

What the tech analysts aren't telling you about AI-based speech analytics

Authored by: Matt Dixon, Chief Product and Research Officer
Published: July 2018

The people side of the AI is more important than most companies realize; artificial intelligence only works when coupled with human intelligence to inform, interpret and leverage it.

The one article every potential buyer of AI-powered speech analytics should read before they do anything—before they talk to a vendor or issue an RFP and certainly before they write a check to anybody—is “[Collaborative Learning: Humans and AI are Joining Forces](#),” a piece that appears in the upcoming July-August 2018 issue of *Harvard Business Review*. In fact, you should read this before you pick up a tech analyst report on the space.

The article, written by Accenture’s James Wilson and Paul Daugherty, is a powerful articulation of what most companies miss when they go shopping for an AI solution to improve sales, service, customer experience or compliance effectiveness: artificial intelligence only works when coupled with human intelligence to inform, interpret and leverage it.

As the authors write, “Many companies have used AI to automate processes, but those that deploy it mainly to displace employees will see only short-term productivity gains. In our research involving 1,500 companies, we found that firms achieve the most significant performance improvements when humans and machines work together.”

Recently, our team attended the Customer Contact Week conference in Las Vegas and we were struck by the palpable interest in AI-based speech analytics from customer service and CX leaders. But, as exciting as this was to be a part of, we were also troubled by the level of understanding about what’s really required to get value out of an AI investment.

This was apparent in the questions these leaders were asking vendors and—almost more importantly—the questions they *weren’t* asking. To be blunt, the questions they were asking seemed to be straight out of a tech analyst report: “Does your platform deliver real-time analytics?,” “How secure is the data?,” “Who’s your cloud provider?,” “Do you integrate with [fill in the blank]?,” “What’s on your technology roadmap?”

To be fair, these are all important questions. But they miss the point. A focus alone on tech functionality—on just the “artificial” part of AI—fails to address arguably the hardest part of getting AI right over the long term: the “intelligence” part rises and falls on an algorithm that, by nature and design, is incomplete.

A focus alone on tech functionality—on just the “artificial” part of AI—fails to address arguably the hardest part of getting AI right over the long term: the “intelligence” part rises and falls on an algorithm that, by nature and design, is incomplete.

[Robin Bordoli](#), CEO of the AI venture Figure Eight—a company that recently raised \$20M from Industry Ventures, Microsoft, Salesforce and others—puts it really well: “We’re at the beginning of a Cambrian explosion in AI applications within the enterprise,” he writes. “The bottleneck for the large-scale adoption of machine learning still remains the availability of high-quality training data and human-in-the-loop workflows to handle the failure states of the algorithm. A machine learning model without [this] is like a rocket ship with a large engine but no fuel and no navigation system. It won’t reach escape velocity, nor will it achieve the trajectory to land on its intended target.”

The machine, quite simply, must continue to learn. How the machine learns—the basis of current knowledge, who teaches it and how it’s taught—remains the key to actually get any business value out of the technology. Any vendor that obfuscates this point is doing you a disservice.

To further illustrate the point, we recently [wrote](#) about how companies shouldn’t be in such a rush to automate quality assurance and that, instead, they should use the power of AI to first solve problems previously out of their reach—namely, to understand, from millions of customer interactions, the precise things that actually *drive* quality.

In that piece, we shared our experience with one Tethr customer we helped identify a new way to measure a specific type of agent advocacy language because of the impact it has on sales conversion. The specific phrasing—when agents say “I recommend this” or “Here’s what I would do”—had a greater lift on sales conversion than actually asking for the sale itself, a common and understandable piece of agent advice.

But where did this insight come from and how would the AI know to look for it? After all, there are millions (perhaps billions) of things that could drive quality... so, how did the AI know to look here? In short, it didn’t. AI needs an Intellectual Property (IP) roadmap in order to be effective. For example, our IP roadmap is based on a robust, 10-year research study I was a part of at CEB, now Gartner (for more on that, check out [The Effortless Experience](#) and our HBR article, [“Stop Trying to Delight Your Customers”](#)). That study—the foundation of which was built on mountains of research in behavioral economics and human psychology (such as [Influence](#) by Robert Cialdini)—demonstrated the impact of language techniques like advocacy on customer outcomes.

In our customer example, once we identified the specific language that impacted customer outcomes, how did we know that was a causal relationship, not just a spurious correlation? Again, it was the interpretation of people that gave us confidence that this language was, in fact, producing this effect. Namely, we knew from other research (see [The Challenger Sale](#) and our 2017 HBR article, [“Kick-Ass Customer Service”](#)) that best sellers take control of sales interactions. They are authoritative—not asking customers what’s “keeping them up at night,” but confidently telling customers what should be keeping them up at night. They are prescriptive controllers.

This is a great example of how AI should work. The AI roadmap provides a solid foundation of knowledge to start; the machine draws conclusions at a rapid clip not because of tech functionality, but because people taught it to look for

specific correlations. Compare this to past insight-creating methods: when Neil Rackham did the research in the 1970s that went into his seminal book, *SPIN Selling*, he managed a team of 30 researchers who spent 12 years analyzing 35,000 sales calls. When we did the research in 2008 that went into *The Challenger Sale*, our team of 12 researchers took more than two years to study 6,000 sales people and 5,000 customers through a series of surveys and interviews. And the insight I shared above? It took two people a matter of hours to find it, using a data set of millions of sales interactions.

The best AI-based speech analytics platform in the world—no matter the cloud platform used, the robustness of the technology roadmap or the bevy of APIs available—is incapable of producing insights like this. Only people who know how to train AI and interpret the results can do this.

So, what does this mean for companies looking at AI-based speech analytics? In short, it means you need to pay as much, or more, attention to the people side of the AI than the technology side.

When you talk to an AI-based speech analytics vendor, here are the questions you should be asking them:

- **Will you equip my team with the “fusion skills” they will need to get value out of the AI platform?**

Smith and Daugherty describe “fusion skills” as “those that enable [employees] to work effectively at the human-machine interface.” A good AI speech analytics provider won’t require you to hire your own data scientists, business analysts or consultants to train the machine, tune it or leverage it to produce insights. Instead, they invest, up front, in teaching your team the fusion skills necessary to do this on their own. With the proper amount of training, anybody in your organization should be able to build machine learning training sets, leverage them against your voice data and produce results that can be interpreted and explained.



- **How is your machine trained?**

AI-based speech analytics solutions leverage machine learning and machine learning works off of training sets, also known as libraries. Think of this as the “navigation system” for the AI. When you buy one of these solutions, it will come with a host of pre-built training sets. These reflect the things that the vendor thinks are the important for you to look for in your voice data. But what informed those decisions? Again, there are millions of possible things to look for, so what drove the vendor to select the specific areas of focus in their libraries? Was it based on a robust body of research or did the vendor simply “mechanize” and automate what its customers were already doing? A dead giveaway is when you see libraries that reflect conventional wisdom—things like “friendliness” or “empathy” (neither of which, it turns out, drive higher quality customer outcomes). If the libraries aren’t built on a robust body of research, ask what statistical impact the categories have on things like sales, service quality or compliance. If they can’t tell you, then they’re likely built on guesswork, hunches and gut instinct. As the saying goes, “garbage in, garbage out.” Don’t use AI to automate mediocrity.



- **Can we build our own libraries and dashboards?**

No vendor will provide a complete set of libraries to you—for the simple reason that your company and your products are different. Even if the vendor has libraries already built for your industry, they still need to be tuned for your specific markets, products and customers. You should assume that pre-built libraries—if relevant—still need about 20-30% tuning to be effective and accurate. And then there are the libraries that are completely unique to your company—or perhaps reflect the creativity and hypotheses of your team about what they think is driving sales, service or compliance. And, of course, once you publish these tuned libraries and start running them against your voice data, you’ll want to dashboard the results. Ask your vendor who does the building and tuning of the libraries. Do they do that or will they teach you to do it? If your vendor tells you they’ll “work with you” to do it, that’s code for you needing to hire your own folks or pay them professional services fees to do the work.

- **What’s your IP roadmap?**

Every technology vendor worth their salt will be able to show you their product roadmap. But what is the roadmap to enable the machine to learn? If your vendor is only showing you a roadmap of future software functionality and technical capabilities, then you’re only getting half of the story. AI providers should be able to also produce an IP roadmap for (and with) their customers. What is an IP roadmap? It’s an overview of the research they’ve taught the machine, the research they plan to continue, the hypotheses they and other CX/Service leaders are actively testing to build new libraries—all with an eye toward making the algorithm that much more complete. If the vendor only shows you a technology roadmap, what they’re really saying is that the “intelligence” part of the AI is on you.

- **Who are your researchers and what are their backgrounds?**

The team that will execute on the IP roadmap—the research that informs new library and algorithm creation as well as content like playbooks, tools and frameworks to interpret and take action on the output from the AI—should have deep research experience and significant domain experience. If you’re looking at AI-based speech analytics for sales, the research team should know the sales space and the body of research that exists in the area of sales effectiveness. The same is true of service, customer experience and compliance. For your vendor to have AI Ph.D.’s and data scientists capable of coding in Python or R is awesome—but are they complemented by researchers with deep social science training and expertise in the domains you’re looking to apply the AI in? Because that’s really awesome. And, if not, is the vendor assuming you’ll just know what to look for and how to look for it?

- **Will you just tell us what’s broken or help us fix it?**

When properly tuned, AI will help you identify problems—lots of them—in your business. It turns out machine learning is really good at that. But what happens next? Does the vendor leave you hanging or do they help you think drive action from the insight? Do they offer you playbooks, tools and templates to actually do something with the insight? What about chances to engage with other users of their solution to learn from their experience? If not, they’re assuming you’ll do all of this work—or maybe hire a consultant to do it for you. Without some support to go from insight to action, you might get smarter, but there’s a good chance you’ll get no better effectiveness using the traditional QA method since agents are fielding hundreds of calls per week. There are much more complex approaches of using AI to enable better HI. For example, Tethr’s AI spotted a pattern with a customer where a specific message used in a sales call converted to a purchase, but within 30 days the same customers called customer care to cancel. Armed with these insights, management was able to adjust the messaging and reduce post-purchase customer churn.

You will find none of these are questions addressed in a tech analyst report, but these are actually the critical questions to address if you want AI to deliver the business impact it’s capable of.

So, use the tech analyst’s report to help you understand the “artificial” part of AI—the technology—but use the set of questions and articles recommended above to assess the “intelligence” of the overall solution.



CONTACT US TO LEARN MORE:

TETHR.COM
(512) 910-4440
SALES@TETHR.COM