



Harvard Business Review

REPRINT H050C4
PUBLISHED ON HBR.ORG
JULY 08, 2019

ARTICLE INNOVATION


What AI-Driven Decision Making Looks Like

by Eric Colson

INNOVATION

What AI-Driven Decision Making Looks Like

by Eric Colson
JULY 08, 2019



DANIEL SAMBRAUS/EYEEM/GETTY IMAGES

Many companies have adapted to a “data-driven” approach for operational decision-making. Data can improve decisions, but it requires the right processor to get the most from it. Many people assume that processor is human. The term “data-driven” even implies that data is curated by — and summarized for — people to process.

But to fully leverage the value contained in data, companies need to bring artificial intelligence (AI) into their workflows and, sometimes, get us humans out of the way. We need to evolve from data-driven to AI-driven workflows.

Distinguishing between “data-driven” and “AI-driven” isn’t just semantics. Each term reflects different assets, the former focusing on data and the latter processing ability. Data holds the insights that can enable better decisions; processing is the way to extract those insights and take actions. Humans and AI are both processors, with very different abilities. To understand how best to leverage each its helpful to review our own biological evolution and how decision-making has evolved in industry.

Just fifty to seventy five years ago human judgment was the central processor of business decision-making. Professionals relied on their highly-tuned intuitions, developed from years of experience (and a relatively tiny bit of data) in their domain, to, say, pick the right creative for an ad campaign, determine the right inventory levels to stock, or approve the right financial investments. Experience and gut instinct were most of what was available to discern good from bad, high from low, and risky vs. safe.

A Decision-Making Model Based on Human Judgment



Source: Eric Colson



It was, perhaps, all too human. Our intuitions are far from ideal decision making instruments. Our brains are afflicted with many cognitive biases that impair our judgement in predictable ways. This is the result of hundreds of thousands of years of evolution where, as early hunter-gatherers, we developed a system of reasoning that relies on simple heuristics — shortcuts or rules-of-thumb that circumvent the high cost of processing a lot of information. This enabled quick, almost unconscious decisions to get us out of potentially perilous situations. However, ‘quick and almost unconscious’ didn’t always mean optimal or even accurate.

Imagine a group of our hunter-gatherer ancestors huddled around a campfire when a nearby bush suddenly rustles. A decision of the ‘quick and almost unconscious’ type needs to be made: conclude that the rustling is a dangerous predator and flee, or, inquire to gather more information to see if it is potential prey – say, a rabbit, that can provide rich nutrients. Our more impulsive ancestors—those that decided to flee– survived at a higher rate than their more inquisitive peers. The cost of flight and losing on a rabbit was far lower than the cost of sticking around and risking losing life to a predator. With such asymmetry in outcomes, evolution favors the trait that leads to less costly consequences,

even at the sacrifice of accuracy. Therefore, the trait for more impulsive decision-making and less information processing becomes prevalent in the descendant population.

In modern context, survival heuristics become myriad cognitive biases pre-loaded in our inherited brains. These biases influence our judgment and decision-making in ways that depart from rational objectivity. We give more weight than we should to **vivid or recent events**. We **coarsely classify subjects** into broad stereotypes that don't sufficiently explain their differences. We **anchor on prior experience** even when it is completely irrelevant. We tend to conjure up specious explanations for events that are really **just random noise**. These are just a few of the **dozens of ways** cognitive bias plagues human judgment and for many decades, it was the central processor of business decision-making. We know now that relying solely on human intuition is inefficient, capricious, fallible and limits the ability of the organization.

Data-Supported Decision Making

Thank goodness, then, for data. Connected devices now capture unthinkable volumes of data: every transaction, every customer gesture, every micro- and macroeconomic indicator, all the information that can inform better decisions. In response to this new data-rich environment we've adapted our workflows. IT departments support the flow of information using machines (databases, distributed file systems, and the like) to reduce the unmanageable volumes of data down to digestible summaries for human consumption. The summaries are then further processed by humans using the tools like spreadsheets, dashboards, and analytics applications. Eventually, the highly processed, and now manageably small, data is presented for decision-making. This is the "data-driven" workflow. Human judgment is still the central processor, but now it uses summarized data as a new input.

A Decision-Making Model That Utilizes Summarized Data



Source: Eric Colson



While it's undoubtedly better than relying solely on intuition, humans playing the role of central processor still creates several limitations.

1. We don't leverage all the data. Summarized data can obscure many of the insights, relationships, and patterns contained in the original (big) data set. Data reduction is necessary to accommodate the throughput of human processors. For as much as we are adept at digesting our surroundings, effortlessly processing vast amounts of ambient information, we are remarkably limited when it comes to processing the structured data manifested as millions or billions of records. The mind can handle sales numbers and average selling price rolled up to a regional level. It struggles or shuts down once you start to think about the full distribution of values and, crucially, the relationships between data elements—information lost in aggregate summaries but important to good decision making. (This is not to suggest that data summaries are not useful. To be sure, they are great providing basic visibility into the business. But they will provide little value for use in decision-making. Too much is lost in the preparation for humans.) In other cases summarized data can be outright misleading. Confounding factors can give the appearance of a positive relationship when it is actually the opposite (see [Simpson's](#) and other paradoxes). And once data is aggregated, it may be impossible to recover contributing factors in order to properly control for them. (The best practice is to use randomized controlled trials, i.e. A/B testing. Without this practice, even AI may not be able to properly control for confounding factors.) In short, by using humans as central processors of data, we are still trading off accuracy to circumvent the high cost of human data processing.
2. Data is not enough to insulate us from cognitive bias. Data summaries are directed by humans in a way that is prone to all those cognitive biases. We direct the summarization in a manner that is intuitive to us. We ask that the data be aggregated to segments that we feel are representative archetypes. Yet, we have that tendency to coarsely classify subjects into broad stereotypes that don't sufficiently explain their differences. For example, we may roll up the data to attributes such as geography even when there is no discernible difference in behavior between regions. Summaries also can be thought of as a "coarse grain" of the data. It's a rougher approximation of the data. For example, an attribute like geography needs to be kept at a region level where there are relatively few values (i.e., "east" vs. "west"). What matters may be finer than that — city, ZIP code, even street-level data. That is harder to aggregate and summarize for human brains to process. We also prefer simple relationships between elements. We tend to think of relationships as linear because it's easier for us to process. The relationship between price and sales, market penetration and conversion rate, credit risk and income — all are assumed linear even when the data suggests otherwise. We even like to conjure up elaborate explanations for trends and variation in data even when it is more adequately explained by natural or random variation.


Alas, we are accommodating our biases when we process the data.

Bringing AI into the Workflow

We need to evolve further, and bring AI into the workflow as a primary processor of data. For routine decisions that only rely on structured data, we're better off delegating decisions to AI. AI is less prone to human's cognitive bias. (There is a very real risk of using biased data that may cause AI to find specious relationships that are unfair. Be sure to understand how the data is generated in addition to how it is used.) AI can be trained to find segments in the population that best explain variance at fine-

grain levels even if they are unintuitive to our human perceptions. AI has no problem dealing with thousands or even millions of groupings. And AI is more than comfortable working with nonlinear relationships, be they exponential, power laws, geometric series, binomial distributions, or otherwise.

A Decision-Making Model That Utilizes AI



Source: Eric Colson



This workflow better leverages the information contained in the data and is more consistent and objective in its decisions. It can better determine which ad creative is most effective, the optimal inventory levels to set, or which financial investments to make.

While humans are removed from this workflow, it's important to note that mere automation is not the goal of an AI-driven workflow. Sure, it may reduce costs, but that's only an incremental benefit. The value of AI is making better decisions than what humans alone can do. This creates step-change improvement in efficiency and enables new capabilities.


Leveraging both AI and Human processors in the workflow

Removing humans from workflows that only involve the processing of structured data does not mean that humans are obsolete. There are many business decisions that depend on more than just structured data. Vision statements, company strategies, corporate values, market dynamics all are examples of information that is only available in our minds and transmitted through culture and other forms of non-digital communication. This information is inaccessible to AI and extremely relevant to business decisions.

For example, AI may objectively determine the right inventory levels in order to maximize profits. However, in a competitive environment a company may opt for higher inventory levels in order to provide a better customer experience, even at the expense of profits. In other cases, AI may determine that investing more dollars in marketing will have the highest ROI among the options available to the company. However, a company may choose to temper growth in order to uphold

quality standards. The additional information available to humans in the form of strategy, values, and market conditions can merit a departure from the objective rationality of AI. In such cases, AI can be used to generate possibilities from which humans can pick the best alternative given the additional information they have access to. The order of execution for such workflows is case-specific. Sometimes AI is first to reduce the workload on humans. In other cases, human judgment can be used as inputs to AI processing. In other cases still, there may be iteration between AI and human processing.

A Decision-Making Model That Combines the Power of AI and Human Judgment



Source: Eric Colson



The key is that humans are not interfacing directly with data but rather with the possibilities produced by AI's processing of the data. Values, strategy and culture is our way to reconcile our decisions with objective rationality. This is best done explicitly and fully informed. By leveraging both AI and humans we can make better decisions than using either one alone.

The Next Phase in our Evolution

Moving from data-driven to AI-driven is the next phase in our evolution. Embracing AI in our workflows affords better processing of structured data and allows for humans to contribute in ways that are complementary.

This evolution is unlikely to occur within the individual organization, just as evolution by natural selection does not take place within individuals. Rather, it's a selection process that operates on a population. The more efficient organizations will survive at a higher rate. Since it's hard for mature companies to adapt to changes in the environment, I suspect we'll see the emergence of new companies that embrace both AI and human contributions from the beginning and build them natively into their workflows.

Eric Colson is Chief Algorithms Officer at Stitch Fix. Prior to that he was Vice President of Data Science and Engineering at Netflix. [@ericcolson](#)
