**Yale Talk: Conversations with Peter Salovey**
*<u>Episode 37</u>: Promoting representation in animation*
*<u>Guest</u>: Theodore Kim, associate professor of computer science in the Yale School of Engineering & Applied Science and co-leader of the Yale Computer Graphics Group*
*<u>Description</u>: Professor Theodore Kim, a former senior research scientist at Pixar and two-time Academy Award winner, discusses the Yale Computer Graphics Group's work to counter racial bias in animation technology with President Peter Salovey.*
*<u>Publish Date</u>: April 12, 2023*

FULL TRANSCRIPT

**Peter:** Hello, everyone. I'm Peter Salovey and welcome to Yale Talk. At Yale, our education and research mission calls us to cultivate multidisciplinary collaboration to bridge different areas of expertise. By gaining new perspectives, we can ignite new forms of scholarship. So insights from one field can drive breakthroughs in another. Foundational to Yale's mission is the faculty. Faculty members foster inquiry across disciplinary boundaries, and Ted Kim is one such scholar. Ted is an associate professor of computer science at the Yale School of Engineering and Applied Science, and his research blends the visual arts, computer science, and ethics, and it holds the promise of countering insidious racial bias in computer graphics. Prior to joining the Yale community in 2019, Ted was a senior research scientist at Pixar Research, where he received screen credits in blockbuster animation films such as *Cars 3*, Coco, *Incredibles 2*, and *Toy Story 4*. The winner of two Academy Awards, the most recent of which was presented at the end of February, Ted has made signal contributions to animation, and he's enriched the national dialogue on emerging technologies such as artificial intelligence. Ted, thank you so much for joining me today on Yale Talk.

**Ted:** Thank you for having me.

**Peter:** I've been wanting to interview you for Yale Talk for a long time.

**Ted:** Oh, great to hear.

**Peter:** So let's begin with your field-changing innovations. As one profile puts it, there's a decent chance that if a movie effect catches your eye, it started out as a series of dense equations that fill the whiteboard in Kim's office. Well, tell us a little bit about this technology and what you're writing equations about.

**Ted:** Sure. We most recently got an Oscar for soft-body physics. So anytime you see something squishy in a Pixar movie, that was probably us. So by squishy, I include stuff like skin or muscle or clothing is the big one, but actually even the passive things that are around you. So like when Mrs. Incredibles just sort of flops down on a pillow or something like that, the pillow was simulated using the laws of physics or in Turning Red when Mei is panicked and she starts pounding on the plush animals on her bed, even the fitted sheet and the mattress and all the pillows, these were all simulated using the laws of physics. And the reason the comedic moment is sold is because they move realistically.

**Peter:** Yeah, and so the result is animation that has a kind of realism that we've never seen before.

That's right.

And starting with the laws of physics and then turning that into algorithms?

**Ted:** Yes, actually. So it goes way back. One of the things that we developed was something called a stable neo-Hookean model. And Hooke actually refers to Robert Hooke. He's a physicist from the 17th century. He was even a contemporary of Newton. The things that were required from the physics back in the 17th century are not the same as what are required today. So we build on it and we discover new mathematical structures that have been hiding in the math the entire time. And of course we're doing it on computers, when back in the day, in Newton's time, they were just writing it out on pencil and paper. So we can do things in a different way that never would have occurred to them, of course, because they didn't have computers.

**Peter:** In addition to Hollywood filmmakers, others are using this technology in a way--medical researchers, aerospace engineering. Tell us a little bit about how this all works beyond animation.

**Ted:** Back in the 1990s or so, when physics was first coming into the movies, we were a bit maybe naively optimistic and we thought, well, people have been looking at this for a long time so we can just go to biomedical engineering, we can go into mechanical engineering, we can just sort of lift the models that they use there and transplant them into the movies. And this didn't quite turn out to be true. So in mechanical engineering, for example, what you care about is, is a bridge going to collapse? So there's what we call small deformation. So if the truss of a bridge is bent by 10%, this is a catastrophic amount of bending and it'll fail. So they don't really compute how something deforms by more than, say, like 10% or 20%. Whereas in computer animation, I like to use the example of Wile E. Coyote a lot. So a giant anvil falls on Wile E. Coyote and he smashes down into a zero-volume pancake. So we consider this like a 200% deformation...

**Peter:** And a violations of the laws of conservation of volume, right?

**Ted:** Absolutely. Absolutely. So then the question becomes, 1) how do we handle the large deformation, 2) what is a graceful way to break the laws of physics just a little bit so that we get the effect that we want? But then once the anvil falls off of Wile E. Coyote, he'll have to spring back. And what is a graceful way to handle that as well? So when we develop models for this people in mechanical engineering or biomedical engineering, they say, Oh, that's interesting. So it actually is robust to these ranges that we never really even looked at before. So they get interested again and the knowledge cycles back.

**Peter:** That's so interesting. Back to computer animation, were you a film buff as a kid and enjoyed cartoons? How did this all start for you?

**Ted:** Absolutely. I tell my students that I actually don't like math that much. I was an art student originally. I liked to draw comics and I liked to watch movies. It was only in high school when the first *Toy Story* movie came out where I realized, Oh, you can use math and computer science to have the computer draw for you. And then that's when I got interested in this version of art.

**Peter:** Computer graphics, I guess.

**Ted:** Absolutely. Yeah, that's right.

**Peter:** So Yale is one of the places where computer graphics is studied in quite a bit of detail. I think of people like Julie Dorsey and Holly Rushmeier, and you fit right into that group.

**Ted:** That was one of the main attractions for Yale. My first job out of undergrad, I was at a studio called Rhythm and Hues, and one of the first jobs that I had was to implement what's known as the Tumblin and Rushmeier tone mapping operator, which is named after Professor Holly Rushmeier, the professor here, and I talked about these elastic models for skin and flesh that we use extensively. And one of the main models that was used before the model that we developed was developed by Professor Dorsey here. So the opportunity to come here and work with the people whose work that I had been using for decades was an immense opportunity.

**Peter:** So in 2019, you make the transition from Hollywood and Pixar to higher education. That's a pretty big career move and maybe a move that some would consider less glamorous. So tell us a little bit about how you made that decision.

**Ted:** So I had actually been in academia for a long time, so I was only at Pixar for about four years. So I had my roots in academia. It was only because a lot of the core technology in computer graphics, just going way back even to the 1970s, 1980s, had come out of the Pixar computer graphics lab that I took up the opportunity to become a part of that legacy. But after four years or so in industry, usually people look out to the horizon, maybe a year or two, whereas in academia that's not really acceptable, right? You have to look out at least five years. That's sort of the minimum horizon. Otherwise you're not really doing it right. And I just wanted to look out a little bit further. So when the opportunity arose to come to Yale, I took it.

**Peter:** That's fantastic. What allows you to look out in longer time periods is you're less focused on the bottom line. So I do understand that. So you're the co-leader of the computer graphics group, and what we see as consumers of animated movies is a kind of realism that is so much more obvious than in early animations of, say, my childhood. But there are problems and challenges that the computer graphics group deals with. Maybe you could tell us a little bit more about some of those challenges.

**Ted:** So currently we are working on depicting all sorts of humans. One of the things that I discovered when I first came to Yale through a collaboration with Professor John MacKay in the film studies department, was that if you really look into the physics being used to depict virtual humans in movies, it actually only focuses on a very small subset of features. And if you look at which features have been selected to have their physics captured in the most precise way, it is white people. So you look at skin and what is the exact physics that's being captured and it is the glow of skin. If you look at a Johannes Vermeer painting or something like that Girl with a Pearl Earring, it is that glow that is being captured by the physics equations that are supposedly for all of human skin. But if you look at the cinematography of, say, Spike Lee movies instead, some of Spike Lee's cinematographers talk about the coatings that they put on people's faces so that it can convey shape in a very storytelling way. And there's never been a paper, from a technical perspective, ever written about how do you capture this kind of physics? So this is something that we think is a significant gap in the literature. We can think of the reasons, of course, that this occurred, but there's also just the raw technical challenge. You know, how do you capture something like this?

**Peter:** Does it mean, at least until now, that Caucasians are more realistically or more accurately depicted in animation than people from other backgrounds? Is that the end result of this?

**Ted:** So it's multidimensional. One component of this is, well, I turn on a video game and there are Black characters in video games, so obviously the algorithms don't fall down entirely. Yes, who looks the most realistic? It is the Caucasian characters. And also just from a representation perspective, who is the easiest to make if you're just starting out with these packages for how to generate virtual humans, the example is always a white person and the algorithms are geared towards generating a European person. So this directly influences who will then appear in the movies and in the video games because it comes down to money. What is easy to actually generate? It is what the algorithms actually support. And this extends not just to skin but to hair. So extremely straight hair is sort of a signifier of status. So the algorithms, they followed suit. Everything is designed for very, very straight hair. In fact, a bunch of the mathematical formulations, they start from assume that hair is an infinitely straight rod, whereas there are lots of types of hair, right? And we have never seen a paper dealing with How do you depict a human with extremely curly hair? Again, you can do it. The algorithms don't entirely fall down when you try to depict a human this way. But it takes a lot of skill because it's just not designed to do that.

**Peter:** Very interesting. That's just fascinating. So one of the things you've been quoted as saying is there are lots of stories out there and we haven't told a bunch of them, so let's go tell those stories. Is that your way of capturing that there's enormous human diversity, but only some of it shows up in animation?

**Ted:** Absolutely. I'm collaborating on a hair project with Professor A.M. Darke, she's at UC Santa Cruz, and she tells me these stories of what happens in family settings. So because straight hair signifies a certain type of status, there are wigs that Black people wear to work and then they take them off when they come home. So my colleague, Professor Darke said, if we develop these types of algorithms, we can actually have that scene, which I've never seen in an animated movie. The mom comes home, she takes off her wig and her afro-textured hair springs out from underneath of it. And I think that it would be such a subtle moment, but also such a real moment if something like that could happen.

**Peter:** Interesting. That's so interesting. Do you view this all as a technical challenge? Clearly, there's a sociopolitical aspect to this work, but is it part of your agenda?

**Ted:** So I've been working in film for about twenty years, and when I started out, everyone wants to just dig into the math. So it's almost like the math is safe. The math is abstracted away from the real world. But as I've worked on it more and more, I've just found that this is not true. And of course there's entire fields, history of science, that looks at this--the types of math that are developed, what are considered valid problems to look at and what are considered not valid problems to look at. These are highly social enterprises. So I want to shift things a little bit so that people can still have fun with the math, but they do need to think about where did the problems that you're looking at come from? Why are they considered valid and what are the solutions that you're coming up with? And why is this considered a solution while other things are not?

**Peter:** So last year, you and your team received $1 million gift from the Bungie Foundation. That was a gift to Yale as part of the For Humanity campaign. But it's supporting work that you're doing. Maybe you could tell us a little bit about that foundation and the kind of work you're doing with those funds.

**Ted:** So Bungie, the video game company, became famous for making the series Halo. But I think actually their big title right now is the Destiny series of games. I'd been talking about some of the limitations of these algorithms for depicting virtual humans for a while now, and Kareem Shuman, who is the head of their diversity employee resource group there, he reached out to me and said, How can we help? So he put me in contact with the Bungie Foundation and we put together a proposal for how do we depict more realistic hair. So the money right now is being used to fund some students to investigate what are the mechanics of highly curly hair. So highly helical hair. What is the interesting new physics that arises that have not been examined very closely? Because this problem formulation just wasn't something that was on people's radar before. And I could talk more about the results, but they're under peer review right now, so I will respect the scientific process.

**Peter:** Watch the space. There are articles coming out.

**Ted:** That's right. We sure hope so.

**Peter:** That is fantastic. So there are other projects that the computer graphics group is working on in addition to animation-based ones. Maybe tell us about 1 or 2 others?

**Ted:** Sure. Professor Rushmeier works a lot on material appearance. So all these textures and if you look at like the surface of a spaceship or something like that in Star Wars movies, it's like, Why does it look so realistic? And it's because of things that Professor Rushmeier does where they construct these incredibly realistic and infinite resolution textures in order to capture the appearance of basically anything that you would see every day. And Professor Dorsey looks at a galaxy of different topics. But one of the ones that I find fascinating is sketched-based interfaces. So the idea of how do people, when they're just drawing, how do they sketch and then how do they find ideas while they're sketching that then crystallize into the actual final drawing? And honestly, I don't know. I loved sketching when I was in high school, I sort of left it behind. But she is still forging ahead, trying to understand what do people actually do?

**Peter:** That's fantastic. There's a psychology side of that too that interests me.

**Ted:** Absolutely.

**Peter:** Yeah. So we are living in a time of great change. The listeners to this podcast, many of them probably have gone online and tried, for example, ChatGPT, a latest foray and artificial intelligence technology that is easily available to the general public, and people are now worrying about ethical issues and bias and that is a part of what you are interested in. It's implicit in everything we talked about today so far, but maybe you could tell us a little bit more about the implications of AI and the way in which I can be a source of bias.

**Ted:** So I should start by saying that I actually worked in some of this machine learning space for almost a decade before I went to Pixar. These were called subspace acceleration methods, also known as compressive sensing methods. This was the immediate precursor right before the big deep learning boom. So very early on, some of the researchers working on deep learning were talking to me a little bit because there was so much overlap in our spaces. So a lot of the stuff that I see with the limitations of these techniques, we saw these exact same limitations about a decade ago. And the big one, which I don't think that people really have a great answer for, even still, is all these machine-learning algorithms, they just ingest everything that, for example, they find in the internet, and then they interpolate between them, right? So it's just minor variations by combining them together. And the question was always, What about extrapolation? So what happens when you see something or you're asked for something that is very far outside of what you've ever seen before? And these things break. We saw it break. We actually made blooper reels of our physics simulations breaking when it was asked to do something that was very far outside of what it had seen before. So I worry about the current AI hype that is happening right now, because there seems to be a presumption that it will extrapolate extremely well. It will actually show you something that nobody would have thought of before, whereas my experience is it just is a combination of things that you have seen before. And from a creative perspective, I find this to be a little bit worrying, especially if people mistake what comes out of these AIs as something new instead of something old. So for example, there are these image-generation algorithms out there and people will say, Show me a picture of London at night, or something like that. And it'll show them a picture of London at night. And whatever they would have come up with if they hadn't been shown at that picture that is essentially destroyed, they will not generate that which came from inside of them. Instead, they will say, Oh, this external source said, this is what London looks like at night, which is just a restatement of what a bunch of other people thought London looked like at night. And I think that this is a shame. I think that there's lots of new, beautiful interpretations that people generate from their own personal experiences inside themselves and these new AIs, they can have the real possibility of just sort of smoothing away all of those interesting things that would have appeared otherwise.

**Peter:** I've played with ChatGPT in anticipation of a poetry group that I'm in.

**Ted:** Wow.

**Peter:** And getting it to write poems in different style, which it can do. They're just not great poems.

**Ted:** Right. Exactly.

**Peter:** They don't have a novel insight in them or a novel metaphor, but they fit the bill. They're in the right style, they have the right meter, they have a rhyming scheme that works, right? So the technical specifications are adhered to, but you don't see that spark of creativity. You don't have a stirring moment of recognition, or Aha!, or wonder, or awe when you read them. More synthesis than

**Ted:** Right.

**Peter:** than creation. And that could be helpful. It's a great tool, perhaps not for breaking new ground in the creative arts. You wrote a column in the Los Angeles Times recently and it had this great title on it. I don't know if it was your title or the LA Times put on there, but the title was "AI Flaws could Make your Next Car Racist." Maybe you can tell us a little bit more about the argument you were making.

**Ted:** So the LA Times did write the headline, but to be honest, I don't think it was that far from the original headline that I wrote. This actually goes back to the biases in the computer algorithms for humans that I was talking about earlier. One of the big topics in the computer graphics field right now is generating synthetic training data for machine-learning algorithms. So we know that getting real-world data, actually sending a crew out there to photograph, for example, all the roads for self-driving cars, this is extremely expensive. Only a few large companies in the world can even afford to do it. So what if you're a smaller company? What do you do? How do you collect this sort of data? And the argument here is, Well, why don't we use the same algorithms that we're used to make *Toy Story* and generate lots of synthetic data?

Because the presumption is that, well, these are the general equations that capture all the variation that exists in the world. And our studies have shown that this is not true. They are geared towards very specific features in the world which they capture extremely well. But this is not what you want from a self-driving car. You want it to actually understand all the variation in the world. So we see articles about this already, like someone in a wheelchair is crossing at a crosswalk. The cars are not that good at understanding these. This goes directly back to the algorithms. I was talking to colleagues at Pixar and if you want to show somebody who does not necessarily have full symmetric control of both sides of their body, this is more difficult to do in the software. The entire character creation pipeline assumes a lateral symmetry of the human body. If you want to introduce asymmetry, you can do it, but it's not what the algorithm was designed for.

**Peter:** So somebody with a disability is just not recognized in the same way. And is at greater risk, right?

**Ted:** Or someone has to be very intentional and go in and insert that data into the training set. Because if you just go with what the algorithms are good at, it's just not going to generate that.

**Peter:** Fantastically important. I always ask faculty who are part of Yale Talk, what they're teaching, and what students are coming to hear you say in the classroom. So maybe you could just tell us a little bit about your teaching schedule.

**Ted:** So right now I'm teaching a graduate course on physics-based simulation. I had the pleasure of announcing to my class that, Well, the material that we're going to learn today, I just won an Academy Award for it.

**Peter:** They must have gone crazy when they heard that.

**Ted:** There were some inklings already that maybe I was up for it. But yes, of course, everyone was very happy to be able to share.

**Peter:** We bask in reflected glory, and your students were doing that.

**Ted:** Yes, I hope so.

**Peter:** And do you know what you're teaching next?

**Ted:** I don't know quite yet. One thing I would like to teach is an undergraduate class, maybe first- or second-year, just on the kinds of art that you can generate from very small programs. So of course you can write very large programs like they have at Pixar to generate movies, but there are some very small programs you can use to make very beautiful things. So for example, fractals. So Mandelbrot, for example, I know was a professor at Yale for a long time, and I teach his algorithm in the advanced computer graphics class. But it actually turns out that even that sort of thing, these very beautiful, intricate structures, you can write it very succinctly. It's less than one hundred lines of code, and there's lots of different programs out there that you can do very similar things. And I would love to teach a class introducing students to these little beautiful snippets of code. And then, of course, the fun part is, is that when you modify it, new, beautiful things start to happen very quickly. And I'd love to see what they come up with.

**Peter:** Many students took Professor Mandelbrot's course on fractals as their quantitative reasoning, to fulfill that requirement, was quite popular actually.

**Ted:** Wow.

**Peter:** You would think it would sound a little esoteric, but students flocked to it. They loved it. So this is really next generation.

Ted, you've given us so much to think about today. Thank you for joining me on Yale Talk, and for sharing your insights with our listeners. Your contributions to your field, including those recognized in recent weeks by the Academy of Motion Picture Arts and Sciences, are a source of immense pride for Yale and, I think, great joy in your students. It's invigorating to consider the implications of your research and of your commitment to unleashing the stories worth telling.

To friends and members of the Yale community, thank you for joining me for Yale Talk. Until our next conversation, best wishes and take care. The theme music, Butterflies and Bees, is composed by Yale professor of music and director of university bands Thomas C. Duffy and is performed by the Yale Concert Band.