



A Chat about GPT3 (and Other Forms of Alien Intelligence) with Chris Dede

Melissa Warr¹ · Punya Mishra² · Danah Henriksen² · Lauren J. Woo²

Accepted: 21 March 2023
© Association for Educational Communications & Technology 2023

Keywords Creativity · Technology · Education · Research · Artificial intelligence · Future · Learning

~ AI as a tool in music-making is fine, but it's always going to be the humanity in music that makes people want to listen to it — Jacob Collier (musician)

~ AI is not a weak form of human intelligence. It is an alien intelligence. It is as alien as anything that you're going to find in outer space — Chris Dede

~ Some people worry that artificial intelligence will make us feel inferior, but then, anybody in his right mind should have an inferiority complex every time he looks at a flower — Alan Kay

Introduction

Since 2016, this column on Rethinking Technology & Creativity in Education has shared insights about creativity, learning, and technology from top scholars in diverse disciplines. In our last two pieces, we celebrated this decade of writing ending with a quote from Neil Postman where he argues that the impact of the advent of a new technology or medium never happens piecemeal (Postman, 1998). It transforms everything. For instance, he argued that the world after the invention of the printing press or after the invention of television was not the same world with some new technology thrown in. Those new technologies and media altered

most aspects of life, politics, family, education, religion, industry, and more—demonstrating how technologies can be disruptive and affect human culture at its foundation. We see similar effects within our lifetimes, with the invention of the Internet or social media technologies. Most conversations about the role of technologies in education often focus too narrowly on classrooms or other formal educational contexts and often pay less attention to how new media change the cultural ecology within which education functions. Given that, we noted how going forward in this series, we will strive to think more expansively, recognizing that changes in the societal context that education exists within are often larger than the immediate or obvious impact on education (Keenan-Lechel et al., 2023; Mishra et al., 2023).

So, over the next few articles in this series, we are seeking to dive deeper into Artificial Intelligence (AI) technologies and their potential impact on education, creativity, and society. This kind of topical deep dive is somewhat similar to the series of articles we did last year around mindfulness and creativity.

This is a particularly appropriate time to do so because over the past few months the world has shifted under our feet in powerful ways. The introduction of AI text-to-image generators (such as Dall-E, Stable Diffusion, and MidJourney) have raised critical questions about creativity, creative agency, and the arts. Around this time, more popular discussions emerged around large language models such as Google's LaMDA and OpenAI's GPT-3. It was the very public introduction to ChatGPT3, (a chatbot trained on large amounts of Internet text data to respond to natural language queries and generate, at least on the surface, intelligible responses) that appears to have truly changed the game.

These tools, whether image or text generation engines, raise a whole range of important questions about authorship, creativity, education, and more. It is within this context that we seek to frame the next few conversations and pieces around AI, education, and creativity.

✉ Danah Henriksen
danah.henriksen@asu.edu

Melissa Warr
warr@nmsu.edu

Punya Mishra
punya.mishra@asu.edu

Lauren J. Woo
lauren.woo@asu.edu

¹ New Mexico State University, Las Cruces, NM, USA

² Arizona State University, Tempe, AZ, USA

In this article, we share insights from our conversation with Dr. Chris Dede who was for 22 years the Timothy E. Wirth Professor of Learning Technologies at the Harvard Graduate School of Education and now is a Senior Research Fellow at the same institution. Chris is one of most respected and cited scholars in the field of educational technology with research spanning a wide range of topics, including, but not limited to, emerging technologies for learning, infusing technology into large scale educational improvement initiatives, developing policies that support educational transformation, and providing leadership in educational innovation. In addition, and directly relevant to this topic, he is currently co-PI and Associate Director for Research of the NSF funded National Artificial Intelligence Institute in Adult Learning and Online Education.

Chris brought a deep historical sense to our conversation. His wealth of experience in the field of educational technology gives him a unique perspective to see the larger picture, identify themes that are truly important, and go beyond the immediate hype of a new technology and its seemingly transformative capabilities. The history of AI is deeply intertwined with the advent of the digital computer in the 1950's and 60's and the advent of the cognitive revolution in the psychological sciences. Most early pioneers in the field of AI saw themselves as computer science researchers who used ideas from information processing to understand the human mind and to develop computer models to make software programs more "intelligent." Because a key part of the cognitive revolution (and early research on AI) involved articulating how the brain was a learning mechanism, these colleagues also explored the concept of learning itself. Their ideas impact how we learn and teach today.

The Early Years of AI

As in any early field, the hype surrounding early AI technologies tended to overtake the reality of the tools and of our knowledge and capabilities. As Chris mentioned in our conversation, his first introduction to AI came over 50 years ago when he read an article published on the topic. He said the article "confidently predicted that within six years we wouldn't need teachers anymore because AI was going to take over." Clearly that has not happened—and in that regard he has mixed feelings about AI, being both "impressed by the progress" that has been made while at the same time disheartened by the many hype cycles that have come and gone, each with its "glowing endorsements that somehow never worked out."

Similarly, he expressed moderated enthusiasm over the newest trend in AI—the recent availability of generative large language models (such as ChatGPT3) and its ability to generate "performances that we have previously thought were limited to human beings." For example, while a search

engine might use AI to compile a list of resources, generative AI actually synthesizes the resources and provides an outcome. Chris explained that it is "leaping to the outcome as opposed to giving the human being the chance to shape the outcome." That said, although generative AI is a major advancement, Chris is skeptical of much of the hype around it, suggesting that it is "somehow less than it's being portrayed by a lot of media."

As is often the case, our conversation ranged over a wide swath of topics around the effects, uses, limits, and ethics of AI broadly, as well as the implications of these tools to learning and creativity. In the rest of the paper, we focus on some key themes that emerged from this wide ranging and insightful conversation.

AI and What it Means to Be Human

Throughout history, artificial intelligence technologies have challenged what we take to be unique human capabilities, undercutting some of our assumptions about what we can do better than machines. AI has surpassed human performance in domains that require speed, scalability, and quantitative capabilities, like chess, the Japanese board game Go, and solving specific mathematical or scientific problems. In each instance, these advances in AI have meant that we have narrowed what at that time is unique to being human. Central to this argument are the questions of "what is 'intelligence'?" and "can it exist outside of humans?" According to Chris, the definition of intelligence has changed over time "because things that we previously regarded as intelligent...we now see how an algorithm can do that." He described useful intelligence as "the kind of intelligence...that understands all the dimensions of the problem." These dimensions include not only facts and logic, but also culture, social relationships, and emotions. This involves a view of intelligence as not only capabilities or knowledge but also understanding how things make sense—understanding in terms of concepts and holistic or connected awareness (or even a kind of "common sense" that, as of yet, AI still lacks).

What makes humans human is their emotional, cultural, and social awareness or consciousness. This understanding of humans hasn't always been evident in the history of psychology or human cognition. Behaviorist psychology viewed humans as slaves to stimulus and response conditioning (Lecas, 2006). Cognitive science moved forward by considering thoughts, beliefs, and desires as being key to understanding how we think and operate in the world, highlighting cognitive functions such as memory, learning, and reasoning (Baddeley, 2010). Over time, we have seen many technologies take over these abilities, most recently with generative AI and its ability to engage in, seemingly, coherent text-based conversation. Until recently, attributes like emotional intelligence, creativity, and language-based

conversations were seen as being uniquely human, but an array of tools and new ways of thinking are bringing these into question as well (Henriksen et al., 2022). Although people have attributed human qualities to machines (such as the software ELIZA and the more recent example of Blake Lemoine, the software engineer at Google who became concerned about a perceived “sentience” on the part of the LaMDA model), Chris made it clear that there are certain aspects of “intelligence” that AI currently does not possess. In particular, humans are uniquely able to effortlessly integrate knowledge, embodied experiences, culture, social relationships, and more: cognition, emotion, and social consciousness are “all linked together richly in the brain.”

This integration of cognition, emotion, and social consciousness allows humans to engage in creative and artistic activities. For example, although a computer program might be able to interpolate from images on the web, the result is a facile form of creativity in which the computer is only repeating patterns it has been trained on. The computer is fundamentally limited by both its lack of understanding of the meaning conveyed by images as well as its inability to step outside of the rules that have already been instituted by the field. It is also limited in its ability to integrate culture and aesthetics in order to make something truly artistically original. A human painter, on the other hand, “who had some knowledge of the history of art and who had some knowledge of what’s considered beautiful in different cultures” could create something that resonates with other humans. Thus, while AI can create simulacra that mimic human creations, there is still an essential lack of intentionality. As the philosopher Frankfurt wrote, “What is wrong with a counterfeit is not what it is like, but how it was made” (Frankfurt, 2005, p. 47). That said, intentionality is often in the eye of the beholder, and AI systems can mimic human responses (albeit at a surface level) to make it appear as if their artistic creations are indeed intentional. For example, even in these early days of AI image generators, there have been instances of AI-created artwork winning art prizes, angering professionally trained artists (Metz, 2022). We might surmise that the frustration felt by human artists is due to the recognition that there is no intentionality, motivating emotion, or understanding within the AI’s efforts, which most artists feel or perceive when they themselves create. However, these kinds of tensions complicate the very notion of authorship or creatorhood in creative domains.

Beyond the idea of intentionality, our conversation also focused on the differential strengths of AI vs humans, characteristics of AI that make it very good at some things and not as good at others. For example, according to Chris, the nature of intelligence of AI systems is fundamentally different from those of humans. He described AI intelligence as being an “alien intelligence” because its workings are fundamentally different from those of humans. And indeed,

even before the explosion of popular AI others have tagged AI as a kind of ‘alien’ intelligence whose functioning may not be fully graspable to humans (Barrat, 2013). One of the fundamental critiques of large language models is that, as they are trained only on terabytes of text, they have no understanding of what these texts mean; they have no internal conceptual models of how the words connect to anything in the real world. This means that while they can string together words based on probabilities, they are also liable to ‘hallucinate,’ or make up content, without any knowledge of doing so. That said, AI can identify patterns and extrapolate them to other contexts in ways that humans would find either difficult or impossible. For instance:

I'm pretty sure that nobody has ever done a study on the effects of mid-range climate change on flooding, specifically on street corners in Des Moines, Iowa. Generative AI can make that prediction because it can pull data from topological databases, from meteorological databases, from other kinds of databases, and come up with a forecast that's beyond just sort of gluing together stuff on the web. It's actually a powerful form of big data and large language models working together to create something.

Yet, Chris was quick to point out that though AI can draw these types of conclusions, the accuracy of its extrapolations is only as good as the data it is trained on. Additionally, it can appear overly confident of its conclusions (even inaccurate ones), because it lacks awareness of biases and inaccuracies in the underlying data. In other words, it lacks a conceptual *understanding* of reality. As Chris described it, AI “has ‘no error bars’” resulting in a false sense of precision, particularly when it comes to topics directly connected to humans and society. He makes an important distinction between the role that AI can play in the hard sciences and the soft sciences. Expanding on this he said:

If you ask an AI to explain something in the hard sciences, like, what are the different cooling mechanisms if you start with hot water or with cold water, you're probably going to get a really good explanation, maybe a better explanation than the typical high school science teacher. On the other hand, if you ask a question about human behavior, how can people hold contradictory beliefs simultaneously and act on one or act on the other without appearing to notice a contradiction between them, AI will be at sea in something like that. And that's because much of social science is also at sea in terms of something like that.

Moreover, though AI systems can draw conclusions based on a specific data set, their expertise is narrow and does not take into account the full complexity of human situations. For example, Chris described how AI might be able

to provide treatment options for a patient dying from cancer and forecast the efficacy of those treatments. However, it does not take into account the human experience:

If someone's dying of cancer, and you want to advise them about treatment options, you're not just making a forecast of, well, with this treatment, you'll live this long and with this other treatment, you'll live that long, which is what AI can do. You're saying, what is your belief about quality versus quantity of life? How is your family going to be influenced by when you choose to die, and what are your spiritual beliefs in terms of life? AI does not have that kind of intelligence.

One of the important points made by Chris was that what AI ultimately lacks is “wisdom”—an ability to integrate context, experiences, beliefs, culture, physicality, spirituality, and ethics of a given situation. He explained:

The wise use of knowledge requires all sorts of things that AI cannot do. [Human wisdom] really involves a kind of knowledge of culture, and the knowledge of what it means to have a body and what it means to have spiritual values that AI is not going to get to.

This type of intelligence is crucial for ethical decision making. Chris described his work in the creation of *Quandary* (quandarygame.com), a game that teaches ethical decision making. In the game, players are tasked with gathering information to make decisions about problems on a new planet. Players sort facts and opinions, consider the perspectives of different villagers, and ultimately suggest a solution to the quandary. In this context, AI might confidently select an optimal solution for a short-term solution. However, it would struggle with making tactical decisions:

AI has no sense of what it may be creating in the Big G game. So in the same way, we wouldn't really want AI making the diplomacy decisions in the negotiation, because even if you get a tactical win, you're likely to get a strategic loss.

AI also struggles with ethical decision making because, as addressed previously, it has a limited ability to consider contradictory beliefs, what Kuhn once termed ‘essential tensions’ (Kuhn, 1977). Much of the conversation around AI focuses on how AI can replace humans. However, Chris takes a different stance: because humans and AI each have unique forms of intelligence, they can work together to accomplish more complex tasks. Connecting to his earlier characterization of AI, he noted that “the strength is that because it’s alien, it truly is complimentary.” This opens up the question of what this alliance between humans and AI could look like.

From Artificial Intelligence to Intelligence Augmentation

Much of the concern about AI comes from a fear of “replacement”; for example, in education, a concern sometimes voiced is that AI will replace teachers. However, Chris believes that, given the limitations of AI we described above, this is unlikely. Rather, AI has the potential to work *with* humans, leading to a new way of working, often called intelligence augmentation (IA). Here, the focus is not on AI down-skilling but actually creating the opportunity for upskilling human abilities. Chris explained:

What I've always been interested in is what people in AI can do together based on complementary strengths and the whole being more than the sum of the parts. What can a human AI combination do that neither the human nor the AI alone can do?

Combining the strengths of AI (its ability to synthesize and draw conclusions from large amounts of data) with human abilities (integrating embodied experiences, social contexts, emotions, etc.) can result in a powerful way of acting. For example, Chris described a current project at the National AI Institute for Adult Learning and Online Education (AI-ALOE), 2023, where he is Associate Director for Research, which is focused on building intelligent assistants for university professors. He explained:

If I upskill to more deeply understand and personalize learning for my students, students from different cultures, students with different kinds of life challenges and so on, then I'm getting IA, then we're able to do more collectively than I could without the assistance from AI. And I think that that's going to be an interesting contribution of this Institute. .. building these assistants, putting them into the crucible of practice, and seeing the extent to which people can and do upskill or whether they just let themselves be de-skilled.

IA draws upon the strengths of AI technologies and humans. Developing these tools and processes has the potential to help us go beyond “just doing things better” to “doing better things” (AIALOE.org).

What implications does this view of intelligence augmentation have for education, learning, and creativity? According to Chris, the main challenge in this new ecosystem is to focus educational systems on the human part of the partnership: what humans do best and how they can complement the strengths of AI technologies. This calls for changes in school curricula and methods because, as machines become better at doing certain tasks, “machines take over some of what people have been doing in the educational system and [the educational system has] to change in response.”

In a previous piece in this series, we had discussed a similar idea writing:

The new forms of expression that are emerging today have significant implications for how we engage and interact with machines. In this “brave new world” machines take on a role in the creative process itself—partners to humans as it were. Seeing the computer in this light has significant implications for how we educate our students for the jobs of the future. Looking forward, we see that human initiative must blend with the capabilities of software programs in ways that have not been possible before (Mishra et al. 2013, p. 10).

We described examples of how algorithms which lie at the heart of AI systems can be used to augment the creative process. We wrote about the musician/composer David Cope who developed a program called Emily Howell that could collaborate with him in composing music. The program used an association network to evaluate musical patterns as good or bad. Cope would ask the program a musical question and the program would reply with a composition which he could accept, reject, or modify, and the process would continue until Cope was satisfied. He described this process as being equivalent to sculpting from a block of marble where the software program worked as a collaborator augmenting his expertise and capabilities.

What this means is that there can be significant augmentation of human creativity if we see AI as a partner to the humans who bring deep humanistic knowledge to the “dialogue.” The addition of the software does not diminish the role of the human, but rather human intuition, knowledge, and agency play a key role in this process.

Unfortunately, most educational systems continue to focus on more basic learning tasks that can be easily accomplished by AI. For example, many educators have worried about the ability of tools such as ChatGPT to write effective and coherent essays. Companies are attempting to create plagiarism checkers that can identify AI-written essays ignoring the heart of the problems: why are we emphasizing tasks that can be effectively completed by AI? We are “preparing people to lose to AI instead of focusing on what people can do differently and better.” Chris emphasized the need to educate children to work *with* AI to accomplish more effective tasks. For example, he suggested starting with an AI-written story, then

Give kids the descriptive story and say, turn this into a really compelling narrative that within your culture speaks to people in a deep way. They see themselves inside of it, they see the story relating to other kinds of stories that may be part of their culture and so on.

In this task, learners are developing unique human abilities: integrating culture, emotion, and identity into a meaningful

story. This task draws upon social emotional skills as well as creativity. Through these types of tasks, learners can prepare for a partnership with AI, doing better things they could not do alone.

One of the challenges of Chris’ perspective is the way that current educational systems emphasize testing, particularly through multiple choice tests which emphasize discrete skills and facts or even rote knowledge. Chris instead emphasizes the need to focus on performance assessments. For example, he described the use of immersive environments for performance assessments, specifically through digital puppeteering such as in Mursion, (mursion.com). These tools provide a type of “flight simulator” for interpersonal skills. In these tools,

AI works at the front end of those systems because it creates a much richer context that's evocative of those skills and authentic in terms of the settings in which you utilize those skills such as a pediatrician learning to elicit knowledge from a young child who's feeling ill. But, on the back end, it's got machine learning and all this rich data flowing into it from human behavior second by second within the simulation that then can be feedback to the coach and feedback to the intelligent coaching assistant, both as a learning mechanism, but ultimately as an assessment mechanism.

Ultimately, we should be preparing learners for doing what humans do best, and the types of assessments we use should reflect these priorities. The focus, according to Chris, should be on key conceptual ideas. For instance, he makes a distinction between basic arithmetic manipulations and understanding the number line. As he said,

You can look up both of them on the internet. Both of them seem to be relatively simple ideas, but the number line is in fact the foundation for a lot of higher kinds of mathematics. And really understanding the number line is something that people absolutely need to do, whether or not they get beyond estimation in terms of multiplication. Another example would be the periodic table in chemistry. Should you memorize Calcium’s atomic number and how many electrons it has? Well, no, you shouldn't, any more than you should be memorizing the capital of every state. Now, if instead you ask why is the capital located where it is? That's a deep question. If you ask in the periodic table, why did the elements above a certain point become radioactive? That's a deep question.

Assessment should, according to Chris, focus on the deep questions of a discipline. This requires not only new forms of assessment but also new forms of pedagogy to align with and suit such assessments. Advances in AI technologies, he believes, offer new tools that can support this type of

learning and assessment, but it will take shifts in educational systems to make their use a reality.

Conclusion

AI systems, Chris believes, have the potential to dramatically impact what we do as educators, but only if we view it the right way. Too often we approach technologies uncritically, accepting that, because they follow mathematical algorithms, what they output is “right.” He cautions against uncritical approaches that see these technologies as black boxes. People may believe that “if you average across everything on the internet, you are going to come out with an unbiased thing. Nothing could be farther from the truth. You are much more likely to come out with a biased thing averaging across the internet.” This is because “there are many forms of bias that are part of innate human decision making,” and they can show up in systems in a variety of ways—from biased algorithms to training sets that contain subtly implicit forms of bias as well as biases that are built into recommendation systems.

Despite these concerns, Chris is fundamentally optimistic about these tools and the role they will play. He does not believe, however, that it will happen by itself. Quoting Alan Kay, he said, “the best way to predict the future is to invent it.”

The same way we need to be very careful about AI saying, “I’ve studied macro history and these are the things that are going to happen.” Well, yeah, they’re probably the things that are going to happen if we drift into them or if we believe them and stop trying. But the human spirit of attempting to overcome what the dominant trends are leading towards is the subject of a lot of our hero stories. And the women and men who are heroes step outside of the trends and the predictions and say, I’m going to do something that appears to be impossible because I believe that it’s really important, and some of the time they succeed.

In other words, there is no predetermined future to how AI will play out in our world. The future is inherently uncertain because it does not exist yet—so we have the opportunity to create the future with our actions in the present. Chris made

the same point, saying that, “the future is what we make of it. We can change our destiny and change our future, even though it’s not easy.”

References

- Baddeley, A. (2010). Working memory. *Current Biology*, 20(4), R136–R140. <https://doi.org/10.1016/j.cub.2009.12.014>
- Barrat, J. (2013). *Our final invention: Artificial intelligence and the end of the human era*. Macmillan.
- Frankfurt, H. G. (2005). *On bullshit*. Princeton University Press.
- Henriksen, D., Creely, E., & Mehta, R. (2022). Rethinking the politics of creativity: Posthumanism, indigeneity, and creativity beyond the western anthropocene. *Qualitative Inquiry*, 28(5), 465–475. <https://doi.org/10.1177/10778004211065813>
- Keenan-Lechel, S. F., Warr, M., Richardson, C., Mishra, P., Mehta, R., Henriksen, D., & Gruber, N. (2023). A decade of rethinking creativity, technology and learning: Reflections with the deep-play research group. *TechTrends*, 67, 7–13. <https://doi.org/10.1007/s11528-022-00817-7>
- Kuhn, T. S. (1977). *The essential tension: Selected studies in scientific tradition and change*. University of Chicago Press.
- Lecas, J. C. (2006). Behaviourism and the mechanization of the mind. *Comptes Rendus Biologies*, 329(5-6), 386–397. <https://doi.org/10.1016/j.crvi.2006.03.009>
- Metz, R. (2022, September 3). AI won an art contest, and artists are furious. *CNN*. <https://www.cnn.com/2022/09/03/tech/ai-art-fair-winner-controversy/index.html>. Accessed 24 Mar 2023
- Mishra, P., Yadav, A., & Deep-Play Research Group. (2013). Rethinking technology & creativity in the 21st century. *TechTrends*, 57(3), 10–14. <https://doi.org/10.1007/s11528-013-0655-z>
- Mishra, P., Henriksen, D., & Richardson, C. (2023). *From crayons to AI: Widening the lens on educational technology and creativity*. TechTrends.
- National AI Institute for Adult Learning and Online Education (2023). <https://aialoe.org/>. Accessed 24 Mar 2023
- Postman, N. (1998). Five things we need to know about technological change. Speech given at NewTech98: The International Conference on New Technologies and the Human Person: Communicating the Faith in the New Millennium. <https://web.cs.ucdavis.edu/~rogaway/classes/188/materials/postman.pdf>. Available at: <https://www.youtube.com/watch?v=CZKUak1fYr0>. Accessed 24 Mar 2023

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.