MIT CSAIL Alliances | Stefanie_Mueller_Project_4

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

[MUSIC PLAYING]

Today, we're going to take a sneak peek into the work of a pioneering researcher at CSAIL. We'll do this from time to time so that you get to find out what's new, what's next, what's coming around the bend. And today, we've got a twist, you could say, on generative AI.

So GenAI has already shown it can do things like create photorealistic images, summarize long documents, maybe even write ad copy. But can it make a coffee mug?

And it actually turns out, if you use existing generative AI technology specifically for anything that should function in the physical world, there are actually a lot of unsolved problems.

That's Stefanie Mueller. She's a professor of electrical engineering and computer science at MIT with a joint appointment in mechanical engineering, and she's head of the Human-Computer Interaction engineering group.

Yeah, it's really funny, actually, because a lot of people have probably seen generative AI where you can say something like, please make me a picture of a mug or a cup or something. And you will actually get a picture of this, and it looks good. But if you actually ask generative AI to make you a mug or cup that you can then maybe 3D-print and use, it's actually not going to work.

That's because today's generative AI models don't understand the laws of physics.

So maybe the handle is going to be teeny-tiny, or it's not going to hold the weight of the cup. Or there might actually be some small holes in the bottom, so maybe your coffee will come out at the bottom. So, yeah, not the best. So generative AI is really not optimized right now to check for these physical, functional aspects.

So how do you teach AI physics? Mueller's team turned to a technique that engineers have been using for decades. It's called mechanical simulation.

You may notice from maybe when a car is being designed, if some wind simulation, how would the car perform on a road? The same when you build a bridge. They would check if a car drives over it, will it break? So mechanical simulation can be really helpful to check for these things. And what we do is we integrate this into the generative AI process so that AI can get access to this additional information on the physics behind what it's generating.

Mueller's team has created a new method called mech style-- mech for, M-E-C-H, mechanical-- that combines generative AI with mechanical simulation. It doesn't build a structurally sound 3D model of, let's say, a pair of eyeglass frames from scratch. But if you've already got a 3D model of them and you want to maybe spice them up a little bit, give them a personalized style, let's say a texture like fish scales, well, mech style will make sure that the fish scales don't introduce any structural issues into the frames.

In the project, we actually got a huge data set of 3D-printable models from an online repository. And first, actually, we checked, do they actually break when we use generative AI to stylize them with a personal style? And, actually, a lot of models broke. I think it was 80% or something along those lines.

And after our approach where we integrate the mechanical simulation, we can show that the results are significantly better. So the majority of the models is now able to withhold much larger stresses on it, and it's not going to break so easily. So you get the benefits of both, right? You get the beautiful stylized output according to your personal preference, while your object remains functional.

What are some of the use cases for companies either that you can imagine or maybe that you hope that this will be used for?

Yeah. So there's lots of potential. If you think about the products we buy these days, typically, as a customer, you can just pick from a set of predefined few options that designer handmade. But this generative AI technology really has the potential that everybody, every single person, can get something that is specifically made for them. That just doesn't scale If you have a human designer sitting down and making it. That would be way too expensive.

So I think there's a huge potential for companies to use generative AI to make more tailored products for their customers. But, of course, we would have to do the same checking that you would normally do on a handmade product to make sure it's actually functional and working.

So the idea is, let's say, that you could make something your style or that you could multiply the number of styles, the number of different textures or patterns that might be out there, that we could see, I guess, so many more options. And right now, I assume, we only see a few, right?

Yeah, that's right. Typically, maybe you have some product customization options these days. If you go to a shop and you buy a phone, you can pick between three colors.

Right.

But that's it. So that's a really sad state that we have right now, and there's obviously way more potential to personalize, especially this is connected to the 3D printing opportunity. Because 3D printing can create these one-off objects which was not possible with mass manufacturing before.

So now this combination of 3D printing technology and generative AI is actually really making this happen right now. So I'm really excited about this opportunity.

And I know in the past, you've had people come up with their own style. So maybe somebody says, let's say, I'd love to see this in a Moroccan mosaic style. I don't really see that out there in the marketplace. But can we do that?

Yeah. So the great part is that in normal product design, the designer really have to know something about 3D modeling and so on. That takes years of training. But with the new generative AI, you can just provide a text input. And speaking out what you would like to have, that's something that everybody can do.

So now it actually gives all the power to the customer, basically, to describe exactly what they want to have and then get the product that they were desiring. So that's great.

So if someone was like, yes, I love that style, that Moroccan mosaic, why would implementing that style make a cup, let's say, any less structurally sound than just a plain red cup that you could get at the store?

Yeah, that's a great question. And the question is, how do you define style? So in our world, we define style as not only color, but also the surface texture. So if you describe you want something that looks more like mosaic glass versus a Moroccan style, they might actually have a different surface texture as well.

So at the moment, the surface changes the walls of the object. If you just talk about a cup, maybe the walls of the cup get thinner in certain locations, and that makes the object more fragile. So there's really this trade-off between approximating the style that the customer wants and still making sure that the functional aspects are preserved.

What do you feel like the biggest challenge was when you were trying to inject some of the rules of physics into how the AI works?

Yeah. So the interesting part here really is the way generative AI works is that it runs many, many iterations where it's, after each iteration, approximating more and more the final desired output. So it actually starts with something very coarse and then works more and more closely to what you actually want it to have.

So if you now actually want to integrate mechanical simulation, you would have to run that after every change and new version that the AI process makes. But simulation in general is very, very slow. So a new generative AI version may run in two seconds, but a simulation may take several minutes. So it's really this question, how often should we run the simulation to not slow down, giving you your desired result, while still running it often enough that we can check if things are still functional as they are being generated?

So your research is pretty new. Have you gotten feedback on it? Since little bits of it have been out there, what have you heard?

Yeah. I'm actually super excited that there are so many people interested in this. So this is a completely new domain in generative AI. There's been a huge interest in the past in images and text because there are a lot of databases on this. For 3D models, and especially 3D printed models, there hasn't been that much data available. So it's really only now that people are looking into this.

So, yeah, we are just tapping into a hot new area, and we're getting lots of requests for collaboration. And this area will grow massively in the next years because there's just so much invention potential still in this space.

This was great. Stefanie Mueller, thank you so much.

Yeah, thank you for having me.

[MUSIC PLAYING]

Stefanie Mueller is the head of the Human-Computer Interaction group. She's a professor of electrical engineering and computer science at MIT with a joint appointment in mechanical engineering. And before we go here, a note about an upcoming cybersecurity course from CSAIL.

Cybersecurity executive positions are among the fastest-growing roles. They're projected to have a 13% growth this decade. We've got a cybersecurity for technical leaders course coming up. It's been developed by MIT CSAIL and MIT xPRO. It's delivered by Simplilearn.

Podcast listeners can save 10%. So if you are interested, email us. It's podcast@csail.mit. That's CSAIL, C-S-A-I-L. So podcast@csail.mit.edu. Tune in next time for a brand-new edition of the *CSAIL Alliances* podcast and stay ahead of the curve.