

Diane Tavenner:

Michael, you've just spent a week at the happiest place on Earth, and I must admit, I'm a little bit jealous.

Michael Horn:

For those who may be confused about what the happiest place on Earth is, it's Disney World. I was just there with my kids. First time for them. It was a blast. Diane, you know what? I came away with a few takeaways, but one of them was the excellence at scale. Disney has 74,000 employees in that park. And almost every single one of them - it's probably like 73,500 of them - are just dedicated to making your experience better than the last person you just interacted with. And it's astounding – however they have managed to do that. So it was a blast. Thank you for asking. But we're not here to talk about my vacation, although that might be fun. Instead, we're looking to continue to dive into some of these sticky questions around K-12 Education. Help people see different ways through what has often been pitted as zero-sum battles between the adults in the room and try to think through how we can unleash student progress and prepare them for the world into which they're entering. And obviously a question that exploded into both of our minds starting last year, Diane was the topic of AI. And, as opposed to the Metaverse, it is still the topic du jour. It is still what everyone is wondering about: artificial intelligence, what do we do with it, and so forth. And you have been teasing me that you have the perfect guest to help us think about this in some novel ways. Take it from here, Diane.

Diane Tavenner:

Well, I have indeed been doing that. You're right. AI so far has a longer shelf life. So we'll see how long that lasts. It's my great pleasure to introduce you to Irhum. And Irhum and I first met a few years ago when he was a freshman at Minerva University. He was coming from Bangladesh to that global university. He's now a senior. He spent the last two summers as an intern at Google X here, just about a mile away from where I live. And at Google X, he's really been focused on large language model, aka AI, research. And you've been hearing about Irhum from me and all of our conversations we've been having for quite some time, Michael. So, what you know is that I've learned a ton from him about AI. And one of the things I love about talking AI with Irhum is that even though he has a ton of knowledge, and, for example, he writes a popular technical blog about AI that I have looked at, and I can't even decipher a sentence of it. So, highly technical, deep knowledge. But he also is a system thinker, and he cares deeply about how technology is used, how AI is used, and what it means for our society. And so he's willing to and able to talk with people like me, lay people like me, and help me understand that and engage in a good conversation. And for our purposes, I think, most importantly, Irhum is 20, and it's so critical to be in dialogue with people in this generation. I think we give a lot of lip service in education to the consumers, if you will, or the students, and then we don't involve them in our dialogue. And so, I'm just really grateful that he's here and you all get to meet. And so, welcome, Irhum.

Michael Horn:

Irhum, it is so good to have you here. Diane has been teasing this for a while, so thank you for joining us. Before we dive into the AI topic itself, I would just love to hear, through your words - because I've heard it a little bit through Diane's - but I'd love to hear and the audience would love to hear about your journey to Minerva University, your journey to diving into topics of AI. And really, how has that school experience specifically been? Like, what has worked? What hasn't?

Irhum Shafkat:

Yeah. Thank you so much for having me. Let's just get into it. I'm one of those people who the standard school system was just not designed for, rather frankly. I grew up inside the national school system back in Bangladesh, up to fourth grade, essentially, and it just wasn't designed with someone like me in mind. We're talking an institution with 50 person classrooms, teachers barely able to give anyone attention. And the school system is geared to make you pass a university entrance exam, and if you do that, you've done it, you've succeeded. And my mindset often was, like, I joked that I was learning full-time and going to school on the side. The way I saw it in my mind, that's amazing. And I'd say the only reason my path worked out the way it did is because when I was in finishing up middle school, entering high school, so, like, eight to nine grades, like that window, the Internet just rapidly proliferated across the entire country within a couple months. Short couple of months, essentially. You went from not a lot of people using the net to a lot of people using the Internet. And I was one of those people, and I was like, "Oh my God, it's not that I don't like math, it's that I don't understand it." And there's a difference between those two things. And I was one of those people. Like, Khan Academy was quite literally designed for me. I'd log onto that thing and I was like, "Oh my God, I actually understand math," and I can teach myself math. I can succeed at it. And the impression that it left me with is that technology is really opened up. The Internet, in particular, is like one of those frontier technologies that just opened up learning to anybody who could access it and go through the right set of tools needed to learn things, like Khan Academy being one of them. And I guess that's the mindset I'm seeing this new generation of technologies in AI, too. I wonder who else is going to be using them the way I did and learn something new that their environment wouldn't otherwise allow them to. I suppose I kept teaching myself things, and that's kind of, I guess, partly how I ended up at Minerva is because I wanted a non-traditional university education because the traditional high school education clearly wasn't a fit for me at all. And Minerva was like, "Come here, we won't bore you with lectures. Our professors barely get to speak for more than five minutes in class. It's all students just talking to each other and learning from each other. And I was like, "Sign me up."

Diane Tavenner:

Well, and Irham, that curiosity that you are describing in yourself is, I think, one of the things that led you to discover AI long before most of us discovered it. And you discovered it on the

Internet. You discovered it by reading papers and sort of following these blog posts while you were teaching yourself. And so, I don't think it was a surprise for you when it burst onto the scene, because you knew what was coming and where it was coming from. And yet, I think that the world's reaction has been interesting. And so, you sent me a quote the other day. You texted me a quote about AI that had us both laughing pretty hysterically, probably because it feels very true. And so, I'm going to share that quote, which is, "Artificial intelligence is like teenage sex. Everyone talks about it. Nobody really knows how to do it. Everyone thinks everyone else is doing it, so everyone claims they are doing it." Let's just start by figuring out what people are actually doing. If anything, this is sort of a ridiculous question, but can you just explain AI to us and why it's suddenly this big deal that it feels like it just spontaneously arrived in 2023?

Irhum Shafkat:

I guess there are like three big terms that we should go over when we say AI because it just means so many things. Like AI is such a vague, hard thing to define. But I guess the way I see it is anything that's just right beyond the edge of what's computationally possible right now. Because once it becomes possible, people kind of stop thinking about it as AI.

Diane Tavenner:

Interesting.

Irhum Shafkat:

Okay, other people would have their own definitions, but I feel like that's like, what's historically been true. Once something becomes possible, it almost feels like it's no longer AI. What has been more recent, though, is how we keep pushing different tier. Like, if it was the 90s, for example, when Casperov played with Deep blue, the chess playing engine. That thing was like a massive, highly sophisticated program with millions of logical rules. That's what's different. What happened over the last ten years is we really pivoted towards machine learning, and specifically a branch of it called deep learning, which allowed us to use really simple algorithms on enormous amounts of data. We're talking several thousand lifetimes worth of data. When you were talking about training a language model, for example, and ginormous amounts of compute to push all that data through a very simple algorithm. And that's what changed in that it turns out really simple algorithms work extremely well when you have a large enough compute budget and enough data to pass through them. And the instantiation that everyone really captured the imagination is against a narrower part of machine learning still is called large language models. The large being like, they're much larger than models for historical language models, in the sense that, at their core, you give them a piece of text, and they just play a guessing game. We're like, okay, so what's the next word? And if you keep predicting the next word over and over, you form entire sentences, paragraphs, entire documents that way.

Diane Tavenner:

So it sounds like, basically, I'm going to bring it back to education really quickly, Michael. Like AI as you're describing it, Irhum, is this concept we have in education and learning, which is I plus one. So it's like where you are just a little bit harder than what you can do, and that's like your zone of proximal development, like where you best learn. So, it sounds like AI is in that space. It's like what we can do plus a little bit more is what keeps pushing us forward, basically, on everything that's written on the Internet with simple formulas.

Irhum Shafkat:

Yeah, I guess what captured the imagination is, I think, language in particular, because ChatGPT came out sometime last December, and that, again, took everybody by surprise, essentially. But the thing is, image generation models like Dali Two and stable diffusion were out earlier that summer, and they're built on basically the same technology, given a ginormous amount of data, produced more samples that look like something that came from the data distribution. I had this with friends, especially outside of computer science, where I show them these generated images, and they're like, yeah, I thought computers could already do that. What's special about them? And I'm like, oh, my God, it's like generating a panda in a spacesuit on Mars.

Diane Tavenner:

And they're like, yeah.

Irhum Shafkat:

I thought Photoshop already does that.

Michael Horn:

So interesting.

Irhum Shafkat:

Whereas with language, language feels...I would almost say people see a conversational interface and almost by de facto assign intelligent attributes to it. Now, this is not to say that it's all smoke and mirrors. Like, these are remarkably practical technologies that I genuinely think are going to change a lot of what we do today, as we'll explain. But at their core, I think there's at least part of the fact that humans genuinely see language differently than other modalities, because, in part, it is unique to us - that we know of at least.

Michael Horn:

It's really interesting, though, because the implication of that, it seems, is that part of this isn't just the power of what's been built, but also our reaction to what's been built. And I think then the corollary, something Diane and I have talked a lot about is, like, right now, you're seeing very polar opposite reactions to AI. Either it's the utopian thing that's going to bring about this glorious future where resources are not scarce and everyone will be fine etc. Or it's very dystopian, and you see a lot of the - I'm going to insert my own belief for a moment here - but a lot of the technology leaders that have developed this being like, it could be dystopian, it could take over the world. As I hear you describing it, it doesn't quite sound like either of those is right. It's like the next step. So, I'd love you to put this in a human context, then, of what does it mean about the roles, particularly as we think about the future of work that humans will play with AI, or how human AI maybe is itself or isn't? How do you think about the intersection with humanity?

Irhum Shafkat:

When I think about the future, one tool I borrow from a former colleague, Nick Foster, who was, I think, head of design at Google X. He would introduce this concept of thinking about the future not as a thing you're approaching, but almost like a cone of many possible outcomes. There is the probable, which is like the current set of outcomes that seem like really probable, but the probable is not the possible. Like, the possible set of features is much larger. And right now, I think we run a real risk of kind of just not seeing the technology for what it is. It's a tool that we could actively shape into a future we want. And instead of trying to imagine it through that lens, we're almost like we're kind of giving up and we're like, yeah, it's going to take over the world or something. It's going to be bad. Oh, well, it's almost like we're sleepwalking towards an outcome to a degree. I see it genuinely as, a technology, a tool. And, yeah, it's up to us to decide what integration of that into our society looks like, but it's a tool.

Michael Horn:

Do you think part of that is because some of the people behind the coding are also surprised by the outcomes that it produces? And, like, gee, I didn't know it could do that. And so they're sort of showing, to your point about this weird passivity that we all seem to be displaying, that maybe it's because they, too seem surprised by some of its capabilities. And so that is surprising the laypeople like me and Diane, and not to mention the people who aren't even playing with it yet.

Irhum Shafkat:

So, I think it's important what we've seen over the last ten years, but really, the last four, I would say, with scaling really taking off, is when models get larger, they have more capabilities. Like a model from three years ago. No, I'd say a little longer. Maybe four or five, if you go to the joke and ask it to explain it, it wouldn't be able to do that. It would struggle. But if you ask ChatGPT now, like, "Hey, here's a joke that my teenage son wrote, I don't get it. Can you explain it to me?" It's going to do a pretty decent job at it. So, one of the things that came with scale is

like, capabilities emerge. But the thing that surprises researchers is, I think I saw this analogy on a paper by Sam Bauman, but it's almost like buying a mystery box. I think you buy a larger box, there's going to be some kind of new capabilities in there, but we don't know what they're going to be until we open up that box. So, I guess that's where a lot of surprise, even from researchers, come from. That said, I will push back on that a little in that there has been really genuine and serious work done in the last two years where we're really trying to figure out. We're like throwing the whole internet at this thing. There's actually a lot of things going on in there and that these capabilities are not actually as surprising as we think they are. Like joke explanation. At the start when it came out, we were like, "Oh my God, this thing can explain jokes." But then when you dug into the data deep enough, you're like, there are tons of websites on the Internet that exist for you to dumb down a joke and explain how it works. So, it didn't appear out of completely nowhere. The fact that it works on new jokes that are hopefully original and it's still capable of explaining them is still cool, but it didn't come out of absolutely nowhere.

Diane Tavenner:

Irhum, you just started giving us some time frames and timelines, and you're sort of, in your mind calculating ten years, four years, two years. I just want to note that ten years ago, you were ten years old, but, okay, we'll set that aside for a moment. But as you were using those timelines, one of the things that comes up a lot for people is that this feels like it's going so fast. If you didn't even understand what was happening in this world, and then suddenly chat GPT came on the scene. You probably didn't look at it in December because you were busy. But then in the new year, suddenly the whole world's talking about ChatGPT, and you log on, and then it just seems like every day something new is coming and faster. And I talk to people who just are like, "I can't keep up." It's only been a couple of months, and it feels like it's just spinning so fast and beyond our control. Is that true? How do you think about that timing, and how do you think about keeping up? How can we conceptualize that, especially as educators?

Irhum Shafkat:

Well, for what it's worth, even researchers have trouble keeping up these days with the sheer amount of papers coming out left, right and center. It's hard. That said, and I'll say it is genuinely surprising, the pace at which GPT, in particular, took off, because these models had existed, the model they're building off of GBD Three that had existed since 2020. We're talking around the pandemic start period. They just put a chat interface on top of it, and that really seemed to take off. I remember reading news articles and even OpenAI seemed, like, surprised that just putting on a chat interface on top of a technology that had lying around for three years caused it to really take off that way. And even I was surprised because, again, I was in Taipei in the spring, and Taiwan in the spring, and I remember being on the Taiwan high-speed rail, and I'm seeing someone else use ChatGPT on the train next to me. And I was like, wow, this thing is taking up a lot faster than I realized. But it's important to understand that when I say again that ChatGPT 3 had been lying around for three years before they put a chat interface on top of it. And again, this should not underscore the fact that these models are going to get larger,

probably more capable, but it should also ground you in two things. One, the specific burst of innovation we saw in this year in particular had been building up for a bit. Essentially, it's almost like a pressure valve went off when they put a chat interface on top of it. The other thing, and this is the thing that I wish people would discuss more often, is that it's not just that the models got larger and we trained them on more of the internet. It's also that we start paying a lot of money to get a lot of humans to label a lot of this data so that you could fine tune what the behavior of these models are. You see, when GPT3 was trained in about 2020 or so, it's what we call a base model. It does exactly one thing. You give it a piece of text, and it produces a completion similar to the text data it saw when it was being trained, which would be raw internet data. And it has a tendency to go off the rails because the Internet is full of people who say not very nice things to each other. What changed was the sheer amount of human data collection that went in during that time frame and that this large model was adjusted over to tame its behavior, teach its skills, such as what explaining a joke even is, and all those things. We could talk about that a little bit more. But the big connection being that the jump we saw in those three years is that of a technology that had already existed, that we really learned to adjust better. But we already burned through that innovation once. It's unlikely we're going to see another leap on the same scale of learning how to use supervised human fine tuning again, because that innovation has now already existed and is already baked into it.

Michael Horn:

That's fascinating. It's something I hadn't understood before either, which is to say that in essence, if I understand you right, obviously the code base continues to evolve for GPT4 and so forth. But in effect, I think what you're saying is the user interface and how we interact with it is what actually changed, like the skin of how humans interact with the code base. I think it leads us to where we love to talk on this podcast, which is the uses in education and how it's going to impact that. And obviously, I'll give you the hall pass if you will. You're not an educator. We're not asking you to opine in some way that puts you in a position you're not. But you are a student right now at a cutting-edge university, constantly thinking about pedagogy. And so, I'm curious, from the student perspective, what excites you at the moment about AI in education?

Irhum Shafkat:

I think shrinking the learning feedback loop is the way I would put it. I'm a systems thinker and I use the lens of feedback loop a lot. And whenever you shrink a feedback loop from, say, learning something, like maybe getting a feedback. What I say by feedback loop in education is like, you write a paper, your professor takes maybe two weeks to get it back to you, and you get almost like a signal, like, "Hey, you got these things right. You don't get these things quite right, though." What happens when you shrink that from two weeks to a couple of minutes or maybe a couple of seconds? It's not just shrinking a number, it's changing how you interact, how your learning experience evolves. And I think it's nice to connect this back to my middle school years. I think the reason I ended up in math and programming in particular is because those two things have really short feedback loops. In programming in particular, if you write

bad code, your compiler just screams that. It's like, "Hey, you're trying to add a number to a set of words, it's not going to work." And you get like really short, tight feedback loops to keep trying your code over and over again until it succeeds. That's not true with learning English. With learning English, you write a paper, you wait a week for your professor or teacher to get it back to you. You maybe pick up something and try again next time. And I would say at least part of the reason I ended up picking math and programming is, again, I didn't have that many great resources in terms of teaching, teachers who could really help me out. So, I naturally gravitated towards things that I would be able to really quickly iterate on: math and programming. Whereas those things - English, the sciences even - I would argue broadly, are not on those same lines. But what then changes with AI, is like, you now have a chatbot, that it doesn't have to be a chatbot. It could be far more than that, really. You just have a computational tool that can actively critique your writing as you're writing it out. You're like, "Hey, what are some ways I could have done this better?" They're like, "Yeah, you're using these passive wordy phrases. Maybe you shouldn't be doing that." And you're like, "Why shouldn't I be doing that? Because it makes it harder for other people to read." And instead of waiting a week for that to happen, you get fed that feedback in real time, and you have another iteration ready, and you then ask it again, "Hey, how could I do it even better?" That loop shrinks significantly.

Diane Tavenner:

That's fascinating. And I think it's also what we, Michael and I, have been working on personalized learning, or whatever you want to call it, for a decade plus at this point. And that is certainly one of the promises of personalized learning, is that tight feedback loop. So you're staying in the learning, and it isn't delayed. It's contextualized, and it's in the moment and immediate. It's more tutor if you will. That's why many people have put so much energy into that. And you have now said a couple of things. One ChatGPT, putting this skin on their product, which is essentially a chat bot. Like, you talk with this bot or this window. And then also this potential of what you just described, like the feedback coming a chat bot, if you will. And in my experience, most people, when they think about the uses of AI, are thinking along these lines. Like, that is what they think is there's going to be someone, whether it's a little avatar or a box or something, that's like sort of chatting with me, and that's kind of how AI is going to play out. I know that you sort of get a little exasperated when that's all people can imagine. What else might be possible help us expand our thinking a little bit beyond just this sort of chatting?

Irhum Shafkat:

I think right now we're at the phase when the iPhone would have been sometime in 2007. I'm not really qualified to comment on that because I would have been what?

Michael Horn:

Don't worry about it.



Irhum Shafkat:

But at least from my understanding of that time period, apps were a new thing. People didn't really know how to fully utilize them. And the first set of apps were kind of like window gag. And I was like, people were building those? Like, a flashlight app where instead of turning on your phone's flashlight, it just showed a flashlight on the screen.

Michael Horn:

I downloaded one of those. So, it's true.

Irhum Shafkat:

So, I feel like we're in that era for these language processing technologies right now, in that we have a brand new tool. We're not entirely sure what using that looks like. And returning back to the quote Diane quoted with the enterprise, I feel like not nearly as many people are asking, "How does this help us solve our problems better?" And a lot more people are asking, like, "How do I put this into my product so my board of directors is happy?"

Diane Tavenner:

Yeah. So, the idea being like, what legitimate problems do I have, and can I use this to solve them? And that's where where a useful app - not a flashlight - is going to come from.

Irhum Shafkat:

With chatbots in particular. Chatbots became the first big use case. So, everyone's like, well, this seems to work. Let's make my own chat bot, but, like, slightly different. And I think it's just so unimaginative. But the other thing with chat bots is that they have really bad what we call affordances in design. And affordance is something that almost cues you into how to use something. Like, when you see a handle on a door, you're like, "Oh, this is the thing I grab." Then, let me ask you, if you had to choose between a big cancel button that cancels your flight that you don't want to go on, versus dealing with a chatbot to cancel your flight, which one would you pick?

Diane Tavenner:

Yes. And you have showed me this and demonstrated this to me. A single button like that is so much more useful. If you know the thing that I'm going to do versus making me talk to it and explain it where something's inevitably going to go wrong.

Irhum Shafkat:

These technologies are so nascent right now, and what people really need to appreciate is that you need to design them in ways where you almost expect they're going to produce some unrelated, unpleasant output somewhere along the process. So, it's almost like a liability if you're giving a user an open box to type anything they possibly want. And it's not just a liability; it also just makes the user, again, it doesn't give the user any affordances. They see a blank box. Like, again, if you need an English tutor bot, you could just as easily have a couple of buttons that could summarize, remove jargon, sorry, highlight jargon, or introduce a couple of buttons that cover a couple of use cases that a standard ninth through twelfth grader may want to use so that they can refine their writing better. Don't give them an open-ended box where they don't even know what to ask for help in.

Diane Tavenner:

And would that still be using AI, those buttons?

Irhum Shafkat:

I mean, behind the scenes? At the end of the day, any implementation of these things that you've looked at is a role play engine. If you have interacted with one of these airline processing, like if they use a language, they're using large language models. Behind the scenes, there's only two bits that anyone else modifies. Once you have the model, which is you have a dialogue preamble, which is almost like a script someone else writes as like you are now about to role play a dialogue, you are about to assist a user with cancel with their airline queries. And then the actual dialogue that happens, all people do when they're creating a new implementation of them is that they switch out the preamble with something else. I'm just saying you can write a preamble for each of those buttons. Essentially, if the button is like, hey, you highlight jargon, then you're like, okay, you are about to role play a bot that takes in some English text and returns the exact spans of text that are abnormally jargony for the writing of a nine to twelve grader. You as the developer should be writing that script because you are the one who knows what this person needs to be using. You shouldn't leave it up to the user to be writing their own scripts, essentially. For the most part, unless it's like a very specific use case where you wouldn't want the user to be writing that, but don't treat it as the default one, which we for some reason do.

Michael Horn:

So, it's super interesting, the implications on education. And I have a couple of takeaways here and I want to test them by you and then get your reaction. So, the first one is on the flashlight app analogy. I think the implication is that if the iPhone, or in this case OpenAI, is ultimately going to very easily incorporate the thing that you've thought of in their own roadmap, your sort of idea is not going to last very long on its own, right? The reason there were these apps is like there was not a button to do your own flashlight or change colors or things like that. And so, they built these apps and then very quickly iPhone realizes, hey, we can just add a quick little button. But I think the second implication is as you start to think about the user

experience and come up with these prompts that a user might want to go through so that they're not guessing what they're putting into that open ended box, that the opportunity for educators might be to start to build around these different use cases, that if I'm understanding you correctly, perhaps OpenAI is not going to develop all these different use cases and so forth. But instead, the opportunity might be for educators to build on top of their code to develop these sorts of things that help get the problems solved that they and their students are actually trying to solve. But I'm curious where that line ends, and if you see it differently.

Irhum Shafkat:

It's not a clear hard line. What's on open AI's roadmap? I don't know, but I guess it ties back into what your company is. If your company is trying to add value to people's lives, that really means sitting down is like, "Oh, am I just going to write a wrapper that takes in a PDF and allows you to just chat with it?" I mean, those all kind of went the way of the dinosaurs last week when OpenAI just integrated like, hey, you can drop a PDF into our thing. But if you really try to aim for deep integration, OpenAI has. I'm making claims here, I don't know for sure, but I would imagine they don't have a deep understanding of the K-12 education system, nor would they be necessarily interested in that, because that doesn't seem to be what their goal right now is, which is to make bigger and better version of these models that are more generally capable. What they're counting on is that other people take these models and integrate them throughout the rest of the economy. And that's where other people come in, and that's where K-12 educators come in, because they are the ones who have a better understanding of what these buttons need to be, what they want their students to be learning at the end of the.

Michael Horn:

Love it.

Diane Tavenner:

What's coming to me, Michael, is, on our last episode, we talked with Todd Rose, who wrote the *End of Average*. That is, in my mind, the foundation of a lot of personalized learning. And one of, I think, the misconceptions people have, or the swings. Education loves to swing the pendulum as we go from totally teacher directed, controlling every second of every bit of learning, and we swing all the way over to like, basically go teach yourself. We're just going to throw you out in the wild and at some level just throwing people into a chat bot of a large language model is throwing them into the wild. And so, what I hear you saying, Irhum, is we need the in-between. There's a real role for educators to narrow. It's not like the whole world personalization is. There's 12345 ways of doing something and we can narrow down to that and create a much more personalized experience that is curated by expert educators. And we should be looking for that. Happy in between.

Irhum Shafkat:

I mean, again, it's been so fascinating seeing these chat bots interact with the education system and kind of wreak havoc. Honestly, to a degree where professors taking stances on everything from like we should go back to making everyone write things by hand, by the way, please don't do that. They are an opportunity. Here's the thing: the factory model of the education system we have where people just go through it, do these problems, hopefully learn something by doing said problems. We don't need to know they're learning the thing. We just need to know if they're doing the thing. They have like a high school diploma. That was going to come to an end because that's already been out. Like that's already ill preparing people for the jobs of today. For a while now, these models, they're not bringing about something that wasn't going to happen. They've just sped it up, essentially, because students already again ask why a student would actually go out of their way to get an essay that I created by these things if they genuinely believed in their own education that hey, if I do this thing I'll learn something new and that will be helpful in the future. They probably wouldn't. So why do they feel disillusioned? Because they know deep inside that what they're learning in their high school is not going to actually prepare them for the world. And you need to actually deal with that disillusionment. And on the counter, I think these models provide an excellent way to actually start tackling that disillusionment by educators seeing themselves almost as designers, as what people need to be learning. I use the example of the door handle because it seems like a simple object. It really isn't. If you've ever been in a hotel with one of those weird, poorly designed shower knobs, you know how much bad design can mess up your day. And when good design works it's almost invisible. Like you don't even notice a shower knob when it actually works. And I think that's what good educational software using these will almost look like the students won't even realize how seamless it feels, like they press a button that tells them, hey, this is the jargon you're using. Here's why it's bad for you. Here's an explanation how you could do it better this time. It should feel seamless and they should feel less disillusioned because they feel like, "Oh my God, I'm actually learning something."

Michael Horn:

Here, so I want to stay just as we wrap up here. And one last question before we go to our sort of bonus round, if you will, of stuff outside of education, which is you just painted a good picture, I think, of how the education system has reacted in very nervous, let's call it ways to this advent of it, because it has immediately sort of thrown into question so many of these tired practices that it holds on to. And I guess the corollary question I'm curious about, I've heard a lot of students, let me frame this a little bit more. I've heard a lot of students say, Professor X, you're thinking a lot about what is the assignment and how am I going to catch you from cheating, but you're spending a lot less time thinking about what do I need to learn to be prepared for this world in which AI is going to be underpinning basically everything I could possibly go do in a career? And so, I guess the question I'm curious, from your perspective is, as you look at these traditional factory model education systems, what's something that they should start teaching students that they don't perhaps today? And what's something maybe that they should lose that they continue to hold on to?

Irhum Shafkat:

I mean, honestly, I don't want to sound like a shill for Minerva, but I am going to. I think freshman year, I had never written a full-length essay prior to freshman year in English, and I was kind of really lost, honestly. But one thing that really stood out is, like, my professor spent so much time just breaking down the act of writing into what does it mean like to have a thesis? What does it mean for a thesis to be arguable and substantive instead of something everybody universally agrees with? Because if everybody agrees with what you're writing, you don't need to write it. Really breaking down the act of writing into these atomic skills that I keep finding myself using even at the tail end of college now, in senior year. I think that is the kind of thing we're going to need to do, is like actually asking ourselves like, this is an instrument, a tool we've built, that we administer to our students in the hopes that they learn something. Does this tool actually do the thing it advertises? It does but a lot of the time it just doesn't. And we just kind of need to be honest about that because, again, it's a lot like, again, the ChatGPT moment in some sense. But also for education, it's been building up like a pressure valve, and that pressure valve kind of just went off in the last year. Wow.

Diane Tavenner:

Well, we could talk for a long time, but that might be the place to land it today. But before we let you go, we always like to, at the end, just mind for what we're reading, watching, listening to outside of our day to day work.

Michael Horn:

Do you have any time for that as a student?

Irhum Shafkat:

Can I talk about a video game?

Michael Horn:

Yeah, that's great.

Irhum Shafkat:

I've been playing a lot of Super Mario Wonder, which is like the new Mario Brothers game from Nintendo. It is fun. It is really the best way to describe it. A lot of media really enjoys being dark and gritty and mature or whatever. Nintendo is like, we're going to make a game that's unashamedly fun and bright and colorful and just playful. And they've been doing that for the better part of, I don't even know, like 30, 40 years now. And they've kind of just stuck to it as a core principle. And I kind of just admire their ability to really set a mission for themselves, which is make things that make people find joyful and fun and actually just stick to it for the better part of half century.

Michael Horn:

Oh, I love that one. Diane, what about you?

Diane Tavenner:

Well, we're going to look kind of boring following that one. I'm going to go to the dark, gritty world. I just finished reading the book *How Civil Wars Start and How to Stop Them* by Barbara F. Walter. I will just say seven of the eight chapters do an excellent job of diagnosing the problem, and it's pretty terrifying. And I do think we should know it. Chapter Eight, where the solutions come in, was not compelling to me, and so it feels like there's work for us to do.

Michael Horn:

And I guess I will say I'm in a similarly dark place, maybe, Diane, because I read the *Art of War* by Sun Tzu. So go figure. Everyone can figure out where our headspace is. But I finished it before Disney. I had tried to read it a couple times before, and this time I made it through, which is setting me up now for reading Klossovitz, which is where I am now. Grinding through is the right verb, I think. But on that note, Irhum, thank you so much for joining us and making a fraught topic - but a topic with a lot of hyperventilation - really accessible and exciting and giving us a window into where this could be going. Really appreciate you being here with us on class disrupted. And for all of those listening, thank you so much for joining us. We'll see you next time.