Artificial Intelligence: The Role of Psychology and How to Navigate Change
Date & Time: Dec 6, 2023 02:00 PM Eastern Time (US and Canada)

Karen Stamm: Hello and welcome to today's webinar, *Artificial Intelligence and the Role of Psychology and How to Navigate Change*. I'm Karen Stamm, Director of the Center for Workforce Studies at the American Psychological Association, and I'll be serving as your moderator today. This webinar is part of a series on generative AI, and we thought for today we would broaden our focus to look at artificial intelligence more generally, as well as psychology can help people navigate the changes that we anticipate will be happening in the future.

Broadly speaking, artificial intelligence is the simulation of human intelligence by computers, machines, and algorithms. It encompasses many different types of tools and models, and it generated rapid interest beginning just one year ago with the public release of generative AI tools such as ChatGPT. Before we get started, here are some important logistics points.

This program does not offer CE, however, we will email everyone watching live today a certificate of attendance for anyone who watches a minimum of 45 minutes. During today's webinar, please use the Q&A feature located in your webinar screen. Have any questions for our presenter? Please type them in using the Q&A at any time. I will be monitoring the Q&A throughout the webinar to identify questions to ask live as time permits.

We are recording this webinar, and we will make the recording available to everybody within two weeks' time. A link to the presentation slides will also be posted in the chat box of your webinar screen. If you missed the slides or the link, don't worry, we'll send out both the slides and the recording with that email within two weeks' time. We're going to start off with a presentation followed by a moderated discussion and plenty of time for audience questions at the end.

Before we get started with the presentation, we wanted to launch two quick poll questions about technology in the workplace. Give us just a moment to launch our first question, and we'd really appreciate it if you would take a moment to respond to these questions from your own perspective to share how much you agree or disagree with the following statements.

[pause 00:02:43]

Karen: We'll leave this poll open for another 10 or so seconds, so please make sure you answer.

[pause 00:02:50]

Karen: All right. Looks like the number of responses are slowing down, so why don't we go ahead and end the poll and show the results? Interesting. A fair amount of agreement with that statement. About 41% look like you somewhat agree that you worry that artificial intelligence may make some or all of your job duties obsolete, and another 5% strongly agree with that statement. Let's go ahead and launch the

second poll question. [silence] Take a moment if you'd like to, to respond to this question.

[pause 00:04:03]

Karen: Let's give it maybe another 10 seconds or so. Looks like the number of responses are still coming in pretty steadily, or starting to slow down a bit.

[pause 00:04:31]

Karen: All right. Looks like we're good to go. We can go ahead and close the pole and share the results. Interesting. A very high level of agreement. 43% somewhat agree that advancements in technology are helping you to work more efficiently and accurately. 45% strongly agree with this statement. Okay, well, thank you all for responding to those poll questions. Those were actually two questions that were part of APA's 2023 Work in America survey, and we'll be hearing more about the results of these survey questions later on in the webinar. We can go ahead and stop sharing that, please, and move on with the rest of the webinar. Let me, there we go.

Now let me introduce our speaker for today. Dr. Mindy Shoss is a professor of industrial and organizational psychology at the University of Central Florida and a fellow of the Society for Industrial and Organizational Psychology. Her research examines the intersection of the future of work and occupational health psychology, including such topics as job insecurity, artificial intelligence, and worker adaptation and well-being. She currently serves as the associate editor of the Journal of Occupational Health Psychology. Welcome, Dr. Shoss. We're delighted to have you here today. I'll turn it over to you now and give you control so you can advance the slides.

Dr. Mindy Shoss: Thank you. Thank you to APA for inviting me to be part of this webinar, and thank you all for joining this webinar on artificial intelligence, which has been described as the new electricity and as the future of work and medicine, education, government, and even as the future of humanity from leaders in this industry. The goal today is to invite a discussion about the role of psychology in the development and implementation of individual AI systems, to talk about some of the applications of psychological science to issues related to advancement in technology, and to consider how psychology can help employees and employers navigate change in the workplace.

The AI expert Azeem Azhar's book, *The Exponential Age*, convincingly argues that we are hurtling toward an unknown future at an exponential rate. There's a wide range of estimates about the rate and the path of AI development. For example, Open Philanthropy estimates what they state to be a non-trivial probability that AI will transform society in the next 20 years, much in the same way that we saw transformations from the Industrial Revolution.

Although timelines are difficult to estimate, most experts see changes coming, and they see these changes coming quickly. What makes this so challenging is that there's so much uncertainty about what these changes will look like, and there's

very, very different ideas about whether these changes will improve or will harm people's lives. This is where I think that psychology comes in.

Psychologists have certainly made the case for psychology as important to the development and the deployment of AI systems. However, I think some of the strongest arguments for the role of psychology come from discussions that are going on outside of our field. Judy Estrin, who is former CTO of Cisco, argued in an op-ed that discussions around AI often seem to treat this technology as inevitable, almost as if the technology itself is going to guide its own path forward, or we're all just going to be along for the ride of whatever technology companies create.

She encouraged leaders and researchers and software developers to utilize human well-being and dignity as a benchmark against which these systems are designed and evaluated. As I'll come back to several times, thinking about issues of human well-being and dignity is something that is part and parcel to the field of psychology, and we have a large disciplinary background in ways to think about, to foster, and to assess these outcomes.

Another quote here from Eric Lamarre from McKinsey echoes this point. The conversations right now make it seem like AI is a technology in search of a problem. This to me, again, suggests that psychology has an opportunity to help guide some of the thinking about the types of problems that AI should focus on in the processes through which AI projects are prioritized, developed, and evaluated.

Here's a broad view of the AI life cycle from the types of AI that are designed, how they're developed, how they're deployed, both in society and in specific arenas such as workplaces, and how these systems are continually evaluated and updated. In the center are values and objectives connected to human dignity and well-being broadly construed. Because psychology has a rich knowledge and understanding around issues of human dignity and well-being, there's a role for psychology in all of these aspects of the AI life cycle, and there's already a lot starting to be done in these areas.

To offer some more specifics, questions about AI design that are human-inspired might center around the types of AI that might benefit humans in society. In other words, what are important problems that would benefit from a tool like AI being applied? Here we're seeing some important work. For example, there's research and applications that are thinking about AI and early identification of workplace hazards. There's research on how AI can create scalable personnel selection methods that can be used to make the tools of biopsychology accessible to a broader range of organizations.

There's also interest in design issues around AI, which include not only the type of problem that AI is trying to address but how the user interface functions and the extent to which there's value in adding more or less human-like elements. Of course, these questions are intricately connected to issues of development and human dignity because models that are trained on biased or unrepresentative data are not going to achieve a useful outcome.

Guidelines suggest that programs should be stress-tested for limited generalizability or potential misuses, and all of this should be supported by ethical guidelines that are integrated into the process and that help developers and potential users to think about the extent to which ethics, fairness, representation, transparency, responsibility, and human well-being and dignity are made central throughout the process.

There's many examples of really great work being done in this area. There was just a recent special issue in the journal *Personnel Psychology* on machine learning applications to personnel selection. Landers and Behrend published an article in *The American Psychologist* with a detailed table of components of the AI system that need to be audited for fairness and bias, and they provide accompanying questions that are developed from psychology as a research base in measurement and prediction. There's also work in the occupational health arena that considers AI and occupational health disparities, as well as uses of AI to improve workplace safety.

Psychology is also critically relevant for the processes of deploying and continually monitoring AI systems. Here there are many different considerations. For example, when deploying AI in the workplace, psychological research, which suggests that there needs to be careful thought devoted to what functions or roles AI is anticipated to take on, how these functions augment or replace people's tasks, how people in the workplace feel about AI, the speed of adoption, the expectations and opportunities for skill development, how decision makers will monitor the impact on the workforce, the processes that will be in place for making changes to the AI system, and so on.

Again, all these things should be thought of with psychology in mind, in particular, with consideration of issues of human dignity, trust and acceptance, identity, quality of work and life, and fairness and ethics. Otherwise, people either won't use the systems at all or the systems will do harm. There's estimates that AI development projects fail about 80% of the time, in large part due to these types of issues.

Of course, systems need to be continually monitored for unintended consequences, especially as systems move from the development stage to being deployed in the real world and are continually being trained on new information. This is what makes AI so different than other technology projects. With other technologies, you generally know the outcome you're going to get, and it's generally the same over time. This is not necessarily the case with AI, and this is why it is so important to build systems and processes for continual evaluation.

Again, there's great work being done in psychology and the social sciences in these areas. These are just a few of many possible examples. The first one, I worked on with some colleagues trying to create a functional identity framework of reactions to AI. Recent research has looked at how managers view those employees who are working under algorithmic management. Those studies have found that workers are seen as less creative and less deserving of resources for innovation when they're being led by an AI boss or team lead.

Recent research has also looked at who benefits more or less from an AI partner and has suggested that AI may be more beneficial for workers lower in

conscientiousness and helping them to develop self-advocacy for their roles. There's also been a number of papers examining how, when, and why electronic monitoring leads to psychological reactions and to negative outcomes in the workplace.

Even going beyond individual AI systems, psychological science is well-positioned to understand and study the implications of AI for society more broadly. I'm sure many of you are aware that polls find that there's considerable anxiety about AI. The American Psychological Association's 2023 Work in America survey finds that 38% of people worry that AI might make some or all of their job duties obsolete in the future. Pew Research has found that 37% of people report being more concerned than excited about AI use in daily life. That number has actually risen to 52% in the most recent Pew survey on that topic. Gallup reports that 79% do not trust businesses to use AI responsibly.

Psychological research would tell us that reactions to AI are part of and contribute to a larger context. Essentially, to predict how the disruption created by AI is going to turn out, people tend to look at what's already going on in their workplaces and in their societies. Therefore, we should not be surprised that those who already see work as a harmful place where they're not supported by their employers also tend to be worried that AI is going to be used in a harmful manner.

This occurs at the country context as well. These graphs illustrate that the countries that have the greatest levels of income inequality also tend to be the same countries in which people are more likely to view AI and robots as workforce threats. I think this suggests that AI anxieties are not irrational. They are reflecting longstanding issues of trust and poor quality work design, as well as existing social systems that have tended to create winners and losers. Here psychology is well-positioned to investigate how even the general notion of AI can shape individual and societal outcomes, such as fear, anxiety, loneliness, job insecurity, sentiment towards outgroups, and career choices. Again, there's already a lot of work that's emerging in these areas.

That brings us to the role of psychologists in implementing change. Although AI is new, psychologists have been involved in change efforts for decades, and I would argue that much of the research applies here too. The literature on organizational change has shown that reactions to changes are driven not only by the changes themselves but also by change processes. For example, participation, communication, justice, supportive change from respected others, management competence, and change readiness. Again, also relevant to implementing change are thinking about the harms and the benefits and the ambiguities associated with the change itself, and how people develop self-efficacy for being able to cope with change.

However, when we're talking about AI, we need to keep in mind that the change here is continuous. The technology itself is changing over time. Societal reactions are changing over time. Regulation guidelines are changing over time, and this is going to have implications for how employees and employers adapt. Again, AI is also a little bit different than other types of change initiatives because there's already this context of excitement and fear surrounding AI, and therefore there's a need to facilitate a discourse that acknowledges and accommodates these views.

Again, I think there's a large opportunity for psychology to apply knowledge that has already been developed in the field and is being developed in the field to address a whole host of issues involving AI from basically AI design through implementation and change. With that, I will stop here and I look forward to hearing comments and questions.

Karen: Thank you so much, Dr. Shoss. That was very insightful. You gave everybody a lot of food for thought through your brief presentation. Hopefully, I can advance it. There we go. In case you have further questions for Dr. Shoss, here's her contact information. We've got a lot of really great questions coming in through the Q&A. If you have those questions submitted, we are going to go through and pick out a few of those. If you haven't typed in your question in the Q&A, now would be a great time to do that.

Actually, one question goes back to some of that Work in America data. What is it that you make of the patterns of why are people both afraid of losing their jobs due to AI but at the same time think that it's going to help them to be more productive?

Dr. Shoss: This is a great question because I think it speaks to just the sheer uncertainty and the wide variety of narratives going on around AI. It's interesting. This similar type of disconnect also occurs in other types of poll data. For example, Pew Research finds that 62% of Americans believe AI will have a major impact on workers generally, but only 28% think AI will have an impact on their jobs personally.

I think what's going on is people are hearing all of these discussions about AI unfolding both at a macro level, so when you read the news or listen to podcasts, as well as at a micro level in your own work. I think people are really trying to reconcile all the different narratives. There's all this discussion about AI can improve productivity, it could take tasks that are boring or rote and that seems like a great thing. Then there's also discussions that AI is going to automate everything and people will be out of work. It is challenging to reconcile these different narratives.

Again, I think that's also an opportunity for our field because none of these narratives are predestined to happen, right? What's going to happen with AI is going to be driven by the voices in the room who have impacts and insights about the way that these systems should be developed and the way that they should be used in workplaces. There is a role for our field. APA has certainly taken a lot of leadership around that in trying to steer the future of AI to tools and technologies that can help people achieve goals rather than to replace them.

Karen: Right. We've got a lot of questions that have covered ethics, and there was an earlier webinar in this series that touched on some of the ethics of regenerative AI, but I wonder if you could maybe address some of the big ethical considerations that we should be considering as psychologists?

Dr. Shoss: That's a great question and I think it's great to have ethics come up in many different conversations about AI because that really is a central concern here. There are issues of built-in bias in the system. Basically, if AI models are trained on data that are systematically biased in some way, that's going to come out in the models, and these models are applied at a broad scale, so we're going to see that

bias amplified. We need to think about the built-in bias in the system implication for disparities in outcomes, we need to think about trust and automation, prioritization, privacy. There's a whole bunch of issues here and APA actually has a fact sheet on AI that lays out some of these issues very nicely.

Also, *The American Psychologist* article I mentioned before by Landers and Behrend have this excellent table of just all the questions that we should be asking when we are either evaluating or developing an AI system.

Karen: Right. We've also gotten several questions about how psychologists can operate themselves into working in AI or what kinds of soft skills and AI knowledge would we need to know. In the future, what kinds of skills do you think will be necessary in the future, especially some of the skills that have a foundation in psychology?

Dr. Shoss: That's a great question. I think so many of the skills that are part of training in psychology very much apply. Understanding of measurements, statistics and ethics and psychometrics, development of psychological theories of behavior, I think those things are going to be very, very helpful in helping AI development teams or evaluators of AI to think about the quality of data and the quality of inference that's made and how to act on those inferences. I really see psychology as a foundational discipline to AI because psychology is really crucial to issues of AI development and whether AIs are used or are useful.

I think we don't all need to go out and get data science degrees. I think going to firm up our knowledge and our ability to communicate what's being gained from psychological research and to try to be in the room of discussions from the ground up and to help organizations think about these broader issues of ethics and representation from stage one.

Psychologists are good at this. This is something we're trained to do, right, to work with our organizations and work with groups and think, okay, who's going to potentially be impacted by this? How can we bring them into discussion? How can we facilitate a model of evaluating and setting up benchmarks for a project that's based on what we know about ethics and diversity and human well-being and privacy. Bringing all that psychological knowledge into this space of AI, I think, is really a huge opportunity for the field of psychology, and again, something I think psychologists are already starting to do.

You certainly see APA have a voice in this conversation. I've seen folks like Fred Oswald who serve on the National Artificial Intelligence Advisory Committee. All these efforts are great and can be done at broad scales and at local scales in folks' own circles and their own organizations.

Karen: That's pretty consistent, I think, with some of what I've observed in talking to psychologists who actually work in AI. They said similar things about they don't have a strong background necessarily in the technical sides of artificial intelligence, but they're brought in specifically because of their expertise in psychology and how they can use psychological science to help create AI tools and models going forward.

We've got a full, plenty of questions. I'm sorry, we will probably not have time to answer all of them because there's some excellent ones here. You mentioned earlier about human dignity and how that should be the center of the AI life cycle. How do you define human dignity?

Dr. Shoss: That's a great question. I view human dignity as really a broad concept, but it really encompasses the value and the esteem and the respect that we afford others. I really like this term because I think it includes concepts such as well-being, and it sets up a nice juxtaposition with discussions about dehumanization that we hear very often. We talk about automation. Instead, we're thinking about how we can promote human flourishing essentially. How can AI serve humans rather than the other way around?

Karen: Right. That's excellent. Let me see. We also often talk about equity, and you mentioned it a bit throughout your presentation. I think some of the questions have to do with non-WEIRD societies, and that's an acronym that I forget exactly what it stands for, but white, educated, Western, those kinds of implications. What do we need to be thinking about moving away from those kinds of populations?

Dr. Shoss: That's an excellent question, and I think it speaks to what are becoming very active conversations about the nature of data that our inferences about human nature are based on. If we're developing AI algorithms or evaluating them, then we want to think about issues of generalizability of representation in data. We also want to think about, on the use side, potential disparities in things like AI literacy.

The National Artificial Intelligence Advisory Committee just released some AI literacy recommendations saying that there are race disparities in awareness of AI in people's daily lives. Again, I think this is where psychologists come in and where AI development would be greatly advanced by including psychologists who tend to think about these types of issues and talk about these issues and thinking about this from square one because it's much better to design a system with these issues in mind and to think about it purposely than to try to reengineer something on the back end.

Karen: We've also gotten a lot of questions about implications of AI for psychological practice. For example, how might it be involved in testing and diagnosis, or where's the decision-making point? Will human psychologists be replaced by AI?

Dr. Shoss: That's an interesting question because I think every discipline is having a moment of reckoning of what do we do with this? Is this going to replace us or how should we view this development? Should we pursue it? What should we do here?

My expertise is certainly not in clinical psychology, but I would offer that when thinking about it, again, we would want to think about, what are the types of tasks that AI could help with, and what are potential opportunities there, for example, in expanding access to psychological care, and then what are the potential risks, and how can we study this in a way that gets good knowledge, but also does it at a small scale initially, and then assuming these systems end up being developed, and I believe some are in development for providing treatment, then what are the metrics by which we need to evaluate this, both on small and large scales.

I also think, again, AI probably presents a whole new set of opportunities for psychological practice because these systems need to be evaluated. They're constantly developing, which means they need to be constantly evaluated. I also think there's a role for folks with backgrounds in psychological science to think about issues of technology readiness, of training, of change management. Again, I really think that although there are very significant risks associated with AI and concerns about it, it also presents really an opportunity for the field.

Karen: Right. I would agree with that as well. We also, going back to something that we wanted to talk about, the workplace and how employees and employers can be prepared. How do you think AI will help employees to navigate change in the workplace? I think we were thinking a little bit more about just navigating change in general, but what's the role of AI in doing that?

Dr. Shoss: That's great. How AI itself can be an agent of supporting change. Again, we could go back to psychological research on what matters when you're making change, the harms and benefits of change, ambiguities around change, and processes. Perhaps there is a role of AI for trying to mitigate some of the ambiguities by providing resources and real-time as-needed training for people to figure out how to use different systems or deal with different changes or navigate complex policies. I think there's certainly a role for AI there.

Again, assuming that there's worker buy-in to, okay, would this AI type of system be useful or not, and is there regular transparent monitoring and evaluation and thinking about consequences, both intended and unintended consequences? The other thing I'll say about that too is that we want to keep in mind the saying that what gets measured gets done.

Especially in an age of datafication or if we're using AI systems to monitor metrics, then there needs to be a real careful consideration of whether this is going to create misaligned incentives.

Karen: We've been talking a lot about employees and employers, but we also know that within the field of psychology, there's a sizable number of psychologists who are self-employed, maybe because they work in independent private practice or they work as consultants or business owners. Are there any unique considerations for psychologists who are self-employed?

Dr. Shoss: It's a great question, especially because not only are there a lot of psychologists who are self-employed, but I believe it's nearly half of the US workforce works in small businesses. For most of these individuals, when we're talking about AI, we're not necessarily talking about AI that becomes internally developed to solve an internal problem. You're talking about using AI systems that are either built for different purposes or off-the-shelf AI systems. They're again, all the ethical issues we've talked about all still apply. Depending on the system and depending on types of regulations that get put in place for AI, it may be more or less difficult to try to evaluate issues of validity or reliability with AI systems.

I think also there's very important questions and discussions going on about issues of data storage, data ownership. There's a move towards open and transparent AI, which may help being able to answer some of those questions as well.

Karen: I think it's really important to use our research skills to try to understand, at least on a basic level, what these tools are doing and what are they doing with the data that you're feeding into them because it's not always transparent what the tools, the AI tools are doing. You have to be, I think, extra careful with confidential, proprietary, or private information because that can sometimes become part of the training data that these tools are using as they are going forward, which is fine if you make sure you've got the right permission for that, but there's certainly some privacy and confidentiality concerns as well.

Dr. Shoss: Yes, absolutely. Those are major, major issues. Again, just knowing how those are resolved at one time is going to be limited as these technologies continue to be trained on new data sets incoming, and as, again, regulations and guidelines change. This really is a moving target and folks looking to adopt AI systems are going to need to be aware and think about these issues.

Karen: Right. I was encouraged by the recent executive order on safety and privacy in artificial intelligence that struck me as being very person-centered. It'll be interesting to see how those proposals actually translate into regulations going forward in the future.

Dr. Shoss: Yes, absolutely. Actually, this morning, the *Wall Street Journal Tech News Briefing* talked about proposals to create a nutrition-type label for AI use. In particular, they were talking about healthcare associated with calls from the, I believe it's the Office of the National Coordinator for Health IT. It does seem that there might be a movement towards this kind of reporting and transparency, or at least I hope that there is. I think that would be something that our field and the research going on in psychology and organizational change and ethics would support.

Karen: Right, yes. Absolutely. I wonder if we might have transparent labels going forward that this content was generated by AI and what kind of effects that have on people. For example, if you have a scalable communication tool that appears to have a personalized message from your university president or a leader in your organization, but turns out it was actually generated by artificial intelligence, and maybe it has your name in it. How do people react to that? I don't know the answer to that question.

[laughter]

Dr. Shoss: No, it's a great question and there's some emerging research and some research grant, I believe it's at University of Southern California, looking at what happens if AI is generating negative feedback because negative feedback is difficult to take from other people versus positive feedback where I think we'd all really like our positive feedback to come from a person, not necessarily a computer algorithm, but I think there are just so many open questions to research on this.

Karen: Right. Lots of opportunities for psychologists to help design those studies to track the safety and the transparency. Not much different than an extension of just the skills we have anyway. Well, we've also gotten a lot of questions about career pathways or what the future might look like and where are these kinds of jobs. Where might the opportunities for psychology to make an impact exist?

Dr. Shoss: Again, these are really excellent questions. As I've done research in this area and I've tried to look at people's predictions about the future of work, one statement has always stayed in my head, and that is whatever the future will be, it's not what we think it will be. In other words, what we think is going to happen is probably the least likely of all the scenarios to happen.

I think if I want to speculate or maybe make a charge to action for the field, the more that we can share the value and the opportunities that the field of psychology and social science can present and offer to AI from design, deployment, development, to evaluation, I think then we'll see more doors open and more conversations happen where folks with a psychological background are invited to be part of those conversations.

Again, not only at technology companies but also with governments, with small business, with big business, I think there's a lot of opportunity because again, for these, there's a lot of money that goes into the building these systems, and for them to be useful and used is going to take a knowledge of human behavior and of, again, issues of validity and reliability and those are skills that psychologists have.

Karen: Yes. Again, I think we mentioned this earlier in talking to psychologists who are working in these areas, they've highlighted the research skills specifically without all that much prompting from us in terms of trying to describe what skills are important. It's the same skills that's applied in a different and really important context.

We've also seen a lot of discussions around how psychologists can help regulate and educate around responsible use of AI, which is certainly very, very important. Has technology advanced to a point where the existing AI tools are appropriate for psychologists to use in their research, teaching, or clinical work? This was posed by somebody who raised some questions around privacy, autonomy, liability, and so on.

Dr. Shoss: Yes, it's a great question. I wish I had an easy answer because part of the challenge is there are just so many different AI systems and they're all developed in slightly different ways and they're changing constantly. Even, for example, ChatGPT, two people could put in virtually the same question and get out two different answers. I think the nature of the domain and the nature of the changes means that there's unlikely to be a universal answer to that question, and it's probably going to depend on the specific application, the research that's been found about the validity for that specific application, and the field's ongoing understanding of, again, implications both intended and unintended.

I apologize, it's probably not a very satisfying answer and maybe we'll get to the point where there is a database full of evidence for specific AI programs that folks can contribute to and share their experiences with and maybe that's how we can start building a better understanding of a field of what we can use and when.

Karen: I am certain of one thing is that there's a lot of uncertainty, that's for sure and that I think we all have to be comfortable with this state where we are right now, where there's a lot of rapid changes that are probably going to happen, but we don't necessarily know what they are and that's okay. We don't all have a crystal ball or know what the future is going to look like.

We've also gotten a lot of questions about implications for education and they all boil down to how can you assess whether students learn or-- so sometimes there's a focus on detecting whether students are using AI tools to complete their work. What do you think about that? What kinds of implications are there for education?

Dr. Shoss: Again, excellent questions. I think they apply not only in education but also in the workplace where employers may have different views on whether or not they want their folks using AI or which AIs to supplement work. From what I've seen, it seems that it's pretty impossible to detect with accuracy, which means then we need to think about other options.

For example, I've seen people create prompts for class assignments asking students to use AI such as ChatGPT to answer a question and then to evaluate the quality of AI's answer and to show what AI has gotten wrong or has misunderstood or has omitted about a topic. I think things like that are probably useful assignments in helping students develop critical thinking, also helping students to develop disciplinary thinking.

How would a psychologist look at using ChatGPT for therapy or for personnel selection or something? What are the sets of issues that the field would suggest that we ask and we try to answer and how can we evaluate that? What kind of data can we collect? Yes, I think we're at the precipice of a change where we don't entirely know, as Karen was saying, where it's going to go. We can try to use the values and the knowledge base of our field to try to educate students, try to educate employers, and try to continue to advocate, again, for thinking about human well-being and dignity as the center of these processes.

Karen: Yes. I'm glad you highlighted the role of critical thinking in some of these prompts and possible assignments and ways to assess students because those are the very same kind of skill that employers are looking for. I think students who can master those skills show that they have critical thinking, can apply the same approaches that we use in other contexts to what's being generated by generative AI and other types of tools. They'll be well positioned for the future, especially in bringing in complex situational information, which may or may not be part of the models that the artificial intelligence is built on. I see nodding, so.

[laughter]

Dr. Shoss: Yes, absolutely. I absolutely agree. Understanding of context, understanding all the critical thinking skills that we really try to try to help our students develop are just going to be so crucial. I see a movement away from rote memorization and a movement towards trying to think about, okay, here's the information. What can what inferences can we draw? What are the ethical considerations? How should we act in whatever context we're talking about on the

information that we have? Also, what should we do if perhaps they were trained to do a job alongside AI? How do we know if something's gone wrong? How do we revert back to doing it ourselves if a system stops working for a period of time?

Karen: Indeed. I think we have time for maybe one or two more questions. There was a good one that was submitted earlier about using AI to find answers and clarity on writing material, but somebody who wasn't really that impressed with its ability to replicate or replace the human touch, such as in the thinking and the human communication style. Do you think that if we keep training AI, it will one day threaten the human existence? Very scary, essential question there, so good one.

Dr. Shoss: Yes, and I think, I personally don't know, but I think the AI developers don't know either. I was listening to a podcast that was playing a panel discussion from the Code 2023 Conference. Ajeya Cotra, who is a program officer at Open Philanthropy, used the analogy of a toaster or a nuclear weapon. Her point was that there's so much unknown about AI that even those in the industry don't know if they're making toasters or nuclear weapons, or if their toaster-making project may inadvertently create weapons.

It might be a, I don't know, very illustrative example of the uncertainty but I think that's, again, the challenge that we're dealing with and the opportunity of trying to have some voice in the direction that these technologies go and thinking about, okay, what would be a very useful application of technologies that are good at predictive modeling? What would be useful for humans and for society and what wouldn't? Maybe we should do the things that are going to help us versus the other things that don't seem to have a direct benefit or value to issues of human well-being. That would be my broad answer to that question.

Karen: Yes, and you keep circling back to human dignity and being human-centered, and I think that's really important as we start to think about going forward. We have just a couple of minutes left today. As we wrap up, do you have any final thoughts for our audience today?

Dr. Shoss: I, again, want to thank APA for inviting me to be part of this conversation and to really applaud the work that APA is doing in this area to try to get messages from psychological science out to a variety of folks in this area, from the technology development side to the government regulations side. Again, thank you all for coming. I, again, think there's really an opportunity and a challenge for all of us in the room to take the knowledge base of the field and to apply it and to find ways for us to have a voice in conversations about AI and to make sure, again, returning to this point of human dignity and well-being and ethics, and really to make sure that that becomes the central focus of these technologies and how we evaluate them.

Karen: Unfortunately, we are running out of time for today. Thank you so much, Dr. Shoss, for joining us. You provided a lot of really, really useful information for our audience today. Thank you to all of you in the audience for your very thoughtful questions. I wish we had had more time to ask all of them, and we tried to get to as many as we can. Thank you to the APA webinar support team behind the scenes, without whom we would not be able to make this possible today. Again, thank you all for being here today.

As a reminder, the recording of this presentation will be emailed to everybody within two weeks' time, and we'll also make sure we send out the presentation slides. You will also be able to view the recording on the APA website. As soon as the webinar has ended, a short survey will appear on your screen. We'd greatly appreciate it if you take a moment to respond to the survey to let us know how we did and how we can improve in the future. Be on the lookout for more APA webinars on AI and generative AI. There's certainly a lot more to cover on this topic. The web address for more information from APA on these topics is on your screen.

That's it for now. Thanks so much for joining. Be well and have a great rest of your day.

[00:57:26] [END OF AUDIO]