


Harvard Business Review

Data Scientists Don't Scale

BY STUART FRANKEL



WIPRO
Applying Thought

brought to you by

Data Scientists Don't Scale

BY STUART FRANKEL

ABOUT THE AUTHOR

Stuart Frankel is the CEO and a co-founder of Narrative Science, a company working on advanced natural language generation for the enterprise.

Big data is about to get a big reality check. Our ongoing obsession with data and analytics technology, and our reverence for the rare data scientist who reigns supreme over this world, has disillusioned many of us. Executives are taking a hard look at their depleted budgets — drained by a mess of disparate tools they've acquired and elusive "big insights" they've been promised — and are wondering: "Where is the return on this enormous investment?"

It's not that we haven't made significant strides in aggregating and organizing data, but the big data pipedream isn't quite delivering on its promise. Despite massive investments in technology to store, analyze, report, and visualize data, employees are still spending untold hours interpreting analyses and manually reporting the results. To solve this problem and increase utilization of existing solutions, organizations are now contemplating even further investment, often in the form of \$250,000 data scientists (if all of these tools we've purchased haven't completely done the trick, surely this guy will!). However valuable these PhDs are, the organizations that have been lucky enough to secure these resources are realizing the limitations in human-powered data science: it's simply not a scalable solution. The great irony is of course that we have more data and more ways to access that data than we've ever had; yet we know we're only scratching the surface with these tools.

A few innovative executives understand this and have sought scalable, automated solutions that interpret data, unlock hidden insights, and then provide answers to ongoing business problems. Artificial intelligence (AI) is beginning to transform data and analysis

into relevant plain English communication. AI is shortening employees' data comprehension-to-action time through comprehensive, intuitive narratives.

Any organization where employees are spending valuable time on manual, repetitive tasks is a prime case for intelligent automated solutions. There may be no better example of this than the financial services industry.

Take for example the necessary but incredibly manual process of producing performance reporting for mutual funds. Typically, marketing teams toil every quarter to document portfolio performance and add commentary (see an example here). Today, some funds are using advanced natural language generation (Advanced NLG) platforms, powered by AI, to automatically write these reports in mere seconds. (My company, Narrative Science, works with multiple financial services clients to do this sort of work.)

These are not simple, static reports, but are data-driven, complex, and dynamic, reflecting the brand voice most appropriate for the firm. The reporting incorporates disparate data sources and performs real-time analysis on portfolio performance. These systems examine the facts, determine which of these facts are most notable, and output these facts as readable narrative text. For instance, the report could be a bulleted list of key findings for portfolio managers who just need the facts, or it could be a lengthy detailed summary to meet an investor's needs.

To do this work, the system starts with the goals of the report (e.g., did this portfolio outperform the benchmark?). Once the business need is clear, the system pulls in the necessary information (portfolio results compared to the benchmark), performs the

The commonality across all of these new technologies is that they offer something additional humans cannot provide: the power of scale.

relevant data analysis (cohort comparison), and finally decides which data is required to complete the task (portfolio attribution data, returns data, expected values, actual values, and any confidence weights on the expected values).

Once the report is created, it updates automatically as the data does, eliminating the need to redo the analysis and reporting each time there is a change.

The use cases for these systems are countless, but they all start with the question: What do I want to communicate that currently requires a significant amount of time and energy to analyze, interpret, and share?

Take medical billing. AI scours thousands of billing records across hundreds of hospitals and generates narrative reports that immediately provide the desired analysis. These reports can highlight changes to a hospital group's insurance-provider portfolio or changes in demographic mix that are affecting revenue and cash flow, while simultaneously identifying room for growth by suggesting changes to a doctor's workload.

There are also a number of examples where AI solutions are improving customer experience. AI is the first technology to make personalized, "audience of one" communication a reality. Companies can speak to each customer as if they are the only customer, offering personalized reports that seamlessly integrate with consumer applications and websites.

Wealth management is starting to see this benefit. Many of us have now heard of "Robo-advisors," the automated financial advisor that can offer a low-cost alternative to expensive, human advisors. AI is being embedded into existing advisory platforms,

delivering personalized portfolio reviews and recommendations in natural language to customers who are unsure as to what their charts and graphs actually mean. Unlike the Robo-advisors that operate in a "black box," these AI systems provide transparency by communicating what has happened and offering recommendations in plain English that people can immediately understand.

The commonality across all of these new technologies is that they offer something additional humans cannot provide: the power of scale. Organizations that do not have a strategic initiative to regularly and organically engage with its customers will be at a serious disadvantage. Soon, AI-driven engagement models that interpret data and intuitively interact with clients will be the norm.

In the near-term, the adoption of AI within business intelligence platforms and customer-facing applications will accelerate. When a business user receives a dashboard, he will also receive an accompanying narrative to explain the insight within the visualization, easing consumption and quickening the pace of decision-making. Eventually, AI will offer even more complex analysis and advice. When a salesperson checks her pipeline status within her CRM application, she'll be able to ask, "How am I performing compared to my colleagues this quarter?" and "Where should I focus my efforts to ensure I reach my quota?" The application will be able to respond immediately, with an explanation and sound advice she can understand and act on, in a back-and-forth, conversational manner.

The key to all of this is the intersection of AI and advanced natural language generation. We're at the beginning of the next phase of big

data, a phase that will have very little to do with data capture and storage and everything to do with making data more useful, more understandable and more impactful.

Big Data: Fueling the Artificial Intelligence Revolution

An Interview with K R Sanjiv, Chief Technology Officer, Wipro

Why aren't companies getting more value out of their investment in Big Data?

Today, majority of the organizations understand the kind of data they might require and have the data sources as well in place. They have built the infrastructure and invested in platforms and tools to analyze the captured data. But the important question is how does one use this data to gain insights; useful information that can be used to take actions. Data itself is not of much value if you cannot make it actionable. And right now, we have all this data around us, but the monetization is too low. Enterprises often ignore the monetary value of data.

There are many reasons for this. People have not put optimal solutions in place. There are legal and privacy issues. But the real issue is that people have not built a strong layer on top of their big data platform which would allow them to realize the real value.

Analytics and business intelligence today is answering all the questions that arise in a scenario. If you don't ask the right questions you don't get the right insights, or if you ask them too late, the opportunity to act on that insight is missed. So to get more value, the whole pattern must be switched, in order to ensure that the use of data becomes proactive and real time. You have to automate ways to extract insights from data and create actions for business users.

How can artificial intelligence help companies with the opportunities Big Data presents?

The convergence of Big Data and artificial intelligence, along with our increased order of computing power, is creating value. Together we are now able to learn from data, detect patterns, and propagate answers in a timely fashion that has real impact. In some ways Big Data is the oil that fuels the engines of artificial intelligence.

We have the ability now to automate a great deal of this modeling and analysis that otherwise would require a lot of dedicated resources. For instance, automation can be used in resolving people's problems at a helpdesk. There are thousands of calls that come in each day, all digitized. It is a phenomenal amount of data, but a fair amount of these problems are repetitive in nature and solving these require a lot of manual repetitive work. Using artificial intelligence, machines can understand the intent of the problem raised to the help desk, review it against a repository of information as to how that problem was solved in the past, and then identify the best answer within seconds.

Another good example would be data from ATM fraud. A person withdrawing a small amount at an ATM would be a normal pattern, but if suddenly they are withdrawing three times the usual amount, this anomaly signals possible fraud and businesses can respond quickly. To be able to do that, modeling is required, a layer built on top of the data to extract information and then the answer—the right action for the company—is propagated in a proactive manner. And this is just one of the many potential use cases where real opportunity lies.

How should organizations get started with artificial intelligence application?

The key is to define the right process that is amenable for disruption. Any process that demands a lot of manual activity but is digitized to a large extent is a low-hanging fruit. Enterprises should focus on use cases where artificial intelligence can improve the quality of decision-making. Enterprises need to start building the automation and feedback engines that can improve learning from the available enterprise big data. An ideal way to start is through pilot implementations; enterprises can then realize the impact and utility of artificial intelligence.

brought to you by

