

SEI Podcasts

Conversations in Artificial Intelligence,
Cybersecurity, and Software Engineering

The Importance of Diversity in Artificial Intelligence: Violet Turri

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.

Palma Buttles-Valdez: Welcome to the SEI's Podcast Series, a production of the Carnegie Mellon University Software Engineering Institute. My name is [Palma Buttles-Valdez](#), and I am the director of the SEI Office of Diversity, Equity, and Inclusion. Joining me today is [Violet Turri](#), a software developer and researcher in the SEI Artificial Intelligence Division. Today, we are here to talk about Violet's experience and work in artificial intelligence [AI] and the importance of diversity in this field. Welcome, Violet.

Violet Turri: Thank you. It is great to be here.

Palma: You have been a guest on our podcast series before to talk about [your work in explainable artificial intelligence \(XAI\)](#). Let's begin by having you tell our audience about yourself and the work that you do here and also include what to you is the best part of your job at the SEI.

Violet: I am a software developer and researcher here in the AI Division. I have been working here for about three and a half years. I have shared

about explainable AI research we have been working on in the past. Currently, I'm leading a project that's looking at how we can improve the transparency of real-world AI systems, especially in the computer vision space. That's been the work I have been doing most recently. But I have also gotten the chance to work on testing and evaluation projects and other kinds of responsibility AI work. I think the best part of this job I get to work with some really awesome people with a lot of different kinds of skill sets who are enthusiastic. So that's just great having that kind of environment and having people like that on my team. Another thing I really like about working at the SEI is being able to investigate questions that I think are interesting and that matter to me taking on challenging problems. Yes, it is a very supportive and enthusiastic environment.

Palma: Before we dig deeper on the importance of diversity and artificial intelligence, I would like to talk about your early years. Who were your early mentors, and what specific actions did they take to build up and inspire you to pursue a career in computer science, or more recently in AI? When did you visualize yourself as being a computer scientist or working in AI?

Violet: I think my parents have definitely been good mentors for me along the way, and countless teachers have been really supportive. My mom was a software engineer for a number of years. Hearing her talk about her experiences and the things she liked about that work, namely, getting to work on interesting problems and solving these little puzzles that come up when you are coding. I thought that sounded really interesting, and it sparked my interest. Also, early on in high school, I took [a course at Pitt](#) that was just an intro to CS class. I think once I had done that, I started really seeing myself in this career and feeling like *OK, this is something I can do that that could be really interesting*. As far as AI goes, I didn't really know much about AI until I was probably a junior in college, and I got to take a class that was about the foundations of artificial intelligence. I also got to take a class about computer vision as well. But I think when I really started to become interested in AI more seriously was [when] I worked on a research project in the [\[Carnegie Mellon University\] Human-Computer Interaction Institute](#) one summer, the summer of my junior year. The project was really about using AI as design material. So thinking about how we build effective systems, using AI, how it can enhance the user's experience. That really exposed me to kind of all the opportunities related to AI, but also how new this area is, and how many open questions there are. Once I learned more about the domain, it really sparked my interest, and I felt like it was something I wanted to dig into deeper.

Palma: You have touched a little bit on the question I was going to ask you about when you wanted to pursue a career in computer science. Was there like a that light bulb moment for you during your classes? Something that just clicked and said, *This is what I want to do as a career.*

Violet: I am not sure if there was an exact moment where I felt that way. I think the more I did it, the more comfortable I got with the material, the more I felt confident and like I was enjoying it. I think it can be a little nerve-wracking at first when you are trying to learn a coding language, and you don't have experience with it. I think over time my interest grew. There was a course that I took in college. It was an algorithms course. Both the professor and this newer professor who were teaching the class were women. I think just having them as the leaders for the course really made me feel comfortable. I remember really engaging with that material in a big way, more than I had in previous classes. Also, a lot of the previous courses I took had been looking at the topic of computer science from a very mathematical way. But, in algorithms class, they give you a specific problem. You can imagine the whole scene. Then you are entering that scene and figuring out the solution. I think that was another aspect of that course that made me realize, *Here is what I am doing in these classes, but here is how it could actually affect a real-world problem.* Actually, reflecting on it, I think that might have been the moment when I thought, *OK, I could really do this. It could be really interesting to see how I can take these skills I have been growing and apply them to real-world problems.*

Palma: It is interesting that you mentioned that your mother was a role model and encouraged you to pursue this. Then you had teachers that were women as well. Oftentimes research shows that when other women, students or women see women teaching or as professors, that it does allow them to visualize themselves as a computer scientist or someone working in AI. Would you say that experience did help to say, *I could do this as well?*

Violet: Yes, definitely. I think having women mentors was super helpful, and it helped me envision my future and made me enthusiastic about this. Yes, I think that opened a lot of a lot of doors for me. Also, as I studied computer science, I built up this really nice friend group of other women who were studying computer science. I got to work closely with them on projects for studying. I think that also helped me feel really comfortable. I felt like we were all kind of on this journey together. It was really cool.

Palma: Now you can serve as a role model for women that are going to be pursuing this as well, so it is coming around full circle as well. Now I am going

to talk a little bit about something, and diving into diversity. Early in 2023, a [Black woman who was eight months pregnant in Detroit was falsely arrested for carjacking by police who had used facial recognition technology](#). At the time of her arrest, she was the sixth such person to be falsely identified by the AI system. All six people had been Black. There has been less catastrophic harm caused by bias in AI systems. Reuters [reported](#) that Amazon discontinued an AI system it was using to review job applicants after learning it downgraded female applicants. All of this leans towards having AI have some bias in the system. Why do you think diversity is so important in the field of AI and knowing some of these examples we just shared?

Violet: When you are working with AI systems, there are a lot of places where bias can enter the system. I think sometimes from the very design of the system itself, there can be some issues there. You really need to think about what the best-case and worst-case scenarios are for the system. In the case of using facial recognition technology or having AI help process resumes, you could say it might improve the efficiency of the process or the accuracy. You might be able to cut costs in some way, so those are all benefits. But there are also a lot of risks that it is important to dive into. What happens if bias seeps into one of those systems, or if you are using historical data to train your models, how are you going to ensure that there are biases within that data set that will be perpetuated by the system that you build? Furthermore, even when you have a trained model, there are trade-offs that can take place. There are metrics that you need to track to make sure that it is going to work for diverse end users or on diverse kinds of data. I think AI systems, with all technology, there is the possibility that the system can work better for some groups than others. But, with AI, it can be a really volatile situation. If something goes wrong from system design to the data set collection to the performance metrics, you can miss things, and you can have issues basically where you have bias baked in in a big way, and AI learns that behavior. It is really important that you have people who are cognizant of these biases and can be advocates to make sure that this is something we are all thinking about, and we are all working on as we build the system. Long story short, I think diversity is extremely important in building AI systems because you need to have diverse stakeholders be part of this process. I think it really helps when you have people of different backgrounds have a stake in that and a seat at the table, so that we make sure that these concerns or considerations are not things that we think about at the end, but that you have people talking about all the way through. Obviously, you don't have to be from a background that's underrepresented in STEM to think about how the system is going to work for diverse stakeholders, but I think it can be really helpful to have people with different lived experiences there to

voice those concerns, and to be cognizant of ways that AI can either serve or harm people.

Palma: In your opinion, how can we address the lack of diversity in the field of AI?

Violet: I have a couple of thoughts. I think it is a challenging topic, just because there are ongoing issues with diversity or lack of diversity and STEM. It can be hard. It can be cyclical. My thought, from personal experience, one is that I think it is really important to build spaces for groups that are underrepresented, somewhere they feel comfortable, and they feel capable. For me, as we discussed, I had these awesome women mentors that really made me feel like I could see myself doing this career. I think that is just really essential is helping people connect with those mentors and see examples of people who are like them, approaching these kinds of problems.

Another thing is, I have read before that women sometimes tend towards jobs where they feel like they are having a positive impact. I know that is something I have felt as well. I think sometimes when you are working on an AI problem or a computer science problem, you might be kind of tackling a really small part of a larger system. It can be hard to tell how what you are working on ultimately has an impact. I think a couple of ideas there. One is to get involved with projects where possible that you are interested in, where you feel like you are building systems that matter to you. Another thing I think that I felt is just reframing this a bit and thinking about how just by being a woman in STEM, you are potentially opening doors for other people, you are showing them what sort of futures they could have. In two regards you can make the work more positive or have a positive impact. One is finding projects that matter to you, and the other is just thinking about your impact as a member of an underrepresented group in STEM and ways you can connect and kind of grow other people as well.

Palma: Great. In your work, you focus on responsible, trustworthy, and [explainable AI](#). Can you talk about your work through the lens of diversity?

Violet: Yes. Responsible AI is looking at how we can build out best practices for engineers and designers, so that they can build systems that serve people basically from the ground up, paying attention to issues like data privacy, thinking about how systems will work for diverse stakeholders, etc. It definitely has clear implications for how well systems will work for diverse groups. There is that expression, *garbage in, garbage out*. If you are not working with quality data on quality systems, you are not going to build

things that are high quality for people. But then, more specifically with transparency work, a lot of AI models, especially ones that can mimic human performance in some way, like computer vision models or large language models, these tend to be deep learning architectures. It is really hard to say what is going on under the hood with these systems. There are a lot of risks that you encounter as a result where you know how the AI performs on certain inputs and what outputs are produced, but you don't understand how it actually made those decisions. Related to the topic of diversity, it is really important to thoroughly vet systems and make sure that there isn't a biased reason why it is reaching these outputs. This work is really kind of a way to counter those issues that we face when we are using data sets that contain bias. It gives us a lot more information about what is going on with the model so that we can make sure we are training models that actually behave as intended for diverse groups.

Palma: Violet, you work at the SEI, which is a federally funded research and development center. What drew you to work within the SEI and at an FFRDC?

Violet: Throughout college I had worked on different research projects at CMU, and I really liked that experience of getting to learn about problems, read past work, think about new ways to approach those challenges. I thought it was a really exciting opportunity to continue doing that kind of work. Then also, I think it has been really interesting and exciting working on projects that have impact for the federal government because it is a great way to give back. It is a great way to learn. I feel like the landscape of AI, for example, is shifting so rapidly, and there are so many opportunities for the government to think about AI, to use ai, develop regulation etcetera. So it is a great place also to be doing responsible AI work. I think both the ability to do research and then also to work on interesting problems that are high impact, that impact the government, is really exciting and drew me to work at an FFRDC.

Palma: We are glad you chose the SEI. Violet, thank you for taking the time to talk with us today. We will include links in the transcript to resources mentioned during this podcast as well as links to your work. Finally, a reminder to our audience that our podcasts are available every place you download podcasts as well as the [SEI's YouTube channel](#). If you like what you see and hear today, please give us a thumbs up. Violet, thank you so much for joining us.

Violet: Thanks so much for having me.

Thanks for joining us. This episode is available where you download podcasts, including [SoundCloud](#), [Spotify](#), 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 do not hesitate to email us at info@sei.cmu.edu. Thank you.