

The Insightful Leader Podcast Transcript

When AI Leaves the Lab

Jessica LOVE: Adam Pah is pointing his smartphone camera at himself, trying to get a good selfie.

PAH: That's an awful photo!

LOVE: Why is he doing this? Because he's about to test out a cutting edge piece of artificial intelligence. It's a new web app from Olay, the skin care company. The app uses an algorithm to analyze users' faces, and then another algorithm to provide them with personalized product recommendations.

PAH: I have to snap my selfie, I have to tell them about my skin, and then I get my results. This is going to be interesting!

LOVE: Pah is a clinical professor at Kellogg. And he's excited to be trying out this app *not* because he's especially concerned about his skin care regimen. Truth is, this app wasn't really designed for him. It's meant for women. But as someone who studies and teaches about the uses of AI in business, Pah like to give new apps like this a spin. And he thinks that this new one from Olay is pretty exciting. Because while the business world has been talking about the potential of AI for years, few companies have really harnessed that potential. So apps like this indicate that, at last, AI is moving out of the lab.

PAH: We're finally getting to the point that companies are making something, actually, with it. It's *just* starting to happen, finally, as enough companies have enough understanding and expertise to execute and launch these products.

LOVE: But Pah points out that these products are still few and far between. Throughout the business world, leaders are missing opportunities because they don't fully understand how AI works, or what it takes to get it right. But Pah thinks, if these leaders have just a little more context about where AI stands...

PAH: They can really understand—not just for today, but for tomorrow—how they can take this and bring it into organizations.

[musical interlude]

LOVE: Welcome to *The Insightful Leader* from the Kellogg School of Management. Adam Pah recently contributed to a new [ebook](#) on analytics and AI for business leaders, available now from Kellogg Insight. So today on the podcast, he shares three surprising things about AI in the product space—problems and trends that savvy business leaders should be aware of. He'll use the Olay app to explain why it's actually getting harder to get good data, why too much data can be a bad thing, and how companies are starting to use AI to solve customers' real problems.

PAH: Not just tying it into the product line, but tying it in to the user experience.

LOVE: The Olay app asks you to snap a photo of yourself, so that its AI can analyze your face. But first you have to agree to some terms and conditions. Now, this is nothing unusual for an app. But there's a line in Olay's terms that's actually pretty revealing. You have to give the company permission to use your data for what it calls "internal research purposes."

PAH: "...including the ability to update our Olay skin advisor algorithm to improve our skin analysis and product recommendations." Eh, I'll take that risk.

LOVE: Why is this interesting? As you may know, AI depends on data. Engineers generally "train" algorithms by feeding them data, and the algorithms then "learn" by picking out patterns in that data. So the training data that's been used in the past is what's going to determine whether or not Olay's app actually tells you something meaningful about your skin.

So where does a company get its training data? Looking around, you might think that this shouldn't be a problem—that we're living in an age of nearly infinite information. If Olay needs pictures of faces, just think of all the selfies that they could mine from Instagram and Facebook! But, Pah explains, just because the data's out there, doesn't mean that everyone can access it. In fact, this is the first surprising thing that business leaders should know.

PAH: The data's become far more siloed as we've started thinking more about privacy.

LOVE: In the wake of high-profile data scandals and breaches, data sharing practices have really changed. In some places, like the European Union, government regulators have cracked down on data sharing, Pah says. And even where laws haven't changed, many companies, including Facebook, Instagram, Google, and LinkedIn, have gotten much more cautious about who can access their users' data.

PAH: And so, in general, things have pulled back. Facebook still knows everything about you. Google still knows everything about you. But in general, data sharing across platforms is harder and less so.

LOVE: Meaning, organizations may need to think creatively about where to get the data they need to train AI. Some companies are turning to their customers—like Olay. That's what that line in the terms and conditions was about. When a user submits a photo, Olay can actually use that photo to keep training its algorithm.

Other companies are dealing with this data crunch in other creative ways. Sears, for instance, has been building what's called "synthetic data." That's where you use a small amount of information about real customers to build tons of imaginary customers, which can give your AI more predictive power. You can then use this synthetic data to predict how all those different customers will respond to an incentive, for example.

PAH: It's basically creating unbelievably, realistic, synthetic us's—hundreds of millions of them, customers upon which to test on. This is where we're moving, into this 'data light'

world where you have a little bit of data, but now it's lots of, essentially, machines teaching other machines.

LOVE: And on the flip side, the data crunch has made some companies realize that they're sitting on a valuable commodity. There's a growing market for what Pah calls "data exhaust," data that companies just happen to collect or create in the course of doing business. For some companies, this data exhaust has actually become a central product. Take the company Netradyne, for example. They produce a fancy camera that you put on a vehicle's dashboard.

PAH: ...That can look out, it can look to the sides, and it looks back at the driver. It can see if you're driving through a stop sign. That's what their product is. The ancillary product of that is, they're constantly mapping the roads. And so they have road map data that's this data exhaust that creates another product.

LOVE: A product that's turned out to be very valuable in its own right. Netradyne recently partnered with Hyundai, which plans to use these millions of miles of mapping data to train its self-driving vehicles. The point is, in an age of data constraints, managers should think broadly about where to get data, and what their data could be worth to someone else.

[musical interlude]

LOVE: Pah's second surprising lesson is about how companies use the data they have. You might think, since AI relies on data, that having more of it will lead to better AI. But you'd be wrong. To see why, let's look at the results that Pah got from the Olay app. First it recommended a product.

PAH: The ultimate eye cream for wrinkles, puffy eyes and dark circles.

LOVE: And then it used the photo to guess what it calls his "skin age." Which is where things got a little wonky.

PAH: It's guessing that my skin age is 42. That's a little bit of an overshoot. Not gonna say how old I am, but it is in the wrong decade. [*laughs*]

LOVE: Now, to be fair, your skin age doesn't necessarily match your actual age. And, again, this product is technically meant for women. But still, Pah doesn't think that he has the skin of a 42-year-old. To understand what might have gone wrong here, it helps to consider his second surprising fact about AI: not all data is equally useful. In fact, training your algorithm using just whatever data you have on hand can come back to bite you. That's because this data typically doesn't represent the entire universe of your possible customers.

PAH: Very rarely did anyone ever say, "I'm interested in creating recommendations. I'm going to go and carefully collect all this data about every different type of person that's ever bought something, and see what they bought and why." That's not how things work. People essentially have the data stores that they have, and they say, "Hey, why don't we use it to make an improvement?"

LOVE: So why is that a problem? Because when you throw all of your data into an algorithm, you actually only end up learning what your most frequent customers want. And while it's nice to know what your most frequent customer wants, this can lead to bad guesses about what everyone else wants. To see why, Pah draws an analogy to trying to teach a person something basic.

PAH: I'm trying to teach you to say whether a picture is of a cat or a dog. And I just show you a picture of a cat and a picture of a cat, and a picture of a cat, and a picture of a cat, and a picture of a cat. And only one out of a thousand times do I ever show you a picture of a dog. The simplest thing that you could learn is just always say "cat." Like, don't bother trying to figure anything out, just say "cat." Because you're actually going to be right 99.99 percent of the time, just saying "cat." And in fact, if you actually try to learn something about when to say "dog," you're actually going to probably start making mistakes and saying "dog" when it's a cat. And you do worse than if you just did kind of the bad thing of saying "cat" every single time.

LOVE: So if most of the people in your dataset are similar—if they tend to like the same products, or are of the same race or gender—then your algorithm won't know what to do with people who aren't like that. This is likely what happened when Pah used the Olay app.

PAH: I imagine that I am one of the few men that have used this so far. So it's applying what a woman's face looks like to judge my face, which is not great when you consider that half my face is hair. Learning about what one group wants effectively ends up harming how good the recommendations are for all the other groups that don't have the same types of interests or attributes as the most frequent customer.

LOVE: So how can you avoid this bias in your AI? For one, instead of just dumping in all the data you have, think about collecting data more scientifically. There may also be resources that can help you. If you're creating face recognition algorithms, for instance, IBM recently made available a huge database of thousands of faces, featuring equal numbers of different ages, genders, and races. It's just the kind of dataset required to train facial recognition software to account for all kinds of people.

[musical interlude]

LOVE: Pah knows that a lot of people have a kind of wonderstruck attitude towards artificial intelligence.

PAH: They're like, "It's new, it's really cool! Everyone is doing amazing things with it! We've got to AI tool this, too!" But if you make something no one's asking for, you haven't really done much of anything at all.

LOVE: So his final surprising thing: for all the talk about AI, few companies have really figured out how to use it to serve their customers. So, if your organization is getting started with AI in the product space, Pah has a few tips. For starters: acknowledge the limits of the technology. Maybe you have a vision for an amazing algorithm that will help a customer troubleshoot based on a photo they snap. But as the Olay app shows, the technology may

not be entirely there yet. You see, in addition to that selfie he took, Pah also had to manually answer a lot of questions.

PAH: “How many skin care products do I use in the morning?” None. Now it’s asking about my skin type.

LOVE: That’s because an algorithm can only learn so much from a photo you snap on your phone. It can’t tell what products you’re already using, for example. And a blurry picture or bad lighting could make it hard to tell your skin type.

PAH: It would be hard to distinguish, if there's a sheen on your forehead that's like an ice rink, if it's from bad lighting or if it's from being extremely oily.

LOVE: So before you start building anything, consider whether AI is really going to meet your customer’s needs. You might talk to your developers, or see what others in the space are doing. Maybe your idea doesn’t call for a fancy algorithm—or, if it does, you may want to wait a couple years for the technology to catch up.

But there’s also something that Pah thinks the Olay app is doing really well: actually helping customers figure out which products will work best for them. Pah think that more companies could follow this example, using AI to address a real customer pain point in a new, clever way.

PAH: Actually looking at me and saying “this would work for you.” Not based on what you've bought. Because let's be real, I have no idea if the past products I bought are good. No idea! Right? I'm like, “What is the cheapest face wash here?” So, yeah, I mean, my past purchases are actually awful if you're trying to help me actually take care of my skin. So, being able to make something that helps me learn a little bit, is actually really exciting.

[musical interlude]

LOVE: This episode of *The Insightful Leader* was produced by Kevin Bailey, Jessica Love, Fred Schmalz, Jake Smith, Michael Spikes, and Emily Stone. It was written by Jake Smith. Special thanks to Adam Pah.

As a reminder, you can find us on iTunes, Google Play, or our website. If you like this show, please leave us a review or rating. That helps new listeners find us.

And of course, we’ve only scratched the surface on AI. So if you want to go a little deeper, just head to <http://kell.gg/book>, where you can pick up a new ebook that we just put together on AI and business. It features Pah and a dozen other experts offering their tips and perspectives on where the technology is heading, what leaders need to know, and how you can stay ahead of the curve. And it’s totally free! Download your copy at <http://kell.gg/book>.

We’ll be back in a couple weeks with another episode of *The Insightful Leader*.