

MANAGEMENT

Can Artificial Intelligence Help Answer HR's Toughest Questions?

Aug 30, 2019

📍 North America 🎧 Human Resources, Podcasts, Research



With companies continuing to shrink or outsource their human resources departments, it is tempting to augment that traditional business function with artificial intelligence. Data science holds so much promise for other fields that it makes sense for algorithms to replace imperfect human decision-making for hiring, firing, scheduling and promoting. But new research from Wharton professors Peter Cappelli and Prasanna “Sonny” Tambe flashes a cautionary yellow light on using AI in human resources. In their paper, “Artificial Intelligence in Human Resources Management: Challenges and a Path Forward,” the professors show how limited data, the complexity of HR tasks, fairness and accountability pose problems for digital HR. The study, which was co-authored by Valery Yakubovich, professor at ESSEC Business School and senior fellow at the Wharton Center for Human Resources, also looks at how to remedy those problems. Cappelli and Tambe spoke about their research with Knowledge@Wharton. (Listen to the podcast at the top of this page.)

An edited transcript of the conversation follows.

Knowledge@Wharton: You make the point that while AI is invading many different industries and sectors, there are some special concerns when it comes to using AI in human resources. Can you talk about what some of those challenges are?

Sonny Tambe: When you talk to HR practitioners who see their colleagues in finance and marketing using these technologies with so much success, part of the question they ask is, why does it seem so hard for us? I think part of the point we wanted to make is that there are systemic and structural differences for HR that do make it harder. For example, when you are building an AI-based system, you need to know what the right answer is, what a good employee looks like, and defining that in HR is already a difficult thing to do.

Peter Cappelli: The language of data science is a language of optimization. The idea is that if we figure out what the goal is and what is associated with that goal, then we can apply this to decisions going forward. The language of human resources, which is coded in law, has an awful lot to do with fairness questions, and fairness and optimization don't often go together very nicely. You could get algorithms that will tell you here are the people to hire, and the people to hire are just like the ones who have done well before. These algorithms are backward looking; they are based on data from the past. The people who have done well in the past are white males. You probably would not want to pursue that algorithm as a hiring strategy, even though the machine learning algorithm would say that's the thing to do.

“The language of human resources, which is coded in law, has an awful lot to do with fairness questions, and fairness and optimization don't often go together very nicely.”

–Peter Cappelli

Knowledge@Wharton: How widespread is using AI in human resources right now? Is it unusual to have a data scientist working in HR? Are HR departments contracting with vendors that have these types of tools?

Cappelli: We have a regular group of companies now that come together to talk about this, and Sonny created the acronym for this: CODHR or Conference On Digital Human Resources. It is talking to these folks who are typically all data scientists in human resources, so they're all there. They're struggling mightily just to get data together to analyze, and this says a lot about

the reality of trying to work with human resource data. Here is a typical problem: We have data on employee performance, and it's filed in this dataset over here. We have data on hiring and the attributes of applicants in this dataset over here. But, by the way, as soon as somebody is hired, we tend to throw that data out. The biggest problem they've got is, can we get these datasets together to talk to each other? That is not so unusual, except we also run into stories about the person running this silo who doesn't want to share the data with this group over here. It's a data management exercise, but it's also a political exercise internally.

Tambe: One of the really interesting things about the two questions you asked — how common is the use of AI, and how often do you see data scientists — is that there is a gap in HR between those two things. There is a lot of hiring, and a lot of people are bringing data science into the HR function. But some of the issues, like the ones Peter raised dealing with data and so on, have also a paralytic effect. Now that we have these data scientists, what can we do with them? You can have data scientists, but the ability to translate that to actual AI has become a struggle in HR.

“You can have data scientists, but the ability to translate that to actual AI has become a struggle in HR.”

–Sonny Tambe

Knowledge@Wharton: How does using algorithms differ from traditional HR practice?

Cappelli: Let's look at hiring, for example. The way hiring was done a generation ago, when they actually took it much more seriously than they do now, is that we would come in as a candidate and you would give us an IQ test, a personality test, structured interviews, and maybe work samples and all of that stuff. We would have five different criteria, and you would get five different scores. The hiring manager would look at the five different scores. Maybe they have discretion over how to weight them, but they would say, “I am going to go with Rachel on this one and here is why. She scored well on three of the five.”

What happens now with these algorithms is they might take those five elements and anything else they can find, too. They look at your resume, everything on it. They look at your background, everything on it. They look at things that you might think are irrelevant. One of the great examples was commuting distance. And they would throw them into a model, which they built to relate as closely as possible to the performance of employees on the job in the past. As a result, you get one score. The difference is that when you get the machine learning-based

algorithm, you don't get five scores on five different items, you get one measure. And unpacking that is extremely difficult to do. What's driving that one measure? Well, that's hard to say.

In some ways, it's much simpler and much more powerful than what we were doing before. It might well be a much better predictor. But it's also complicated because, if you think about these fairness issues, somebody asks you, "What's driving that measure?" You say, "Well I don't know. There are these 10 things."

Tambe: Another point of tension is what practitioners call explainability. You have all of this data going into a prediction. With the old system, maybe it's easier to say, "These are the reasons I arrived at this conclusion." With the sort of basket type of approach, it's harder to back out and say, "This is how we got there." You just have a prediction, and it's hard to figure out how you got there sometimes.

Knowledge@Wharton: Might that also create some legal challenges where applicants could file claims of unfair hiring practices?

Cappelli: Yeah, especially if gender is one of those factors that is in the mix. Even if you take gender out, there are attributes that are sometimes correlated with gender, like the courses you took in college and things like that. Can we really be sure that those are not the factors driving the score in some big way? It is hard to know.

Knowledge@Wharton: You studied this in part using a workshop that brought together HR practitioners and researchers. Can you talk a little bit about that and some of the findings?

Cappelli: We've done a couple of these, and the first one was we brought our data science colleagues in from the engineering school. People who are data scientists as opposed to statisticians who work with human resources. We brought those folks in, and then we brought in people who have data science jobs inside companies. I think there was enormous sympathy for the problems that the data scientists have trying to do this kind of work. I think also there are some things that the data scientists know that have not worked their way down into the practice, so there's a gap there. That was interesting.

Tambe: I think because this is an emerging frontier and a science, you can learn a lot from just getting people in the same room who are working on the same problem. We talk about data being the beginning of the pipeline for AI, data being the new oil and so on. The notion that just getting the data together, as Peter was mentioning earlier, can be just so difficult for a variety of

reasons. Part of it has to do with old databases, legacy systems, GDPR type regulations. But it's such a high hurdle for a data science team to have to cross that it can really be a big, big constraint.

Cappelli: As an example of that, one of our colleagues, Susan Davidson in data science, works on the problem of small datasets missing data. One of the big problems that people have in organizations is that data is often pretty messy. There are often pieces of it not there, or they skipped over something. The data scientists have thought about that a lot and have figured out how to work with missing data, very small datasets in particular, and how to do sophisticated data science stuff with data that you would think was too small to do it. That would be pretty useful to these folks who actually have to do it, but it hasn't quite made it over yet.

Knowledge@Wharton: In addition to complexity, you identified in the paper four key challenges facing the adoption of AI in HR. Could you go through those?

“One of the big problems that people have in organizations is that data is often pretty messy.”

–Peter Cappelli

Tambe: I think we talked about the definition of a good employee. As Peter said earlier, what you are trying to optimize on? That in itself is a difficult thing for organizations to define, especially when so much of work is team based. That was the first one.

Then, we talk about data size. When we think about most AI applications, it's pretty short work to generate millions of observations. If you are thinking about advertising, for instance, you have click stream data — lots of people viewing ads and clicking. Finance, you've got trades going on all day. But for HR data, you've got maybe one data point per employee if you're talking about churn or compensation. So, limitations to data are another obstacle.

Fairness was another challenge. You have these algorithms and sometimes they make predictions in ways that disadvantage certain groups. Sometimes it is hard to figure out that the algorithm even did that because of the explainability problems.

Cappelli: The fourth one was employee reaction. Here's an example of employee reaction, and it comes to this explainability thing. Right now, supervisors have a fair amount of power in the workplace. They might control your schedules, for example. If you move strongly in this direction of algorithms, the algorithms start to make those decisions instead. Let's say we figured out in our work group who has to work on Saturdays. The algorithm [creates a rotation],

and I get two Saturdays in a row. Who do I complain to? Well, I can't complain to my supervisor because she didn't do it, right? Would she say, "I don't know. Here's the name of the software programmer in Silicon Valley who came up with the algorithm." I can't complain to my supervisor, nor is my supervisor able to do anything about it. A supervisor can't say, "I understand this was not fair. But we'll take care of you next week."

One of the issues we might think about is the extent to which you move in this direction really centralizes authority and disempowers supervisors, which makes their work much more difficult. We are increasingly recognizing that the connection between direct reports and supervisors is the heart of the organization. It's the heart of how employees feel about their employer. You are taking power away from them and making that intersection weaker. Have we thought through what that is going to mean? I think the answer is no.

Knowledge@Wharton: What are the key takeaways for companies about getting past those challenges or confronting them at all stages of the HR cycle?

Cappelli: Be careful, I would say is the main one. Try to think through what are the consequences of this when you're done. Let's say we're going to have an algorithm to predict who gets which assignments. How do we do it? What is going to change if we do this? What do you imagine the complaints are going to be? How are we going to deal with those complaints?

I would say another piece of advice is be really careful about vendors because it's hard to know what they are doing. Where do these algorithms come from? They're never built on your own data, so the fact that somebody has a whiz-bang story about what this algorithm will do for you doesn't mean you know it's true. Before you invest in any of these algorithms from a vendor, if they are not building it for you, if it's come from someplace else, you want to test it with your own data to see if it works for you. Among other things, you have no legal protection against adverse impact claims unless you can show that this algorithm predicts in your own workplace. I would say those are my two big ones.

Knowledge@Wharton: You mentioned that one of the biggest issues for many companies is that the data gets trapped in silos, even within larger HR departments. Is there any advice for getting the data out of the silos so everyone can use it?

Tambe: One of the most interesting discussions that we heard about was the difficulties that the companies were facing when it came to just federating different datasets to make even simple predictions. I think part of the challenge here is that when you are talking about data science, the outcome is inherently uncertain. It's exploratory. It's experimental. Meanwhile, the costs

are big because you've got data living in different parts of the organization and there are always difficulties in accessing datasets. The feedback from practitioners was that it is useful even to just establish a common understanding that there is something substantially different here.

**“When you are talking about data science, the outcome is inherently uncertain. It’s exploratory. It’s experimental. Meanwhile, the costs are big
....”**

–Sonny Tambe

Peter makes a point that we have a common understanding that, from a marketing perspective, companies are trying to make money. If they have a better machine to do it, they will use it and that makes sense. We as customers are okay with that. Socially, we don't have the same understanding about HR. There are a lot of issues around what best practices should be, and these issues cut across functions or organizations. These are organization-wide decisions to some degree; they have to sort of float above the HR data science team.

Cappelli: Also, I think that to manage expectations better is probably useful. As Sonny was saying, these are big bets to do even the simplest thing because the biggest problem is getting the data together. Spend your money on data-based management before you start hiring fancy data scientists, right? Because unless the data is together, it's going to be worthless to have all of that power. Somebody above at a pretty high level has got to decide the answers to these questions really are important, so you do have to share the data.

We had heard from a couple of companies where they were making vendor decisions in a council-like way. The performance management people could not buy a software system for performance management where the data could not be shared with the hiring people. That is trying to figure out what the meta goals are here that lie above each of these individual decisions and making sure that we can, at the very beginning, conduct these bigger exercises. Somebody at the top should be saying, “They are really worth doing, and let's prioritize them.”

I think there is often a view that you just turn these data scientists loose and they are going to find all kinds of cool stuff. That is unlikely. First of all, just to be able to do anything is quite difficult. And the questions in human resources are ones that people have studied for a very, very long time. The idea that you are going to find some simple a-ha breakthrough that no one has thought about before, which maybe in marketing you might find because nobody has looked at a lot of these relationships before, is much less likely. On the other hand, current practice in a

lot of human resources is so lousy, to be honest, that the opportunities for making some progress are pretty good. Even if you're not going to generate the silver bullet solution, being better than what we're doing right now might not be all that hard to do, and that's a reasonable thing to shoot for.

Knowledge@Wharton: How can companies educate employees and bring them into the decision-making so this is not like a slap in the face?

Cappelli: This is the explainability thing that Sonny was talking about — that is, can you explain to some new hire why they didn't get the promotion because the algorithm score said 86 for them and 92 for somebody else. What does that mean? How do we unpack this algorithm, and how was it built in the first place? We need to know the answers to those questions before we talk to the employees about it.

Part of the problem with explainability is the answer you might have to give them is not one that the employees particularly like. You fit with the work needs, so you just happen to have to work three Saturdays in a row. The employee says, "That's not fair," and the algorithm says, "What's fair?"

Knowledge@Wharton: Does that speak to having some sort of system in place for the human touch to override the algorithm?

Tambe: That's a really interesting area of AI design right now which is thinking about how you separate the algorithmic piece of this from the high touch piece of this. I know in medicine they are thinking about this as well. They are thinking about the notion that you can make a decision algorithmically, but you probably don't want the machine to deliver that decision to a patient or a patient's family. When we think about the design of systems in HR, there might be some very similar design issues that start to arise.

Knowledge@Wharton: In the paper, you say companies shouldn't start with the big questions but maybe start with smaller bites of things that are a little more easily graspable. That seems counterintuitive.

Cappelli: Yes, I think some of that is because of explainability. We are trying to decide, let's say, what to encourage you to take in terms of training next, and we are doing this based on what people like you in the past have done next, so this kind of training makes sense for you. Something like that about advice, which doesn't have quite as big a bang as "you get this promotion, you don't," might be the way to start because it's an easier thing to swallow.

Starting out with simpler questions, partly because many organizations are not so sophisticated at those simple questions anymore, is probably a reasonable thing to do. With the really complicated questions, it's harder to get an answer that is going to be useful and even harder to get an answer that you can explain.

Tambe: Another point in favor of starting with something tractable is that so much of this data science process is learning and exploration. You learn about your data, you learn about your capabilities, you learn about your employees, so it's helpful to at least understand the question you are asking before embarking into less structured unknown questions. You get up to speed on what you can do, what your data can tell you, what your people can tell you.

Cappelli: I think it's useful for everybody to understand how such use of data actually works. As much as I would like to have people think that we're like wizards with this stuff, that we go into the computer, we take this data and then magic comes out, it's an awful lot of just plugging away: There's a problem with this data over here, there's some missing observations here. It takes a very long time just to get the data into a format where you are pretty sure what it is. Then you look at things, you don't find anything, and you look again in a different way. There's a lot of trial and error, and there's a lot of experimentation going on in terms of coming up with something that is credible. It takes a long time to do it. I think there is a view that this is all kind almost like magic. You get this data together, you hire a data scientist, they go behind the curtain, and poof, they come back in 10 minutes with the answers.

“As much as I would like to have people think that we're like wizards with this stuff, that we go into the computer, we take this data and then magic comes out, it's an awful lot of just plugging away.”

–Peter Cappelli

Knowledge@Wharton: What's next for this research?

Tambe: That is a good question. I think one of the things it underscores is that there are so many big questions — the way algorithms in HR map to legal frameworks, the way they map to social frameworks. A lot of these challenges that are emerging just for HR are uniquely big and uniquely difficult. They probably point to some pathways for a focused research agenda required to really get to those answers. A lot of this is new territory, and a lot of it would benefit from more exploration.

Cappelli: In the short term, I think what we are hoping to do is just see what companies are doing. Are there new algorithms that are being done in a different way? Are there new techniques being tried out? Just getting a sense of what is sensible. The good thing about the vendor world, for our purposes and some employers, too, is they are always throwing up new and different solutions. Many of them don't make sense, but occasionally some do. We're in the happy position of just being able to watch them and not have to live with them. I think just seeing what happens next is probably what's next for us.

All materials copyright of the Wharton School (<http://www.wharton.upenn.edu/>) of the University of Pennsylvania (<http://www.upenn.edu/>).

Report accessibility issues and get help (<https://accessibility.web-resources.upenn.edu/get-help>).