



Listen to this story

--:--

18:26

Introduction to Decision Intelligence

A new discipline for leadership in the AI era



Cassie Kozyrkov

Follow

Aug 3 · 13 min read ★

Curious to know what the psychology of avoiding lions on the savannah has in common with responsible AI leadership and the challenges of designing data warehouses? Welcome to decision intelligence!



Source: xijian/Getty

Decision intelligence is a new academic discipline concerned with all aspects of selecting between options. It brings together the best of applied data science, social science, and managerial science into a unified field that helps people use data to improve their lives, their businesses, and the world around them. It's a vital science for the AI era, covering

the skills needed to lead AI projects responsibly and design objectives, metrics, and safety-nets for automation at scale.

Decision intelligence is the discipline of turning information into better actions at any scale.

Let's take a tour of its basic terminology and concepts. The sections are designed to be friendly to skim-reading (and skip-reading too, that's where you skip the boring bits... and sometimes skip the act of reading entirely).

What's a decision?

Data are beautiful, but it's decisions that are important. It's through our decisions — our actions — that we affect the world around us.

We define the word “*decision*” to mean any selection between options by any entity, so the conversation is broader than MBA-style dilemmas (like whether to open a branch of your business in London).

It's through our decisions — our actions — that we affect the world around us.

In this terminology, labeling a photo as cat versus not-cat is a decision executed by a computer system, while figuring out whether to launch that system is a decision taken thoughtfully by the human leader (I hope!) in charge of the project.



What's a decision-maker?

In our parlance, a “*decision-maker*” is not that stakeholder or investor who swoops in to veto the machinations of the project team, but rather the person who is responsible for decision architecture and context framing. In other words, a creator of meticulously-phrased objectives as opposed to their destroyer.

What's decision-making?

Decision-making is a word that is used differently by different disciplines, so it can refer to:

- taking an action when there were alternative options (in this sense it's possible to talk about decision-making by a computer or a lizard).
- performing the function of a (human) decision-maker, part of which is *taking responsibility* for decisions. Even though a computer system can execute a decision, it will not be called a decision-maker because it does not bear responsibility for its outputs — that responsibility rests squarely on the shoulders of the humans who created it.

Decision intelligence taxonomy

One way to approach learning about decision intelligence is to break it along traditional lines into its quantitative aspects (largely overlapping with applied data science) and qualitative aspects (developed primarily by researchers in the social and managerial sciences).

Qualitative side: The decision sciences

The disciplines making up the qualitative side have traditionally been referred to as the decision sciences — which I'd have loved for the whole thing to be called (alas we can't always have what we want).





The **decision sciences** concern themselves with questions like:

- “How should you set up decision criteria and design your metrics?” (All)
- “Is your chosen metric incentive-compatible?” (Economics)
- “What quality should you make this decision at and how much should you pay for perfect information?” (Decision analysis)
- “How do emotions, heuristics, and biases play into decision-making?” (Psychology)
- “How do biological factors like cortisol levels affect decision-making?” (Neuroeconomics)
- “How does changing the presentation of information influence choice behavior?” (Behavioral Economics)
- “How do you optimize your outcomes when making decisions in a group context?” (Experimental Game Theory)
- “How do you balance numerous constraints and multistage objectives in designing the decision context?” (Design)
- “Who will experience the consequences of the decision and how will various groups perceive that experience?” (UX Research)
- “Is the decision objective ethical?” (Philosophy)

This is just a small taste... there are many more! This is also far from the complete list of disciplines involved. Think of the decision science side as dealing with decision setup and information processing in its fuzzier storage form (the human brain) rather than the kind that’s neatly written down in semi-permanent storage (on paper or electronically) which we call data.

The trouble with your brain

In the previous century, it was fashionable to praise anyone who stuffed a fat wad of math into some unsuspecting human endeavor. Taking a quantitative approach is usually better than thoughtless chaos, but there’s a way to do even better.

Strategies based on pure mathematical rationality are relatively naïve and tend to underperform.

Strategies based on pure mathematical rationality without a qualitative understanding of decision-making and human behavior can be pretty naïve and tend to underperform relative to those based on joint mastery of the quantitative *and* qualitative sides. (Stay tuned for blog posts on the history of rationality in the social sciences as well as examples from behavioral game theory where psychology beats mathematics.)

Humans are not optimizers, we're satisficers, which is a fancy word for corner cutters.

Humans are not optimizers, we're *satisficers*, which is a fancy word for *corner cutters* who are satisfied with *good enough* as opposed to *perfect*. (It's also a concept that was enough of a shocker to our species arrogance—a punch in the face of rational Man, godlike and flawless — that it was worth a Nobel Prize.)



Image: Source.


In reality, we humans all use cognitive heuristics to save time and effort. That's often a good thing; working out the perfect running path to get away from a lion on the savannah would get us eaten before we've barely started the calculation. Satisficing also reduces the calorie cost of living, which is just as well, since our brains are ridiculously power-hungry devices as it is, gobbling up around a fifth of our energy expenditure despite weighing approximately 3 *lb*. (I bet you weigh more than 15 *lb* in total, right?)

Some of the corners we cut lead to predictably suboptimal outcomes.

Now that most of us don't spend our days running away from lions, some of the corners we cut lead to predictably rubbish outcomes. Our brains aren't exactly, er, optimized for the modern environment. Understanding the manner in which our species turns information into action allows you to use decision processes to protect yourself from the shortcomings of your own brain (as well as from those who intentionally prey on your instincts). It also helps you build tools that augment your performance and adapt your environment to your brain if the poor thing is Lamarckably slow to catch up *a la* Darwin.

If you think that AI takes the human out of the equation, think again!

By the way, if you think that AI takes the human out of the equation, think again! All technology is a reflection of its creators and systems that operate at scale can amplify human shortcomings, which is one reason why developing decision intelligence skills is so necessary for responsible AI leadership. [Learn more here.](#)



Perhaps you're not making a decision

Sometimes, thinking through your decision criteria carefully leads you to realize that there's no fact in the world that would change your mind — you've selected your action already and now you're just looking for a way to feel better about it. That's a useful realization — it stops you from wasting more time and helps you redirect your emotional discomfort while doing what you were going to do anyways, data be damned.

“He uses statistics as a drunken man uses lamp-posts... for support rather than illumination.” -
Andrew Lang


Unless you would take different actions in response to different still-unknown facts, there's no decision here... though sometimes training in decision analysis helps you see those situations more clearly.

Decision-making under perfect information

Now imagine that you'd dealt very carefully with setting up a decision that *is* sensitive to the facts and you can snap your fingers to see the factual information you need for executing your decision. What do you need data science for? Nothing, that's what.

The first order of business should be figuring out how we'd like to react to facts.

There's never anything better than a fact — something you know with certainty (yes, I'm aware there's a gaping relativist rabbit hole here, let's move along) — so we always prefer to make decisions based on facts if we have them. That's why the first order of business should be figuring out how we'd like to deal with facts. Which of the following uses would you want to put your ideal information to?



Your author particularly enjoyed this wall in Jamaica.

What can you do with facts?

- You can use facts to make a single important pre-framed decision. If it's important enough, you'll need to lean heavily on the qualitative side of things to frame your decision wisely. Psychologists know that if you allow yourself to be ambushed by surprise information, it can manipulate you in ways you wouldn't like, so they (and others) have lots to say about how to approach selecting the information you'll accept in advance.
- You can use facts to shore up opinions ("I expect it's sunny outside" becomes "I know it's sunny outside").
- You can use facts to make a single important existence-based decision. Existence-based decisions ("*I just found out there exists a case of ebola right next door, so I'm getting out of here pronto...*") are decisions where the existence of a formerly unknown unknown shakes the foundation of your approach so much that you realize in hindsight that your decision context was sloppily framed.
- You can use facts to automate a large number of decisions. In traditional programming, a human specifies the set of instructions for converting fact inputs into appropriate actions, possibly involving something like a lookup table.
- You can use facts to reveal an automation solution. By seeing the facts about the system, you can write code based on them. This is a better approach to traditional programming than coming up with the structure of a solution by thinking really hard with no information. For example, if you don't know how to convert from Celsius to Fahrenheit, but you could use a dataset to look up the entry in Fahrenheit that goes with the Celsius input... but if you analyze that lookup table itself, you'll discover the formula that connects them. Then you can just code up that formula ("model") to do your dirty work for you and lose the clunky table.
- You can use facts to generate an optimal solution to an automation problem that is perfectly solvable. This is traditional optimization. You'll find many examples in the field of *operations research*, which covers, among other things, how to wrangle constraints to get the ideal outcome, such as the best order in which to complete a series of tasks.
- You can use facts to inspire how you'll approach future important decisions. This is part of analytics, which also belongs in the section on partial information. Hold that thought!
- You can use facts to take stock of what you're dealing with. This helps you understand the kinds of inputs you have available for future decisions and design how to curate your information better. If you've just inherited a big, dark (data)


warehouse full of potential ingredients, you won't know what's inside until someone looks at it. Luckily, your analyst has a flashlight and rollerblades.


- You can use facts sloppily to make unframed decisions. This is efficient when decisions have sufficiently low stakes that they do not warrant the effort required to approach them carefully, such as, *"What should I eat for lunch today?"* Attempting to be rigorous all the time on all decisions gives suboptimal long-run / lifetime outcomes and falls into the category of pointless perfectionism. Save your effort for the situations that are important enough for it, but please don't forget that even if it's efficient to use a low-quality low-effort approach, the optimal decision approach is still of low quality. You shouldn't thump your chest or be overconfident when that's your method... If you cut corners, you're holding something flimsy. There are situations where flimsy gets the job done, but that doesn't suddenly make your conclusion sturdy. Don't lean on it. If you want high-quality decision-making, you need a more rigorous approach.

With training in the decision sciences, you learn to reduce the effort that it takes to make rigorous fact-based decisions, which means that the same amount of work now gets you higher-quality decision-making across the board. This is a very valuable skill, but it takes lots of work to hone it. For example, students of behavioral economics form the habit of setting decision criteria in advance of information. Those of us who took a beating from sufficiently demanding decision science training programs can't help but ask ourselves, for example, what the maximum that we'd pay for a ticket is BEFORE we look up the price.

Data collection and data engineering

If we had the facts, we'd be done already. Alas, we live in the real world and often we must work for our information. Data engineering is a sophisticated discipline centered on making information available reliably at scale. In the way that going to the grocery store for a pint of ice cream is easy, data engineering is easy when all available relevant information fits in a spreadsheet.





Things get tricky when you start asking for the delivery 2 million tons of ice cream... where it's not allowed to melt! Things get even trickier if you have to design, set up, and maintain a huge warehouse and you don't even know what the future will ask you to store next — maybe it's a few tons of fish, maybe it's plutonium... good luck!

It's tricky to set up a warehouse when you don't even know what you'll be asked to store next week— maybe it's a few tons of fish, maybe it's plutonium... good luck!

While data engineering is a separate sister discipline and key collaborator to decision intelligence, the decision sciences include a strong tradition of expertise involved in advising the design and curation of fact collection.

Quantitative side: Data science

When you've framed your decision and you look up all the facts you need, using a search engine or an analyst (performing the role of a human search engine for you), all that's left is to execute your decision. You're done! No fancy data science needed.

What if, after all that legwork and engineering jiu-jitsu, the facts delivered are not the facts you ideally need for your decision? What if they're only partial facts? Perhaps you want tomorrow's facts, but you only have the past to inform you. (It's so annoying when we can't remember the future.) Perhaps you want to know what all your potential users think of your product, but you can only ask a hundred of them. Then you're dealing with uncertainty! What you know is not what you wish you knew. Enter data science!

Data science gets interesting when you're forced to make leaps beyond the data... but do be careful to avoid an Icarus-like splat!

Naturally, you should expect your approach to change when the facts you have aren't the facts you need. Maybe they're one puzzle piece of a much bigger puzzle (as with a sample from a larger population). Maybe they're the wrong puzzle, but the best you have (as with using the past to predict the future). Data science gets interesting when you're forced to make leaps beyond the data... but do be careful to avoid an Icarus-like splat!

- You can use partial facts to make a single important pre-framed decision with statistical inference, supplementing the information you have with assumptions to see if you should change your action. This is Frequentist (classical) statistics.
- You can use partial facts to reasonably update opinions into more informed (but still imperfect and personal) opinions. This is Bayesian statistics.
- Your partial facts may turn out to contain facts about existence, which means you could use them in hindsight for existence-based decisions (see above).
- You can use partial facts to automate a large number of decisions. That's traditional programming using something like a lookup table where you convert something you haven't seen before into the closest thing you that you have, then proceed as usual. (That's what k-NN does in a nutshell... and it usually works better when that nutshell has more things in it.)
- You can use partial facts to inspire an automation solution. By seeing the partial facts about the system, you can still write code based on what you're seeing. This is analytics.
- You can use partial facts to generate a decent solution to an imperfectly solvable automation problem so you don't have to come up with it yourself. This is machine learning and AI.
- You can use partial facts to inspire how you'll approach future important decisions. This is analytics.
- You can use partial facts for understanding what you're dealing with (see above) and to accelerate the development of an automation solution with advanced analytics, for example by inspiring new ways to blend information together to make useful model inputs (the jargon for this is "*feature engineering*") or new methods to try in an AI project.
- You can use partial facts sloppily to make unframed decisions, but be aware that the quality is even lower than when you use facts sloppily, because what you actually know is one step removed from what you wish you knew.

For all of these uses, there are ways to integrate wisdom from a variety of previously-siloed disciplines to approach decision-making more effectively. *That's what decision*

intelligence is all about! It brings together diverse perspectives on decision-making which make all of us stronger, together, and gives them a new voice that's free of the traditional constraints of their originating fields of study.



To return to the kitchen analogy for AI, if research AI is building microwaves and applied AI is using microwaves, decision intelligence is using microwaves safely to meet your goals and using something else when you don't need a microwave. The goal (objective) is always the starting point for decision intelligence.

If you're curious to read more, most of my articles here on Medium.com have been written from a decision intelligence perspective. My guide to starting AI projects is probably the least subtle, so I'd recommend diving in there if you haven't already chosen your own adventure by following the links in this article.

[Data Science](#)[Artificial Intelligence](#)[Technology](#)[Leadership](#)[Tds Narrated](#)[About](#)[Help](#)[Legal](#)