

MIT CSAIL Alliances | George Westerman Project 7(1)

Welcome to MIT's Computer Science and Artificial Intelligence Labs Alliances podcast. I'm Kara Miller.

[MUSIC PLAYING]

On today's show, how do you know when to use AI and when not to?

If you're going to diagnose a patient, you've got to be right. If you're going to send one more marketing email, it doesn't really matter.

Plus, should a new set of tools change how managers manage? MIT Research Scientist George Westerman says absolutely. And he argues that change is coming, but not fast enough.

I believe there are two false narratives that people are telling themselves on a regular basis. And nearly everybody I know, when I talk to HR people, nearly 100% of people will give one of these false narratives.

But there are companies that are getting this transition right.

And at this one company, it's a healthcare company, not only can they tell you, out of the 1,000 people that were in your job before, here's the percentage that went to different jobs, but also, if you push a button, we'll send an email to one of them and they'll call you up and tell you what that job is like.

That's all coming right up.

[MUSIC PLAYING]

George Westerman advises corporate and government leaders around the world on AI. He knows there's huge interest and he's pretty excited about the transformation that's underway. But when you get talking to him, caution starts to creep into the conversation.

So I'm seeing a lot of people applying AI where they don't need to apply AI.

Westerman is a Senior Lecturer at the MIT Sloan School of Management and the author of multiple books, including *Leading Digital-- Turning Technology into Business Transformation*. And he advises leaders to essentially look before you leap, as in really look at those tools in the toolbox before you choose AI.

So things that they could do with basic statistics or basic rule-based systems, which are cheaper to run, you actually can explain them better, but because AI is the thing, we're going to apply AI to it.

What happens when you employ a fancy solution to solve a simple problem? Well, says Westerman, a few things.

Algorithms that require much more expensive talent to create them, often much more data to do it and not explainable models. So one of my colleagues, Sam Ramsbotham, when he interviewed-- I believe it was one of the head data scientists from Orangetheory, he said, what's the best-- your favorite example of when you did AI?

And he said my favorite example is when I solved the problem with a regular regression. So that's one of the things I'm seeing. I'm seeing a lot of people with hammers running around for nails.

And I mean, when you talk about these more traditional data analytics, I think the issue of cheapness is not to be thrown away too quickly because when you-- I mean, just when you look at the energy usage of a ChatGPT query versus a traditional search, it's quite a lot more. I think it might be on the order of 10 or 12x. So you multiply that across doing it all the time and you're really looking at cost differences. So it makes sense if the problem needs AI, but maybe not if it doesn't.

Yeah, it's been interesting. I'm seeing this with my students also, that we have such powerful tools out there, that we use the powerful tools without thinking whether we can do it with something less. Just for example, one of my students, he's doing some really interesting work with hundreds of millions of data points. And he's got so much data that it's breaking his PC.

Now it turns out when we sat down, he's doing hundreds of millions of data points with some really powerful tools that are good for using small data sets, but it just uses too much memory if you've got a giant data set. And so we've had to go to first principles and say, if you were to code this without all those help, how would you do it?

And we're using something like 100 times less memory because we're not using those sophisticated tools. But he's not trained to do that. He's trained to use the most sophisticated stuff available. And I think we do that in many cases now.

We got the tools. We'll use them. And like you said, we don't really think about what it's costing us in terms of energy or even in terms of money because it's out there in the cloud. That's somebody else's problem to manage that problem.

So if you're an executive and you're thinking about, "OK, do I use AI for this or do I not? Do I use something more traditional and maybe cheaper and maybe simpler?" what's the question you ask yourself or the litmus test for that?

So I think the first thing to do is if you're starting to talk about "Will this company use AI or not use AI?" that's the wrong conversation. The question is going to be "Where will we use AI? Where is it appropriate?" And so start with the problem you're trying to solve.

So, for example, at Home Depot, they talk about a delightful customer experience. And whether that's got natural language involved in it, whether it's got image recognition involved, whether it's got optimization algorithms involved, that doesn't matter.

It's "What is this doing to create a delightful customer experience?" that's the problem. And we'll apply the right tool for it when we get there. So number one, don't talk about AI or not AI. That's the wrong question. Talk about where AI, where not AI.

Number two, figure out-- when you're talking about where not AI-- what's the problem you're trying to solve, and then we'll get at that. And then three, then there's some questions you can go through for each problem to figure out whether it makes sense.

So what are the questions that you'd ask yourself to figure out what AI to use?

So these different-- first of all, we want to think about-- there is the really-- there's a generative AI that's creating new things and it can also be used to sense patterns. There's the deep learning that looks at optimization and looks at pattern matching.

Then there's the statistical work, just finding trend lines and other statistical econometric work. And then there are rule-based systems. And each of these, some are more explainable than others. Some are more accurate than others. Some are more repeatable than others.

And so you want to think about those kinds of issues when you're thinking about the different things. So it kind of breaks down to a few questions you can ask. First of all, what is the cost of being wrong in this situation? If you're going to diagnose a patient, you've got to be right. If you're going to send one more marketing email, it doesn't really matter.

Number two is does it have to be explainable? If you're in medicine, if you're in finance, you need to be able to tell the regulators why you did that. Other cases, it may not be the same. Repeatability is another one. Do you need to get the same answer every time?

If you need to get the same answer every time, don't do generative AI because it, by its very nature, will not-- it's not supposed to. And then this one's probably going to be less of a problem over time, but confidentiality. So to what extent do you need to be careful that if the information you have gets out, it will get into trouble?

We all have heard what happened with the generative AI, where you want to use it to talk about that merger you're about to do it, and now everybody knows you're about to do a merger. That's getting better, but maybe not there. So I think if you ask about those questions, that can at least help you narrow down not only what kind of AI to use, but also a lot smarter about the problem you're trying to solve anyway.

Now I'd add to that. Then, of course, you want to get back to the data. Do we have the data? How well do we know true or good from bad on that data? Good loans, bad loans, that kind of thing. And then also, this is what we were talking about earlier, about just bias and generalizability.

So if I tested this medicine on a bunch of adult men, will it work for women? Will it work for children? How generalizable is the data you've got?

Well, this is a little bit of the garbage in, garbage out thing. If you don't know which loans turned out to be good and which turned out to be bad, how can you feed decent data into the system? You have to understand.

Exactly. And that's why--

[INTERPOSING VOICES]

And that's why every time we go on a website, it says, "Tell us where the stop signs are or tell us where there's a bicycle," because we're creating that data for other algorithms to use. And the one that gets-- the ones that drive me crazy because my relatives do it all the time is show us the young you and the old you. It's the Facebook challenge.

Well, we're training Facebook now on how to age algorithms. Show us the sibling challenge. And you show the family pictures, now we've given all that data over. And if you care about privacy, now the algorithms have all of that.

Let's talk a little bit about jobs. You alluded to this before. One of the striking things is this Pew poll asking Americans, "Have you used ChatGPT?" In March 2024, not that long ago, they updated the poll. They've been asking this question.

Just over a quarter of people, between 30 and 49, so prime working years, have used ChatGPT. Just over a quarter. That's not that many. I wonder what you make of the kind of disconnect between, I'm guessing, what you hear, the boardroom, and the public because it seems like there is one.

Well, I think it's just another example of the kinds of polarization we've got going on in our country. We got people talking to people like them and assuming the rest of the country is that way, through no fault of our own.

So I talk to professors, I talk to business executives, I talk to highly educated people who do technology all the time. Everybody I know has done ChatGPT. But I forget about my mother-in-law who had trouble learning how to use her cell phone a long time ago. She got really good at it, but it took a while. It wasn't her first instinct.

And so I think that's the challenge we've got, is that it's another example of those filter bubbles, where we're making judgments based on our own experience when we really need to broaden that out. But the other thing to think about in this situation is that although only 25% of people have used ChatGPT, closer to 100% are worried about ChatGPT.

And so there are two problems there. One is that we've got people worrying about something without having used it. And number two, we need to, as a population, as a public, as a media, help people understand this better.

If we take this to the employee side, people are nervous. What will this do to my jobs? One of the big things a company can do is they can help people understand how this might help, how this may hurt, and most likely, how it's going to be some combination, where some of your job will change, where others' job will actually get improved.

You have argued that AI is going to transform knowledge work, which I think is really-- I mean, it's a really profound statement because more and more of us do knowledge work. Certainly, not everybody, but a lot of people do. How do you see it transforming knowledge work?

So certainly, we know that from past work on automation, past work on AI, if you're doing routine work, it's highly likely that routine work will be automated over time. So if you do routine work, it's time to think about what you could do that's less routine. What's the extra value you add beyond the routine because that may be going away soon?

So let's assume you're doing knowledge work that has a combination. We're seeing, actually, a lot of augmentation opportunity there. It doesn't have to replace you. And first of all, it can do a lot, as we've all seen, with just taking the routine work out.

Here's a 300-page PDF. What did they really say? Or, hey, I didn't get to go to that meeting. Can you give me a paragraph on what happened in that meeting? It's easy to do now and it just saves you all that cognitive effort.

Well, and it may also make you smarter because maybe you end up reading ten-- not fully reading, but kind of reading the summary of 10 research reports on stocks. Let's assume you're in finance or something. You never would have really gone through all those pages, but now you actually have a sense of what was said in those research reports.

And what's awesome about that is all the effort you would have spent and all the brain cells you would have tired out reading those things, you don't have to now. So you've got more time to put that together and make it make sense out of it in your brain, just like you said.

And that leads to the next thing, is it can help you enhance your cognitive abilities. And certainly, what you just mentioned is one of them. The other is just for creativity. So, many times, I will put something in and say, "Can you give me a creative title for this?"

But imagine, hey, let's try to figure out-- "Here's three ideas. How might this fit together in an image? Give me 10 versions until I find one I like," or even better, "OK, I'm trying to make sense of this idea and this idea and this idea. What do you see as some commonalities and differences there?" It can do that now. So it can help me be better, not only removing the routine work, but actually better cognitively there.

And then there's a third one. And the third one is just what we're seeing in computer as tutor, computer as coach. And we're seeing that everywhere from call center, pretty routine call center work, all the way up to like at strategic consultancy, these computers tend to be helping people do a better job.

They're coaching people in ways that you'd hope that people could coach, but you just don't have that many people to do it. So the tools can help you become better at what you're doing, also, help you learn better.

So there's this gap between how many people have actually used AI and how many people are fearful of AI. Does that mean that executives are having to deal with this kind of existential fear from people? And I don't know if that actually puts up a little bit of a roadblock to implementing more technology in the workplace.

People don't like change. Most workers do not love change. And so there's this constant process of not only changing the structure, but helping people realize that there's something in it for them for making this change happen, or else they're going to fight you either passively or actively.

And like you said, AI is another one for managers to think about. There are some companies going out and giving basic AI training to everybody in their company. And it's not teaching them how to code, but it's just "Here's what it can do. Here's what it can't do. Here's how it can help you with exactly that goal."

That people are nervous. They may or may not consciously realize it. They may or may not be actively voicing it. But let's just help people understand that we're not trying to get rid of you. This should be a tool to help you. And let's help you understand more than just me telling you it's going to be good for you. Help you understand a little bit more on why.

You've thought a lot about skills for the jobs of the future and how you build those skills. And you've argued that we don't work nearly enough on developing people's careers. And I think this is tied in with this question of AI because when you think about what's next in technology and how do people adapt and how do people feel actually more calm about what's ahead because they're ready for it, it seems like there's tension there.

Yeah, I think companies think about what they're doing with their high-potential employees and somehow think that that's working for everybody. I believe there are two false narratives that people are telling themselves on a regular basis. And nearly everybody I know-- when I talk to HR people, nearly 100% of people will give one of these false narratives.

And as number one, we make it every boss's responsibility to develop their people. It's a false narrative because even if the boss wants to develop you, which is not a given, they may not know how. Especially if you're looking to move horizontally or diagonally, then you want.

So that's a false narrative. Just because you say in your annual review, "Talk about it," doesn't mean it's really being talked about in a serious way. The other false narrative is "We give you a lot of tools and a lot of training courses. you can assemble them together and chart your own career progression.

And that's a false narrative because although it sounds believable, it misses the whole point that most people need more than just skills. They need a person to help them. They need a person to take a chance on them. All of us in our careers have had that lucky break with somebody.

And the computer is not going to give you that lucky break. And so what we need to get is beyond that false narrative. If the high potentials are 1% of your population, what about the other 90% that you might want to develop below there? We need to get much more intentional about how to do that. Stop telling ourselves these easy-to-believe stories and start to get better at that.

And how do you do that?

So number one is help people understand what's possible for them. The routes upward are limited and not everybody may want them. I know, for example, a postal carrier who got promoted to manage the post office.

And he said it was the worst decision that ever got made because he said, George, I had to manage people like me and I forgot how bad that was. Not everybody wants to move up. Not a lot of teachers want to become principals.

I was just thinking of somebody I was talking to a few months ago who said-- she actually is in administration, but she said "The last thing-- originally, the last thing I wanted to do." Because being a classroom teacher and being a principal, they're not the same job. I mean, they're in the same sector, but that's about it.

Yeah, so not everybody wants to move up. Where else might I go? One company I know has this really-- and actually, multiple companies now. I learned about them a few years ago, but now many people are getting there. I am in this job. Where have people like me gone before that?

And at this one company, it's a healthcare company, not only can they tell you, out of the 1,000 people that were in your job before, here's the percentage that went to different jobs, but also, if you push a button, we'll send an email to one of them and they'll call you up and tell you what that job is like.

Isn't that amazing? Wouldn't that be great? That's what you hope your boss would do for you but they can't. That career navigation opportunity is huge. Next up is we just need to make the learning just faster and more. And we have a whole work to transform our chief learning officer and how to make this stuff work better.

But then last, I think we need a better way to help give people rapid feedback and good mentoring that we just don't do. If we're limited to an annual meeting with an HR person, an annual meeting with our boss if the boss is willing, that's not enough to really mentor people in where they're going.

So first of all, people need to know what's possible. Second, they need faster ways to get the skills and practice those skills. And third, I think we need more ongoing coaching and mentoring. What's good about that is AI can help with all those things.

AI can help you understand what are some good paths. They can help you figure out what learning might help and may even give you an opportunity to practice those skills in a scenario-based or a game-based world. And they can also give you more rapid feedback because, especially, generative is very good at coaching as we move through.

The only thing I'll just play the skeptic here is that even if you use AI to facilitate these different paths-- you can take in your career, at least exploring those paths-- I mean, it still takes some investment of time and money from the company.

And if your manager or HR or whoever, they're very busy doing their own job and having their own meetings and getting done what they need to get done for their manager and their list of to-do items. So what's the incentive for a company to help you to learn things, to connect to what people in this job before have done. Who cares?

So in 2022, a really well-done study found that for people changing jobs to another employer, 2/3 of them said the reason was they didn't have adequate career development opportunity. So just from a pure retention, it's hard enough to hire people right now. You don't want to lose so many of them just because you're not giving them opportunities to develop.

And also, you could argue-- I know Zeynep Ton at MIT has argued this, that it's very costly when somebody who has a lot of knowledge about your company walks out the door. And we don't really know how to put a number sign on that, like how many dollars is that walking out the door. But when somebody new comes in and you have to teach them from scratch, that's costly.

And we do have numbers on it. By the way, Zeynep's research is fantastic. If you have not seen Zeynep's research, you should look at it. She does a great work. Yeah, we do have numbers we hear all the time, that it's two or three or four months salary to bring a new person on.

But I think what you're getting at is that less measurable point, that a knowledgeable worker may be able to do things better than others. Knowledgeable worker-- a loyal worker is going to be willing to do things that non-loyal workers might not be able to. And that is harder to measure.

And if we think about the tasks that AI is going to start to do, more and more of the tasks that are left for people are going to require that knowledgeable, that loyal worker, because, it goes beyond the routine, goes beyond the basic stuff to the really in-depth knowledge of what your company's about.

Do you worry-- this is kind of big picture, but do you worry that AI will widen the divide in companies between workers and managers?

I think it depends on the company. So if you've got a company that's pretty exploitive, it's going to allow more exploitation. I think when we study the best companies, they know that the workers matter. They know that engaging the workers there-- and so you'll see less of that problem happening with the employees and the workers in those companies.

Now, where we could see the divide happen is between the rich companies and the rich countries and the less rich companies, the ones that can invest in AI than have all the data and have all the algorithms, and those that don't. That may be a big divide.

And the big, potential saving grace will be a lot of the vendors, from Adobe to Salesforce to SAP, to whatever, are going to start putting the AI in their tools so it can be used by more people. Could it hurt competition? It's possible. And I think we need to be careful about that.

Yeah, finally, when we talk about jobs, and thinking about your job differently, you've looked at a bunch of companies-- Fidelity, UBS, GE Digital-- that have tried to change, I think, their approach to the way they look at career development, maybe. Do you want to maybe pick one of those and talk about it, or what kinds of things are the companies doing that you see as at the forefront here?

Well, the first thing they're doing is they're not believing those false narratives. They're not going with the easy answer that "We think this makes sense because it makes sense to us at the top of the company." We're not going with that anymore. They're really trying to figure out where they are.

And what they're doing is they're saying, "How can we do what we do for our high potentials but at a lower cost?" [INAUDIBLE] high potentials, we help them with navigation. We help them matching their skills to their jobs. We help them get the learning they need. And we help them also get the feedback on a regular basis. How do we do that?

One of my favorite examples of this is not exactly a high tech company at all. It's a company called Amsted Rail that does really serious big iron things. They make the undercarriages for rail cars and things like this.

And you wouldn't think of them as the most high tech company using the most high techniques, but they have a really interesting career development process. What they do is you have your boss, who manages your performance. And one of your boss's peers is responsible for managing your development. And together--

Like assigning you a mentor.

It's like it's assigning a mentor, but now that mentor has authority. And a mentor, part of their annual bonus is going to be based on how they developed their colleagues, employees.

Interesting.

And so you go up and you talk to that mentor about what do you want to do. They put together a performance plan that they work out with you and your current boss, even if it means leaving your current unit. You and your current boss, that performance plan has goals that you have to meet to develop yourself and goals they have to meet to help you develop.

And at the end of the year, they talk about how your performance was, how you developed, and whether that mentor did their job in helping you get there. And so what they've done is they've aligned all of this. Because mentors are great but the mentors have no power in many companies. In this case, they do.

And it works because I'm going to mentor your people, but then you're going to mentor my people. And so they made this work. Now the question is how much time does it take? And from what I hear, it just doesn't take that much time to do this. If you're a good manager, you set that time aside and you can make it happen. But if you're mentoring 10, 15 people, you can easily do that during a year.

Well, I think what's really interesting about it is people, generally, they manage up, they manage down, but having a diagonal line allows you access to a different person with a different set of experiences. I mean, you talked a little bit before about how do you progress in your career.

And you have limited knowledge. You have limited anecdotal knowledge around you. Opening up a whole other person gives you maybe double-- instead of one manager you talk to regularly, there's two managers, and that doubles the amount of experience and knowledge of the world that you're potentially exposed to, which could matter.

It could really matter. And you want to think about it from the other side too, from how can this be good not only for the employee but also for the companies? There's a tremendous amount of talent in your company that just doesn't have the skills you need.

And so, for example, in medicine right now, it's hard to find the right lab techs. So what two companies I know of are doing-- and most recent, I heard it from Beth Israel here in Boston-- they'll take a receptionist, they'll take a cafeteria person who's a great worker, and I'll say, "Do you want to become a lab tech?"

And they'll say "I can't do that. I don't have the school." "Well, we will help you." And they'll pay them to develop the skills and they'll give them a career. And as they work through the way with community college, whatever, the employee is taking the courses and they're getting better jobs and better pay as they move along the line.

And the idea is you're a great worker. Let's turn you into a great lab person, as opposed to let's hire you as a lab person and hope you're a great worker. So this attention to development is a huge opportunity not only for the employees, but also for the companies themselves.

And this leads, well, I think, into some work that you're doing. You're a founder of an initiative at MIT called the Global Opportunity Forum. It centers around a better approach to the jobs of the future. Do you want to talk a little bit about what the forum is and what motivated you to create it?

So I've been doing digital transformation research for over a decade. And as we get into AI, we start talking about jobs changing. And the thing that's been really clear to me is what's not changing as fast as everything else is the way that HR is training and developing people.

And the good thing about that is that if I do HR better than you, that's not necessarily a competitive secret, usually. And so what we're trying to do with the Global Opportunity Forum is get companies talking to companies about how they're developing their people. What are they doing for learning, for careers, for mentoring, these kinds of things?

And of course, we will develop the case studies. We present research on what it is and let people talk about it. It's an opportunity for every talent and learning leader who is part of this to learn from what we're doing, but also for everybody to share among themselves so everybody gets a little bit better at making this happen.

So it's companies talking to companies about the challenges they've got and the solutions they've got. And we're also taking on some challenges. How do we find the right people? What are data science jobs and how do we find the right people for them?

How do we train these soft skills, these human skills. This is what we're doing in Global Opportunity Forum. Let's try to help every company develop people's careers by getting companies talking to each other.

And I mean, I'm sure you're hoping for great things, but I wonder if it feels at all-- is it like pushing a boulder uphill? Is it hard to change how HR departments think or how leadership thinks?

It's very hard, but there are a lot of HR departments out there. And the people that want to change, we want to make this available to them. People that don't want to change, they can wait and see what happens to everybody else.

George Westerman is a Principal Research Scientist, Senior Lecturer at MIT Sloan School of Management. He's founder of the Global Opportunity Forum. George, thank you so much.

Oh, it was great. Thanks for a good conversation.

[MUSIC PLAYING]

And we'd love to hear from you. Email us with topics that you'd like us to explore in future shows. And let us know if your organization is interested in joining companies from Fidelity to Microsoft in working with MIT researchers on how best to use AI. You can email us at podcast@csail.mit.edu. And CSAIL is C-S-A-I-L So podcast@csail.mit.edu.

I'm Kara Miller. Our show is produced by Matt Purdy and Nate Caldwell with help from Audrey Woods. Tune in next month for a brand new edition of the *CSAIL Alliances* podcast and stay ahead of the curve.