Is Your Organization Ready for AI?

Featuring Dr. Rachel Dzombak and Carol Smith

Welcome to the SEI Podcast Series, a production of the Carnegie Mellon University Software Engineering Institute. The SEI is a federally funded research and development center sponsored by the U.S. Department of Defense. A transcript of today's podcast is posted on the SEI website at sei.cmu.edu/podcasts.

Rachel Dzombak: Hi, and welcome to the SEI Podcast Series. My name is Rachel Dzombak. I am digital transformation lead in the SEI's Emerging Technology Center. Today, I am joined by my colleague Carol Smith to discuss how to determine if your organization is AI-ready, and how to make sure that you are putting things in process to become AI-ready as an organization. Carol, welcome to the podcast.

Carol Smith: Thank you. I am happy to be here.

Rachel: Why don't we start off by telling the audience a little bit about ourselves and the work that we do here at the SEI? I can start just to make it a little bit easier. As I said, my name is Rachel Dzombak. My background is in engineering systems and human-centered design. Before joining the SEI, I was an innovation fellow at UC Berkeley where I helped to grow and define the field of development engineering. Here I am helping our partners think about what is needed from an organizational standpoint to adopt new technologies, as well as helping to lead the SEI's effort around AI engineering, which I am sure we will talk a little bit about today. So, very excited for this conversation here with you today, Carol.

Carol: Excellent. I am Carol Smith. I am here at the Software Engineering Institute as well in the Emerging Technology Center. My focus is on looking at humans and machines. So, I am a senior research scientist working on human-machine interaction, and my background is in human-computer interaction, where I have a master's degree. I have been working in this field for 20 years. Since 2015, I have been working on artificial intelligence and autonomous vehicles and other types of emerging technologies and really looking at that human experience, the user experience, and doing research to improve those situations.

Rachel: Excellent. Thanks, Carol. You bring so much experience to the table and have worked in many different types of organizations. So, I think it will be a great conversation today about really what goes in, from an organizational standpoint, into getting ready to adopt and enable AI systems. First off, I think we should define a little bit about what we mean when we talk about AI-ready in the context of today's environments and the organizations that we work with. I am curious, from your perspective, what are some of the hallmark traits of an AI-ready organization?

Carol: Organizations really need to first understand the problems that AI can solve and look at it from that perspective to make sure that they are actually approaching a problem that is solvable with those systems. Then that next step is really identifying data and making sure that they have the information that would be needed to build this type of system. And, then, looking at the people, of course, and figuring out what people they have, what resources they have, the skills that are needed to make this system a realization or the partners that they can bring into the organization to help them to be supported in their approach to making AI systems.

Rachel: Great. Thanks for that. One of the things you really highlighted in your answer there, Carol, was about that it is not a binary that you are magically AI-ready or you are not. It is more of a journey that you are starting from problem definitions to organizing your data to starting to experiment and prototype some AI systems, all the way up to implementation. You need different people and organizational structures along that journey. I am curious, from your perspective, we talk about AI being this now pervasive technology, but do you have a pulse of where you think most organizations are today in that journey of becoming AI-ready?

Carol: Yes, I think many organizations are very curious and interested in these technologies. I see lots of them trying out new things and trying out solutions for the more common problems. But, I do think that most organizations are still at those early stages and really trying to understand what these systems can do. The AI systems themselves are very early in development and effectiveness, so that also factors into how useful they can even be once they are instituted. Probably the most helpful ways I see AI working right now are in the simpler chatbot-type solutions. Certainly in analysis and looking at connections between pieces of information and things like that, where both the problem is very clearly defined and the amount of data that is needed is available and is there and usable, and that those systems are really good systems for AI to solve, that there is nice patterns to follow, that sort of thing.

Rachel: Absolutely, I think something I know we talk a lot about is this mismatch between people's expectations of where the technology is today and the applications they want to use AI for. And, so, I imagine a big challenge that a lot of companies have today is just unpacking where are they actually and being honest with themselves. Where are they in their organizational readiness? Where is the technology in its organizational readiness, as well to make sure that there is alignment between those two sides of things.

Carol: Definitely. Yes.

Rachel: One of the reasons we are having this conversation is because we know through our work in AI engineering that having systems in place to support AI's development and adoption is critical. It just can't be a system like you purchase a cell phone from the Apple Store, and you are ready to go. You really need to be paying attention to how you are designing, developing, and implementing AI systems. It relates to our work in AI engineering, for those that aren't familiar. AI engineering is a field of research and practice that integrates the principles of software engineering systems, computer science, and human-centered design, to create AI systems in accordance with human needs for mission outcomes. That field of AI engineering that we are trying to grill and spark conversation around, is about exactly this. That it is not just the technology itself, it is the implementation of it and what that looks like. That is very critical to pay attention to. I feel like it is worth keeping this conversation grounded in what is different about implementing AI versus other software solutions.

Carol: That goes back to that long-term issue that you brought up earlier with the maintenance piece. Maintenance probably isn't even the best term, but these systems aren't ones that you can set and forget. These aren't systems that simply are installed and maintain themselves over periods of time. Rather, they are in a constant state of change, because new data needs to be added to them. Once that new data is added, the actual system will change and be affected by that. That may significantly affect the results that are made, or the recommendations the system might make. Another big difference is just the amount of infrastructure that is needed. The types of systems that are needed to maintain the speed at which people want these systems to work—there needs to be a lot of computing power behind them. Also, just maintaining a system that is robust, that is secure, that has the trust because of its ability to respond in a timely manner, that sort of thing. Those types of activities require people who know how to continue to tweak and maintain the system. So, there is a lot of ongoing and scaling that needs to occur within the organization to keep these systems healthy and maintained.

Rachel: It is almost as if the organization needs to be growing, adapting, evolving right in parallel with the way that the system is growing, adapting, and evolving over time.

Carol: Yes.

Rachel: It sounds like, from what you are saying, most organizations are focused on the technology's evolution and less about their own scale on the workforce side to accommodate that.

Carol: Yes. I think in a lot of cases, people make the assumption that an AI system is actually going to reduce the work. What instead it can do is augment certainly. Certain activities can be simplified, sped up. And, those individual's jobs may change dramatically, but it also means that other individuals need to pick up more work and do more maintenance, do more maintaining of the data, and the system architecture, all those sorts of things. Those skills are in more need as the systems continue to grow and be used.

Rachel: Yes, and especially short term, there is a huge investment, both in terms of money, people, that needs to go into getting it to a place where it is even ready to be operationalized in some way, shape, or form.

Carol: Most of these systems take months to build, weeks to months at least. For highly specific and broader sets of information or problems that need to be solved, they can take even longer. That in itself is a big commitment when potentially adding more people may solve the same problem much more quickly and at much lower cost.

Rachel: Interesting. On that note, you had spoken earlier about the triaging that is necessary for organizations to do to really reflect on, *Is this a problem that's important to me?* And, *Is this problem one that makes sense for AI to solve instead of having*, as you just said, *hiring additional people and have them address that challenge as opposed to kind of task shifting it over to technology?* So, when you think about implementing AI and organizations, how do you approach identifying problems that can be solved using AI?

Carol: Yes, the first question is about data. It should be about data anyway; really looking to see that that data is in existence, is accessible, is something that the organization has control of, and access to, making sure that it is the right data. That it is not overly biased in one direction or another with regard to what they are actually trying to do, and then thinking through also the implications of the system. Is putting the system together, is bringing these sets of data together, going to put people at risk in a variety of different ways? And really being imaginative about those kinds of implications as well. So, lots of early thinking before any programming, any algorithms are put together; really looking at that, to see if the initial pieces are there. Then, once we are assured that the data and the problem is the right one and everything is in place from that point of view, then it's, *Do we have the people who can actually build the product available?* And, *How do we match them with, potentially, the subject matter experts that are needed to be able to ensure that the system is giving the right information back and that it's answering questions in the right way, whatever that topic area is?* And continuing along, thinking through the actual building of it, *Are we building it for the right people? Are the people who are going to be using the system going to understand the system? Are they going to be able to use it in a way that really is going to augment their work and make their lives easier, hopefully, and protect them as well as the people who might be affected by the system?*

Rachel: I really loved the comments you made about the hard conversations that need to happen up front, you know, in organizing your data, really questioning, *In what ways is our data biased?* Not, *Is our data biased?* Because we know that all data is biased in some way. And questions about the user of, *Are we building it for the right people?* as well as questions about unintended consequences? We know from our experience that a lot of times those conversations are skipped, that people jump right into the building instead of kind of doing that up-front work. And I am curious, from your perspective, why do you think that is?

Carol: Yes, I think part of it is just enthusiasm. People are really excited. AI seems like it is the utopia. *It's going to fix things. We won't have to worry about the bias or the problems that we are having because the AI will fix it*. But, unfortunately, it just doesn't do those types of things. AI can't solve people problems. It can't solve social problems. We have to do that work ourselves, and it also can't fix problems in the data specifically. That is why those critical questions, those difficult conversations, the really hard and sometimes really unpleasant conversations, have to occur early so that you don't unintentionally build something that is really hugely problematic and a waste of money too and time, but rather thinking through ahead of time to make sure that you are building the right thing for the right reasons.

Rachel: Absolutely. So much of that is dependent on, *Is there a psychological safety in the organization to facilitate those conversations? Are there people that are in roles that know how to facilitate those hard conversations?* We know that is often overlooked or taken for granted in organizations but something that is just so critical for success. Did you have something?

Carol: I was just going to add to that.

Rachel: Go for it.

Carol: Right. And, along that same line, having a diverse group of people who are going to come at the problem in different ways, with different ideas and different experiences. Bringing people with not only different skill sets, but literally different education experiences, different life experiences. And bringing them all together is a point at which you need to have a set of ethics to guide them, to bring them together at the same point, and from the same perspective. But also making sure, as you mentioned, that psychological safety is there so that they do feel that they can be themselves. That they can ask tough questions, that they are not going to be seen as challenging in a negative way, but rather, doing the right thing. Asking the questions up front. Making good trouble, if you will, and doing the things that are needed to be done to make really good software. But, that does require leadership to be in place that is supportive of that work and, ideally, that is diverse itself and that can reflect the diversity of the workforce, so that they stay around. Retention is a big problem in technology in general. Certainly, with a more diverse workforce, you are going to have even more challenges there. And so having that leadership available and visible is really important.

Rachel: Absolutely. I love all of what you just said. I think something that is really important too, is that, going back to our conversation about what types of problems you are going to apply AI to, that diversity is so important always, and particularly important when you are working on higher stakes applications for AI. I think a lot of us forget that these technologies were developed for a lot of low-stakes applications, things like recommending where I should get my lunch from. If that decision goes poorly, it is a bummer, but it is not necessarily going to affect human outcomes, or at least not on the surface. As we are getting into these higher stakes applications, as organizations are starting to look at new areas where they can apply AI, all that diversity is even more critical to make sure you are able to imagine the full scope of possibilities.

So, something that you brought up earlier was that the starting place for organizations when they are really thinking about implementing AI is getting their mind around the data and thinking clearly about what do they need to do to collect data, organize data. I am curious, from your perspective, have you seen any strategies implemented that really set organizations up for success when they think about data?

Carol: Yes, so certainly understanding that problem, understanding the people who are needing to use that data, understanding where they are coming from, and the types of questions they are going to be asking about that data, or the types of recommendations that they are going to need. How they are going to be in interacting with that system and using what is created from that system is really key and does take a significant investment in that learning, but pays dividends when it comes to the solution. Because if you build the right thing, then the people will use it. They are more likely to trust it because it is providing them with the information that they expect. Along with that comes their viewing that data in as early a stage as possible and making sure that really, that is the right data. Having the subject matter experts really look at the source of the data and confirm that it is the best possible source or at least the accuracy is there that they would expect.

Rachel: In many ways, what you are saying is also broadening the notion of what data is worth collecting. That it is collecting the data, kind of the set of examples that you are going to use to train the system, but also the qualitative data of what it actually looks like in practice to make the decision or achieve the outcome that you are wanting to achieve, and investing in both sides of that to make sure that you are collecting both in parallel.

To change directions a little bit. The recent [National Security Commission on Artificial Intelligence] announced the goal of an AI-ready force, stating that AI is poised to transform every industry and is expected to impact every corner of the department, spanning operations training, sustainment, force protection…goes on and on. Carol, your work has shown that for the DoD to succeed in AI implementation, new skills and new roles are needed. Could you talk a little bit about your concept of the curiosity expert and what that looks like?

Carol: Yes. Sure. These are the individuals who really are focused on understanding people and problems. It is not that everyone else isn't curious, but rather that these are the individuals whose job it is to really focus on those areas and to really look at, *What is the individual's situation? What are the needs that they have, and how can the system potentially meet those needs?* So, these people tend to have backgrounds such as I do, in user experience and human-computer interaction. Some of them are cognitive psychologists or digital anthropologists. But, the point is that they really are thinking about how the system is going to be used. How the system is going to work in relation to other people. Looking at human-machine teaming, for example, when it comes to robotics and other types of systems where humans may do part of the work and the system may do the other part of the work. So, a situation where an analyst is doing part of the work, but the system is also doing some of that work, and how are they partnering together? *What kind of information does each of them have? And how can you build the system in a way that makes it trustworthy?* Not overly trusting, but rather appropriately trustworthy and that the people using the system find it to be helpful. And that they're going to accept the system and use the system as its intended versus rejecting the system for a variety of reasons.

Rachel: Something I love about the notion of the curiosity expert is also that it is defined broadly. That that work can be done by people under multiple different roles or different types of training. I think that also helps the DoD because it is not just, you have to hire one very specific type of person, but it is about looking to the workforce, looking to applicants, and saying, *Who can do this work?* But not losing sight of what the true need is, to unpack, *What are human experiences? What outcomes do they want to have?* And the different ways you might be able to get there.

I am curious, from your perspective, what needs—what linkages need to be in place between the curiosity experts who are doing that user-focused work and the development team that is building the system, writing the code, and trying to instantiate it? How do you make sure that there is connection between those sides, and is it necessary to have?

Carol: Yes, definitely. It is definitely necessary for them all to be working together and to all understand the problem that they are solving, so that they really are working towards the same goals. Part of this is really having these people be embedded in the team, at least on a partial basis. So, sharing the information that they learn, having some of the machine-learning experts and system architects—whoever else is on the team—observe the research, so that they understand what is being done and the types of questions that are being asked. And also seeing the responses from those end users. Another great activity is usability testing of early prototypes.

So, making an early prototype of an AI system is difficult, but it is doable. Creating that system in such a way that can be tested by the individuals who will end up using it. And, having the people who are actually building the system observe those sessions as well is extremely helpful and can build a sense of sympathy, maybe even empathy, for the way that they may struggle using the system and the types of questions they are asking, the types of information that they don't see the system providing and, hopefully, that lead to improvements in the system.

Rachel: So, even if you are not a curiosity expert or haven't been trained in HCI [human-computer interaction]. Say I'm a developer: if I can just observe that research happening, it can shift how I see the development informed design decisions down the road as well.

Carol: Definitely, definitely. There are some techniques that anyone can try. Certainly, someone who is trained in the series could do a better job and get more effective and impactful results. But, certainly, anyone who is interested should be going out and understanding the users and not just showing the work that they are doing, but rather listening and observing and asking questions to understand the other individual's experience, so that they are not just building what they think should be built, but rather what is needed by the person who will be using that.

Rachel: Absolutely. I love that notion that it is not just up to the curiosity experts, and there is no magic skill set that they have, but it is really that everyone…. Certainly there are things you can learn from people who have been deeply trained in it. But it can be successful to everyone to just go and ask, *What kinds of—what types of problems do you have? How are you looking to use this?* to grow their familiarity with users themselves. As we look ahead and think about implementation of AI across the DoD, what do you think people aren't thinking enough about? Where do you see, kind of, existing gaps in conversation about AI implementation?

Carol: Yes, certainly that long-term piece, the idea that these systems need to be monitored. The human needs to be in the loop and responsible for the decisions, the recommendations that these systems are making. It is certainly being talked about more, but I think that that is not something that's always understood initially when people are thinking about making an AI system. Often, I think people assume or hope that an AI system is going to stand alone and be responsible for itself, but systems don't have rights and responsibilities. Giving that kind of responsibility to an AI system is wrong. A human really needs to be responsible, particularly for significant decisions. Any significant decision that would affect a human's life, their well-being, their way of earning money—anything along those lines certainly should be a situation where a single human is responsible for those types of decisions that the AI system might make.

Rachel: And it is really then for the workforce, shifting their mindset from kind of a linear mindset of, *If I do x, y will happen*, to being open to there being emergent effects also. That we can predict most of what will happen, but we always need to be looking for unexpected consequences or other outcomes that could emerge from the system itself. I think that mindset shift is really hard for a lot of people and a major challenge to growing this AI-ready workforce as well. To start thinking about, *OK, how do I just get into a place of anticipating those outcomes?*

Carol: Part of the challenge also is building the systems in a way where if there is an unexpected result, it is explained. That it is clear what happened, that there is an obvious way for someone to investigate that more fully. It may not be the end user. It may be something that someone with more machine-learning knowledge or something along those lines needs to look into. But, there should be ways to find that information out. People continuously use terms that indicate that these systems are somehow magical and that they can't control them. That is not the way the system should be built. They need to be built with that idea of responsibility as well with regard to what might happen. If you are not sure what is happening, then you need to take responsibility for that and make sure that that is maintained in a way that people do, then understand what is going on. Or, the system is shut down if it is not being reliable.

Rachel: Absolutely, so much of what is talked about is, *Oh, the system's opaque. I have no control over that. It is just the way it is.* I think a lot of what our work has shown is that that is just insufficient for high-stakes applications. If that is the case, then it is not going to be adopted. If we can let go of that a little bit, we can start to challenge ourselves to say, *How do we develop systems that are transparent, that I can find the off switch for*, and not just say, *I have no control, therefore, this is going to happen no matter what*. So, let's look to transition. If I am leading an effort to adopt AI within my organization, are there any go-to resources that you would recommend starting with?

Carol: Yes. There are some wonderful resources out there. There are papers, such as datasheets for data sets, and model cards for models that really help people to think through the types of questions they should be asking about the data, the provenance of the data, the way the data was collected, the breadth of it, and really starting that work to look at the information that they are putting together in a better way. The model cards are very similar in that they have you really asking very specific questions about the work that is going to be done to build this AI system and how those decisions are being made and what the end goal is for that work.

There also sets of ethics that I feel are very important to help align various people from different backgrounds, people with different experiences, bringing them all together on a set of ethics. And there are checklists, including the checklists that we'll link to from this podcast, that really support that work. Often, the sets of ethics may be very vague, and the people using them need some support in order to really be able to enact them, to be able to apply them to the problem that they're trying to solve. So, by using these types of tools, you can begin having those really important conversations, the critical conversations that are going to lead to better work. A lot of this really is just opening up that communication line, making sure that people are actually talking about the work that they're doing. Not just the typing, but rather the intent, the work that they are trying to do and the effects it will have on the world. Because these AI systems are just so much more powerful, bring together so many different types of data, and potentially can affect so many more people. We need to be more critical about the work and take responsibility for those conversations and the work that they create.

Rachel: Absolutely, I love your notion on the resources being a support tool to start critical lines of conversation. The other ones that you made me think of are—there are some great resources available on just the basics of design and systems thinking that I think are super accessible. And, if you have teammates that, or you, are trying to start AI and people are super focused on the solution, not on the problem space as much, or the unintended consequences that Carol was really talking about. I think just gaining some language and familiarity with those two fields can really help people to understand why it's necessary to do that up-front work. Why you need to be reflective of the ethics, the questions in your data sets, the questions in your model. Why you need to be interrogating those up front. We can link a couple of resources there as well.

Carol: The mindset that people have as they are approaching these problems is a key. Really, being curious, being speculative about the potential dangers that could occur, and really thinking thoroughly about the work that they are doing.

Rachel: Absolutely. Well, with that, I want to thank you so much, Carol, for being here with me today and engaging in this conversation. To our listeners, thank you for joining. We will include links in our transcript to all resources mentioned in this podcast. Please don't hesitate to reach out to us if you have any questions, either to Carol and I or, to email us at info@sei.cmu.edu. Thank you.

Thanks for joining us. This episode is available where you download podcasts, including SoundCloud, Stitcher, TuneIn Radio, Google Podcasts, and Apple Podcasts. It is also available on the SEI website at sei.cmu.edu/podcasts and the SEI's YouTube channel. This copyrighted work is made available through the Software Engineering Institute, a federally funded research and development center sponsored by the U.S. Department of Defense. For more information about the SEI and this work, please visit www.sei.cmu.edu. As always, if you have any questions, please don't hesitate to email us at info@sei.cmu.edu. Thank you.