changes unfolding: AI digital assistants such as Google Assistant, Alexa, and Siri. Watch closely as they move from giving you facts in response to your questions to making unprompted recommendations based on your behavior and moods.

Fact reporting such as, "Here are pharmacies near you" or "The name of that song is ABC" will begin to shift to customized and unprompted recommendations. They may look and sound something like this: "XYZ Restaurant has added one of your favorite dishes to its menu. Would you like for me to book the opening in the reservation schedule on Thursday at 7pm and put it on your calendar?" Or, it may say something like this: "Would you like for me to place your favorite coffee order for the pickup window? The one a block from your meeting place has less than a 10 minute wait."

The AI assistant will also produce files for meetings and other handy actions as the user moves through their day. As sidekicks in a user's personal and professional life, the augmented activities will be far more productive than had the human personally tended to all the details and micro decisions.

In digital decision-making, AI will improve at everything it now does — and then some. For example, it will improve at writing algorithms to rapidly meet an organization's or researcher's desired outcomes. That means that, for today and far into the future, AI will be in a position to continue its role as sidekick, producing everything you need to win the day. It's unimaginable that AI won't have some role, small or large, in most Decision Intelligence processes.

Seeing How Decision Intelligence Looks on Paper

Though the decision intelligence framework is perfect for guiding AI to consistently produce business value for you, the methodology can be used with no digital data or machines. For example, you can use AI to make decisions on a spreadsheet, on the back of a napkin, on a single sheet of paper, or even on a wall (using a crayon, of course). That's because the process you use is up to the decision-maker to choose. The Decision Intelligence process itself can be quick and short, or it can be quite complex and take some time to complete. You may want to start with a SWOT table listing the *S*trengths, *W*eaknesses, *O*pportunities, and *T*hreats when making your initial decision. From there, you can determine the steps you need to take to make your decision render a desired impact in the real world.

The process is similar to determining a destination and then mapping out the best route between where you are and where you want to be. It's the impact you desire, however, that will define which route is best. Need to be there fast? The direct route is best. Want to see more along the way or stop at tourist attractions? Then a scenic route is the best way. Want to use your hotel rewards points or your gas rewards card on the trip? Then mapping a route based on the location of certain hotel and gas brands is the best route.

REMEMBER

In decision intelligence, the impact always matters most, for it is the manifestation of your decision.

Working within a decision intelligence framework forces you to become more aware of how the decision-making process works. For example, many of the mental processes you use are intuitive — that's what makes it possible for you to come to conclusions quickly. But make no mistake: Whether you realize it or not, your brain is calculating the same mathematical formulas as a machine would use to help you reach the same conclusions. There's a simple reason for that: Machines copy how people think. As such, machines are definitely the sidekicks in decision intelligence processes, there to assist and augment your efforts.

REMEMBER

Superheroes don't always need a sidekick, and you won't either. Choose the processes, tools, and information according to the needs in executing your decision. Don't default to the technologies and queries with which you're most familiar. The point is not to repeat the same acts, but rather to produce consistent value in personal, professional, or digital decisions.

Tracking the Inverted V

You may be wondering how the processes in decision intelligence differ from those used in data analytics. After all, it's obvious that decisions are also made first when using data analytics in the usual way. For example, someone decides what the business rules are before they apply them to data analytics or AI. Someone also decides what data to use, what data sources to join, and what queries to make. Further, someone decides what projects to launch and whether to send them to production. And so on.

With all these decisions upfront, what does "Put the decision first" in decision intelligence mean? And how does it change anything? It helps to remember that the process in machine-based decision-making is linear, meaning that it moves consecutively from data preparation and selection to algorithm inputs and, finally, to an output. The output is typically an insight or a recommendation delivered as a visualization, as narrative text, or as both, from which a human can decide what action to take. Sometimes, the output is connected to an automated process that then takes an action as directed by the output.

In any case, the path is a straight line.

Now tilt that line upward so that it's the first leg of an inverted V. At the bottom is the starting point, which is the data to be analyzed. At the top is the decision to be made based on the analyses. That's your path upward.

Ignore that path and work your way back down from the decision to the data. Rarely will you follow the same path down. Instead, you'll create a different path that will be more specifically tied to the decision. The two paths together resemble an inverted V.

The first leg of the inverted V begins with mining the data, and then an analysis follows. If you think about it, this process is now defining the decisions you can make. By contrast, in the V's second leg, the decision is defining the data, tools, and queries.

The first leg is a discovery mission. The second is a mission with a purpose.

Which leg do you think will consistently deliver a payload?

And that, my friend, is why and how you put the decision first.

Estimating How Much Decision Intelligence Will Cost You

Ah, yes, the bottom-line question on everyone's lips is cost. Certainly, cost is a major consideration in nearly every business decision. This time, however, it isn't much of an issue. Because decision intelligence is a rethink and not a redo, you likely already have in place many of the technologies and tools you need. (Think of it as leveraging those items to produce a higher return on investment, or ROI, on what you already have.) Of the tools you may not have, many of the products you need offer free versions or at least free trials so that you can see how they work and whether they're a fit for your organization.

That may leave a few tools to buy, depending on your current mix of technologies. All told, it's rarely a huge expense to switch from data mining tactics to decision intelligence.

The following checklist can help you form an idea of some of the technologies commonly used in decision intelligence. That way, you can quickly see what you may need to put on your shopping list — or which functions you might want to hire a third-party who has these things and the experience to use them to do some of this for you.

- **Decision modeling software** is a part of decision management systems that represents business decisions via a standardized notation — often the Business Process Model and Notation (BPMN) standard — that is used by business analysts and subject matter experts (SMEs), rather than developers. Examples include ACTICO Platform, Red Hat Decision Manager, and FICO Decision Management Platform.
- **Business rule management software** to manage the business rules in decision-making. Sometimes these are standalone software products and sometimes they are part of a decision management system as well. Examples include VisiRule, Red Hat Decision Manager, SAS Business Rules Manager, InRule, and DecisionRules.
- **An AutoML stack** *or* another collection of software capable of automating all or part of the building of ML models. AutoML simplifies the machine learning model developer process by automating many of the more laborious steps, such as feature engineering, hyperparameter optimization, and creating the layers in the neural architecture. Don't worry if you don't quite grasp what these automated activities entail because the point in having AutoML is to do all that complicated and time intensive stuff for you. The cool thing is that

while AutoML is a useful tool for data scientists, it's just as useful in democratizing AI. Yes, you too will one day make the AI you want to use in your DI process — by telling AI to make it for you. See, not as hard a concept as you thought. Examples of AutoML vendors include DataRobot, H2O.ai, and Google Cloud AutoML.

- **A good data platform** which is a technology that bundles several big data applications and tools in a single package. Preferably get one that supports both the creation of algorithms and the delivery of transactional data in real-time. Examples include Google Cloud AI platform, RStudio, TensorFlow, and Microsoft Azure.
- **A BI app** with natural language processing, AI assistance, and a built-in visualization tool. Examples include Qlik Sense, Domo, Microsoft Power BI, Yellowfin, Sisense Fusion, Zoho Analytics, and Google Analytics.

Members of your data science team will spend most of their time and effort (at least at first) learning how to capture your newly made decision's requirements using decision model and notation standards such as the Business Process Model and Notation (BPMN), Case Management Model and Notation (CMMN), and/or Decision Model and Notation (DMN) standards.

For decisions where digital data has less of a role or no role, look to the standard tried-and-true array of decisioning tools, like the ones described in this list — and others:

- **Mind mapping tools** are used to create diagrams to visually organize information, typically from brainstorming sessions or collaboration sessions. Examples of mind mapping tools include Coggle, Mindly, MindMup, MindMeister, Scapple, and Stormboard.
- **SWOT tables** consist of four quadrants labeled ***S***trengths, ***W***eaknesses, ***O***pportunities, and ***T***hreats. Users list line items in each quadrant to clarify considerations (and what might be at stake). SWOT tables can be simple or very complex. There are numerous templates available online if you want to use one.
- **Comparison tables** are also known as comparison charts. These are typically line charts, bar charts, pie charts, or other types of charts used to compare or contrast data about any number of things such as data fields (expense categories, for example), competitors, or any other item needing a comparative analysis. Examples of these are everywhere online and off and templates and tools to make such charts are available in visualization and BI tools like Microsoft Power BI, Google Charts, Tableau, Chartist. js, FusionCharts, Datawrapper, Infogram, Canva, and ChartBlocks.

- **Decision trees** depict cascading questions where the answer to one question leads to the formation of the next question. Decision trees are particularly effective in making very complex decisions. They can be simple or very involved depictions, depending on the level of complexity of the problem to be solved. Templates are plentiful online, but there are also tools that will help you make and use them. Examples include Smartdraw's Decision Tree Maker, Lucidchart's Decision Tree Maker, and Creately.
- **Spreadsheets** are those all-too-familiar tools that exist in paper and digital forms, such as Excel and Google Sheets.
- **Paper and pencil** are the tried-and-true standbys. A simple list of pros versus cons on the back of a cocktail napkin has solved many decision dilemmas and they still work today in some instances.

For smaller organizations and start-ups looking to leverage technology in their decision intelligence processes without investing much money, try starting out with a cloud- or browser-based business intelligence app, or one that's embedded in software you already have and use, like Microsoft Power BI, which is embedded in Excel in the Microsoft Office suite. You can find many BI apps with free versions as well. If that's more firepower than you need, check out one or more of the online visualization tools listed above (some are even free!).

REMEMBER

One important caveat: Business intelligence (a BI app that produces reports on current and predicted performance of various aspects of the business based on business data analysis) is *not* the same as a Decision Intelligence process, though BI apps can be used as part of the DI process. A good BI app is simply a quick and reliable way to analyze the data that supports your decision.

The bottom line here is that monetary costs should be relatively small. You may need to spend more on training, however, because your tech people may need additional training on decision theory and the decision sciences — as well as on decision intelligence tactics. Conversely, your business leaders may need that training, too, as well as some training on BI apps to gain a working understanding of data analysis and its full potential.