uberknowledge

Helen Patton, CISO, Ohio State University

Bias in AI

Helen talks about the dangers of assumption and generalization in AI and ML and how diversity is key in avoiding this.

02:17	Before the technology gets well developed, we have the opportunity to avoid assumption and bias.
05:10	AI and ML provide a risk that we'll create lack of choice for people.
07:12	You cannot generalize about a generation.
09:12	Diversity is key in avoiding bias: have a diverse workforce and have a variety of data sets.
11:10	We need a human QA check in the process, not for technical accuracy, but for societal accuracy.

Ashwin Krishnan: [00:00:24] So we are at another edition of the Cyber360+ podcast with UberKnowledge. And today I have the pleasure of inviting Helen Patton again, who is the Chief Information Security Officer of the Ohio State University. Welcome back, Helen.

Helen Patton: [00:00:38] Thank you for having me. I'm glad to be here.

Ashwin Krishnan: [00:00:41] Yes, it's good to finally talk to you after our face to face at Black Hat, it seems like months ago.

Helen Patton: [00:00:47] [Laughs] Yes, it does.

Ashwin Krishnan: [00:00:49] So, as we were chatting briefly before we delved into the podcast, clearly there is a lot of discussion, as well as actual evidence, about bias when it comes to training data in artificial intelligence. But there isn't as much talk about, maybe it's happening, maybe it's not getting the coverage, how this kind of bias would affect the areas of privacy and security, if at all. I wanted to kind of touch upon that and get your thoughts.

Helen Patton: [00:01:15] First of all, I would agree. I think this is an area where there are some people who are absolutely geeking out on this in very deep ways; I am not one of those. But when it comes to the community talking about bias in artificial intelligence and security and privacy, I'm not seeing very deep threads around that conversation topic yet.

[00:01:41] For me, I think there are a couple of areas where we need to keep our eyes out for things, for bias in AI. One is as we're developing products and tools that help people work with technology. And when I say that I'm thinking about folks with disabilities or folks with different cultural backgrounds or whatever. What kinds of algorithmic assumptions are we making about how the tools will be used or how the tools will work? That's one area.

"We've got some opportunities before the technology gets really well developed."

[00:02:17] Then on the security side itself, I am concerned that particularly as we go into user behavior analytics and then we're using those analytics to make decisions around access, I think we are going to have to really watch out to make sure that we aren't again embedding some inherent bias or some overt biases into the way we work. So some examples would be things like, you know, if we're going to make an assumption that young people are techno savvy, for example, and therefore are going to be comfortable with a certain kind of authentication mechanism, I think there's going to be problems with that. If we make assumptions around where people work or when people work or how they work, I think there could be some challenges with that, too. So, I think we've got some opportunities before the technology gets really well developed to think about that now, and make sure that we're not setting ourselves up as a society and as a profession for failure later on.

Ashwin Krishnan: [00:03:30] Yeah, that's a great point that you mention. I was thinking myself when you mentioned the demographics, how people like to authenticate

themselves or not based on age, and that's something I've heard so much from the conversations I've had with the vendors.

[00:03:44] Privacy, as an example. I think. Okay, so Gen Zs, as an example, for demographics. Well, they don't really care about privacy and therefore we can treat them as if they don't care about that concept. And then I was like, how do you make that leap of faith? Maybe you have a teenager at home and maybe they just act in a certain way, but that doesn't give you the prerogative or any authority to treat their entire demographic the same way.

Helen Patton: [00:04:05] Yeah, absolutely.

Ashwin Krishnan: [00:04:07] Can you talk a little bit about what you've seen when it comes to — again, this is probably convenient for vendors to do just because this generation, this demographic, we're going to use UBA and therefore make some decisions. So, like you were saying, right now is the time, before we delve into this cesspool, to understand the fact that we have a lot of unlearning to do before we can even go down the path of, let's say, AI-driven security and privacy.

Helen Patton: [00:04:29] Yeah. It's interesting to me at the moment because the things that I notice — which are, of course, just the things that I happen to notice — are where people are justifying making certain technology choices in the way they're designing a product or a process with very much sort of a big brother, or big sister depending on your bias, point of view, like a paternalistic point of view. So, you know, we're doing this because it protects the user or the person at the end of the line, without even asking does that person want to be protected?

"We run the risk of using artificial intelligence and machine learning to put in place lack of choice."

[00:05:10] So, I think we're taking, and quite reasonably, we're taking a very corporate approach to things. We're the corporation, and so we can say what we want to do and really the end user is just going to have to deal with that, particularly if the end user is an employee. If you're a customer, you may have a little bit more say, but often not. But I think we run the risk of using artificial intelligence and machine learning to put in place lack of choice. And I think we have to be very minimalist in understanding when choice should not be required. We've had sort of this approach in security to say the best security is one where people don't have to choose security. It's automatically given to them or automatically required of them or whatever the right adjectives are. And I think when it comes to particularly, again, user behavior analytics and making assumptions about who people are or how they work or where they work, we need to be able to give them choices around that interface. And I don't think we do a great job of that as a profession right now.

Ashwin Krishnan: [00:06:25] Yes. Helen, you have the advantage of being in a very large institution and having the ability to actually interact with a lot of these up and coming generations. So from your perspective, just looking at the same discussion from the lens of the user, what are some of the behavioral changes that you're seeing over the course of, let's say, before your college degree, somebody who's coming as

a freshman and somebody who's graduating? I mean, is there a pattern that you're seeing in terms of their acceptance of technology, the acceptance of privacy practices there?

[00:06:50] I've heard, for example, a few of my podcast guests talk about the fact that there is no allegiance to brands anymore. They go by certain values; they go by the values of the company and the value that they offer to the customer themselves. And then if the value disseminates or disappears, they just move on to the next thing. So again, given your vantage position, getting your perspective would be extremely valuable: How does this next generation of leaders look at this?

Helen Patton: [00:07:12] So, without being a market research scientist, I'm only giving you Helen's opinion on this one. What I see as we communicate about security or technology across the university campus is that you cannot generalize about a generation. I think it comes down to, are people change averse or not. There are certainly Gen Z students who will stick with a product because that's the product they know, and the hassle of moving to something else is too high. Even if the product that they're using just sucks. So yeah, there are some really tech-savvy people who are going to bounce all over the place. But I would not extend that generalization to the entire generation. And similarly, there's a trope that sort of says, they grew up with a phone in their hand and so they must be very comfortable with technology. It doesn't mean that they know technology. It means that they know how to use a phone. But that doesn't mean if we took phones away and replaced it with something else, they'd be super comfortable with that. They just know how to use a phone. So, I think there's some challenges to that. And the more we use artificial intelligence, machine learning to speed up the speed of change, the more uncomfortable I think it's going to make everybody, including the most tech savvy groups we've got.

"As we communicate about security or technology … you cannot generalize about a generation."

Ashwin Krishnan: [00:8:43] Right. Kind of extend that further, so let's say, I'm the CEO of a startup that is looking at this and suddenly comes to the realization, thanks to listening to our podcast, that I need to go back and relook at how I'm looking at these broad demographics and not just look at everybody through the same lens. What suggestions do you have, especially as you talk to so many vendors as well as end users? From your perspective, what are some of the things that, let's say, a startup needs to put in place right now to make sure that they get this foundation right? So as they build these models, they've already thought about these things beforehand.

Helen Patton: [00:09:12] Yeah. So one, I'd say they need to start with a diverse workforce right from the beginning. I know diversity is hard when you've got a workforce of two people in a garage or something like that. But in general, if you're doing a startup, you don't have years of legacy weirdness that would prevent you from having a diverse workforce. So have a diverse workforce so that you can get as

much variety into your thinking as soon as you can, as well as you can.

[00:09:42] I think the other thing is, as we're building on large data sets to generate our AI algorithms and our ML work, be very intentional about where you're getting those data sets from and who created those data sets in the first place and so forth. If you're getting a — I'm going to be very bad here and generalize in a really awful way — but, you know, if you're getting all of your faces from the prison population, that's not necessarily representative of the society at large, and it's certainly not representative of the world at large. So how could you expect not to have inherent bias in what you're doing and what you're teaching your computers to learn, if the sources of information they're getting are biased from the beginning? So, I think those are the two main things that I'd be thinking about.

Ashwin Krishnan: [00:10:37] Yeah, it's a great perspective because the opposite is really what startups are all about, fail fast, and now we have pretrained algorithms that are available, open source. It used to be a hurdle few years ago: startups have this big disadvantage of not having enough data sets. But now you have pretrained algorithms, so it's almost counterintuitive for founders to slope around and start digging deeper, looking at where this prediction is coming from. And that's, I think, super advice, because otherwise you go headlong into this path and figure out you've got it completely upside down or you're making decisions that embarrass you.

Helen Patton: [00:11:10] Yeah. When you see conclusions being reached by the software, I think putting eyes on it that is not itself the technology, that allows you to say, OK, we asked this question up front. We've run it through our sausage-making machine. It's come out with this answer on the other end. Is that a reasonable answer based on what we know and based on what we see the world around us looking like? There needs to be some — and this is Helen with no tech background in this space at all — there does need to be a QA check somewhere in the process for sanity, not for technical accuracy, but for sort of societal accuracy. And I think if we can start building that kind of quality assurance into what we're doing, again, sooner rather than later, the end product will be better for everybody.

"There does need to be a QA check somewhere in the process … for societal accuracy."

Ashwin Krishnan: [00:12:14] That's great, so we can come up with the term, human assurance at the end of the podcast?

Helen Patton: [00:12:18] Yeah, absolutely. What annoys people universally, I think, is when a vendor makes an assumption about them, that is wrong. For whatever reason, someone markets the wrong material to us, or we get served the wrong ad in a web browser, or someone thinks we're one person, but we're really another person, that kind of thing. Nothing's going to lower your market value more than having bad AI. So, you know, we're going to have to fix it.

Ashwin Krishnan: [00:12:53] Correct. And that's a great point. We've come full circle to where we started with user behavior analysis. If you misidentify me and you're actually

poking holes at me as an individual, then nothing angers me more than that.

Helen Patton: [00:13:05] Yeah, right.

Ashwin Krishnan: [00:13:07] That's great. So once again, I think this is a great conversation. I know not a lot has been talked about AI, especially when it comes to security practices, hopefully we scratched the surface and our listeners got some value out of it. I'm looking forward to our conversation next month.

Helen Patton: [00:13:22] Terrific. Thanks so much.

Ashwin Krishnan: [00:13:24] Thanks.