Artistic Context, GenAI, and the Dilution of Intention: An interview with Ted Chiang by Andrew Erickson

June 18, 2025

The much-awarded science fiction author and essayist Ted Chiang recently joined Andew Erickson and his colleagues from Schaufler Lab@TU Dresden via Zoom to discuss his writings on GenAI and making art in the context of the Schaufler Lab@TU Dresden videocast format on "Data↔Worlds." The session explored the interconnected themes of worldmaking, language, and generative artificial intelligence (GenAI), drawing on the science fiction and essays of the guest speaker. We discussed how storytelling and language are foundational to human experience and understanding, with particular attention to the ways language shapes perception, decision-making, and the construction of reality. To open the conversation and introduce Chiang's work, Erickson asserted that Chiang's novella "Story of Your Life," highlights how language can transform awareness and facilitate new experiences of time and existence. This is exemplified in that work by the protagonist's encounter with an alien language that alters her perception of time, causality and free will. "Story of Your Life" underscores the power of storytelling to humanize science and argues for the ongoing importance of language and narrative to shape meaning, even as technology and data increasingly influence human experience.

This entry point allowed us to dive into the primary intervention of Chiang's own remarks, addressing the significance of creative decision-making in art and daily life, distinguishing human intelligence from artificial intelligence.

Note: some parts of the conversation that took place have been edited for brevity and clarity.

Ted Chiang: Thanks, Andrew. I'm going to say a few words about generative AI and art, expanding on the piece that appeared on the New Yorker website last fall.[1] In that piece, I talked about how I think of art as a concentrated form of intention and how generative AI is fundamentally incompatible with that, because generative AI is all about small inputs and big outputs. It basically represents a dilution of intention, and I don't believe that is compatible with art as I understand it.

There are many people who are profoundly moved when they see a painting. If you are such a person, I want you to consider a hypothetical scenario. Can you imagine a machine that generates an image that profoundly moves you and that can make a new image every time you hit a button? Let's forget about how this hypothetical machine might operate. Purely in terms of your subjective experience as someone who's perceiving the images it creates, is it possible for you to react with rapture to every image? I am going to suggest that this is not possible. I think that, at some level, the

[1] See Chiang, "Why A.I."

knowledge that this machine generated the image at the touch of a button and can keep generating more, no matter how often you hit the button, would actually prevent you from having the same reaction to any of these images that you would have if you, say, saw Picasso's *Guernica*.

If you are a lover of painting, you can have a transcendent experience when you visit a museum that contains a bunch of paintings that you revere. I think that what makes that possible is your awareness of what an exceptional sampling of human creative output has been collected into a single building. I don't believe that running an app on your phone to generate one thousand images per hour can give rise to a transcendent experience, because the knowledge that you can immediately hit the button again and again, ad infinitum will cause you to regard its output as disposable. Some AI boosters might say that this reveals an irrational prejudice on your part, because in this hypothetical scenario, there would be no quantifiable metric they could find which distinguished museum paintings from the images that the app on your phone generates. I would say that they are drawing the wrong conclusion. I think the correct conclusion is that the value of art is not solely inherent to the object. Whatever it is about a painting that moves you, it is not something that exists only as a measurable property of that painting. It also derives from your knowledge of the context in which that painting was created, and your reaction to that is dependent on your individual life experience. Art is all about context, both the context in which it was created and the context in which it is perceived, and to think that art can be separated from those contexts is to misunderstand the nature of art.

Now I'll talk a bit about the production side. There is, of course, the matter that there are a lot of copyright infringement lawsuits underway. Copyright is a legal issue that is separate from the issue of plagiarism. Plagiarism is the act of passing off someone else's work as your own; if you are unwilling or unable to make all of the choices needed to create a novel or a painting or a film, plagiarism is one way to save yourself the effort. This is different from copyright infringement. When Ansel Adams's photographs go out of copyright, anyone would be free to sell a collection of his photos for their profit and not pay anything to his estate, but that person could not legitimately claim to be the creator of those photos, because if they did, they would be a plagiarist. It would not be a violation of copyright law, but it would be a moral and creative failure that disqualifies them from being an artist.

Here's an analogy for thinking about generative AI: imagine someone offers a service where a person types into a kiosk "I have an idea for a book, and the plot goes something like this" and behind that kiosk is a vast library, and a librarian finds the book that comes closest to meeting the specified criteria, and then the librarian changes all the proper nouns in it before supplying the text to the requester in such a scenario. Can the requester claim to be the author of the resulting text? I would say no. If the library consists of copyrighted books, there is a legal problem, but even if all of the books in the library are in the public domain, the requester is still a plagiarist. The requester might

not know anything about how the kiosk works, so the plagiarism might have been unintentional, but even then, they have a dangerously naive idea of what being an author is. Having an idea for a book is not enough to make you the author of a book, any more than having the idea for a flying car makes you the inventor of a flying car. In the same way that turning your idea into a vehicle that actually flies is what qualifies you as an inventor, turning your idea into an actual book is what qualifies you as an author.

Large language models and text-to-image generators may not return results as close to the originals as the kiosk in my thought experiment, but I don't believe that changes the ethics of the situation. What would change the ethics of the situation for me is if, instead of making a request at a kiosk, a person reads some books from the library, thinks "I have an idea for a book," and then writes a book. That person can call themselves an author, even if the books that they read are still in copyright and even if the book that they write is derivative or clichéd.

I'm not trying to argue against the use of generative AI as a brainstorming tool. The author Philip K. Dick famously used the *I Ching* to plot his novel *The Man in the High Castle* (1962).[2] Whenever he needed to make a decision about what would happen in the novel, he would throw some coins, consult the appropriate passage in the *I Ching*, and then use that to determine what happened. So, the reason that that novel exists in its current form is because of the *I Ching*. However, it's important to note that the *I Ching* did not enable Dick to write novels faster; it was not a force multiplier. No one could reasonably claim that the *I Ching* enables people who have never written before to create a novel that they have had in their head for years. Some writers use the *I Ching* or tarot decks or various other brainstorming tools as sources of inspiration. It's certainly possible for software to serve a similar function, but in all of these situations, the writers are using a brainstorming tool to prompt themselves to write, which is entirely different from prompting a large language model to generate text. So, I think, if writers or artists want to use generative AI to prompt them the same way that they might use a tarot deck, that's creatively acceptable. However, I'm pretty sure you don't need a program that costs hundreds of millions of dollars to build to replicate the functionality of a tarot deck.

I also want to clarify that I'm not arguing that there is no way for machine learning to be useful for artists. One way to think about this is to bring in an idea from the world of user-interface design. A good user interface offers as many degrees of freedom as the system that the interface is intended to control. An airplane has more degrees of freedom than a car, so a pilot's yoke has to offer more degrees of freedom than a car's steering wheel. A digital painting program like Photoshop is a good control system for image creation, because images are a very high dimensional space, and Photoshop offers many degrees of freedom. Text prompts are a comparatively low dimensional interface, and if they are short, they cannot offer enough degrees of freedom to enable fine grained control. Any program that is useful to an artist will have to support many

[2] See Williams on Dick's use of the *I Ching*; see also Mountfort on the significant influence of the use of this technology for Dick's creative process.

degrees of freedom. They will have to offer a very high-dimensional interface. A good user interface is also predictable. Once the user becomes accustomed to the interface, they will have a good idea of the effect on the end product of any operation they perform. Existing generative AI programs do not meet any of these qualifications because, as I said, they are all about providing small inputs and getting large outputs.

The philosopher Tobias Rees has observed that the electric guitar was not a means for playing guitar faster[3], and I think this is a useful way to think about the issue. The electric guitar opened up new possibilities for musicians, and it facilitated the creation of new genres of music, but it did not make musicians more productive by any qualitative metric. An electric guitar is not easier to master than an acoustic guitar, and so it does not enable non-guitar players to suddenly become musicians. A person might prefer a certain genre of music that is associated with electric guitars rather than acoustic guitars, but they will still have to practice for years to become proficient in it. So, what does the software equivalent of electric guitar look like? I have no idea. It certainly might involve machine learning, but it will not involve generative AI. I should also point out that if a private company invests hundreds of millions of dollars in building a tool in the hopes that they will have the next electric guitar, you better believe that they will demand a royalty on every song performed with it. I think that alone would prevent whatever they build from having the influence that the electric guitar did. It's only if the open-source community is able to replicate this hypothetical machine-learning equivalent of the electric guitar that we might one day see good art result from its use.

Andrew Erickson: You hinted at something like an appropriate role for machine learning technologies in the production of art, literature or in scientific thinking. If there's an appropriate role, in your opinion, what would it be? How would we use AI, if we should, in the way that musicians engaged with the electric guitar, as in the example that you gave?

Ted Chiang: This question is easier to answer in the context of scientists or researchers instead of art. I was talking with a physicist who said that he thought machine learning could be a tool for scientists in the same way that the electron microscope is. Just as an electron microscope lets physicists see things that they cannot see for themselves, machine learning might be a way to discern patterns in data which humans cannot discern for themselves. That analogy makes a lot of sense to me. I think what we were looking for are machine learning tools that have some analogous function to electron microscopes, as tools that enhance researchers' ability to perceive things.

Generative AI doesn't operate this way. The argument I made about the importance of decision making in art is not limited just to art; I think it applies to doing good work in any field. If you are a professor of philosophy or literature and you're going to write a paper, do you want to think "I have this idea for a paper" and then have a machine write it for you? Would you feel like the author of that paper? Would you feel like you had done

[3] See Rees.

good work? I think that in order to feel like you've done good work, you need to have thought about every sentence in the paper, so every sentence is the product of a decision you made and elaborates on the argument that you had in mind.

I don't know if there's an analog to an electron microscope that artists use. We can talk about the possibility of an analog to the electric guitar for other fields of art, but no one has any idea what that looks like before it arrives. No one predicted the electric guitar before it was invented, and it's not because the electric guitar is so radically different from acoustic guitars; we can still easily recognize it as being a guitar. But even that non-dramatic alteration created whole new genres of music. As with the invention of probably every musical instrument that gained acceptance, people were sort of muddling around, trying to come up with something new and, eventually, someone devised something that people really responded to, and it took off. But it's not something that you can plan; it just has to happen. So, I don't think anyone can say what the software analog of the electric guitar will be.

It's also worth noting that for decades, people talked about how the computer would change fiction. Some of you might remember hyperfiction, or hypertext fiction, which people were excited about in the ‘80s and ‘90s. Hyperfiction didn't really take off. Nowadays we have interactive fiction or visual novels, which are related but, with no insult intended, remain relatively small genres. The interactivity that computers offer did not—despite some people being very enthusiastic about them and doing legitimately interesting work in that space—fundamentally transform fiction. You can't predict what’s going to capture people's imaginations and take off. I don't want to say that interactive fiction has been a failure because it’s not as successful as the electric guitar has been. Making the next electric guitar is a pretty high bar. I don't know if we should be using that as our standard.

Andrew Erickson: Your fiction imagines all kinds of really interesting technologies, and you're thinking through different ways with which people can engage the future. Science fiction more broadly imagines all kinds of future technologies. Which of the technologies that you have imagined in your fiction do you anticipate—or have you perhaps seen—to become real in the world outside of the text?

Ted Chiang: I don't think there's any anything in my fiction that which this really applies. The purpose of science fiction is not predictive; science fiction writers are not trying to anticipate future technological devices. I'm mostly drawn to science fiction because it is a good way of dramatizing philosophical thought experiments. I'm interested in philosophical questions. When philosophers propose thought experiments, they can often seem very abstract and hard to understand to non-philosophers, and it’s easy to wonder, why should anyone care about this thought experiment? One thing science fiction allows you to do is to create a world where this thought experiment is an actual thing people have to deal with. Then you can create characters who are profoundly invested in the issues that you are posing. If the piece of fiction is effective, your reader

can feel some identification with these characters, and then they will have a more visceral appreciation for why this thought experiment is interesting.

I also have to note that science fiction is a very big tent. It can do lots of different things. This is true for art in general. Art is not primarily instrumental. Art is not a tool for accomplishing any particular social goal, so it doesn't fail if it doesn't accomplish that goal. Art is valuable for its own sake. I don't want people to reduce art to any specific function. Art can promote certain political values, but that is not the reason that art exists. Likewise, I don't think science fiction can be reduced to any particular function. It is a genre of art, and any genre of art is its own justification.

Andrew Erickson: How do you view your works, if at all, as interventions into where things are heading—as cautionary tales? Or maybe a different version of that question is something like, how much does the science fictional mode allow us to model the future and act upon that model of the future to materialize the kinds of presents that we want to see. And on the other side of that, when should we not do what science fiction stories imagine.

Ted Chiang: One of the things that art in general does is remind us of what it is to be human. At a very broad level, all art has the potential to incite or encourage action, because by reminding us what it means to be human, it hopefully makes us aware of dehumanizing trends in society. When we are reminded of what it is to be human, then we are better able to see things which take away our humanity. All art can be a kind of call to action, a reminder of what is worth preserving. Hopefully, if we see threats to the things that we wish to preserve, then we will take steps to oppose those now.

As I said, I don't think that science fiction has a more clearly instrumental function than art in general. It can be more specific in depicting scenarios where certain people are dehumanized, and having that spelled out in an explicit fashion can be helpful. But there is also the fact that there's no way to control how people interpret art. As we've seen, people like Elon Musk or Mark Zuckerberg have clearly read a lot of science fiction, and they've taken exactly the wrong lessons from it, the opposite of what the authors intended. Jeff Bezos is a fan of Iain Banks's culture novels; those are pretty explicitly socialist, and he seems to be blind to that. So, there's no way to guarantee that people will take the intended lesson from your art. You can try to make art that allows only one possible interpretation, but that's probably pretty bad art. And even if you try, I don't know that it is possible to create something which can only be interpreted one way. I don't think there's a way that to construct a list of sentences that you could be sure Elon Musk would interpret the way you intended. This is just to reiterate the point that we should not be reading science fiction as a kind of roadmap or a blueprint. In fact, the people who do look at it as a roadmap or a blueprint are the ones who are pretty much, as we can see, taking the wrong lessons from it. Their misinterpretation might have something to do with the fact that they are looking at it as a roadmap or a blueprint.

Andrew Erickson: Some of some of the technologies that you've created in your stories call into question free will and personal choice. What is the role of free will in a world that is increasingly mediated, surveilled and controlled by technologies, many of which people haven't actually chosen to engage with? And how is the capture of personal data for the implementation of future technologies sort of handing over the future to forces beyond our will?

Ted Chiang: I firmly believe that humans do possess free will. And when I say that I'm talking about it in the context of the free will versus determinism debate in philosophy. In that context I am persuaded by the compatibilist school of thought. Compatibilism is the philosophical stance that the most meaningful definition of free will is not only compatible with determinism, but it also actually requires determinism.

That question of free will—and compatibilist free will—in a deterministic universe is a separate question from the question of free will in the modern age. Do we have free will in a universe bound by physical law? I would say the answer is yes, we have the capacity for free will. But you can have free will and be imprisoned; society is limiting your freedom and restricting your options. If we agree that agency is being taken away from you, that means it's possible for you to have agency. Having the capacity for free will doesn't mean that you are immune to manipulation. For manipulation to be a meaningful concept, there has to be the possibility of not being manipulated. You need to have the capacity for free will in the philosophical sense to have the discussion about what sort of options are being taken away from you by modern technological society. It is undeniably true that big corporations have a huge amount of data on us and are using that to influence our buying decisions and our political views, but it is also within our power to resist.

Once you are aware of these forces, you can reduce their influence on you. You can take steps to reduce the amount of data that you give them and be more cognizant about the things that you are being shown on social media platforms. You are not without agency in this situation. I'm not going to say that it's easy, and I'm also not going to say that there is some ideal of pure freedom that's achievable. In discussions about free will, people are often asking for a kind of freedom that's ill defined. The more that we know about the forces acting on us, the better equipped we are to make choices about them.

Andrew Erickson: I have to admit that I find many technologies seductive. They're so they're attractive and often in ways of which we're not totally conscious. How do you mediate that space where you want to engage with technologies that might have problematic aspects, e.g. data collection or screen addiction or whatever? This is something that reverberates deeply in your fictions; some characters struggle with the decision whether to give up some free will in order to engage with technologies that are transformative.

Ted Chiang: Free will is a separate matter from the question of how much one embraces or refuses to engage with technology. If you're a gambling addict, you're very susceptible to casinos, and you can try and avoid going into casinos. The problem is that the same sort of design that went into casinos is now on your smartphone.

Most apps and platforms have the addictive design that casinos and slot machines have. They are preying on people's weaknesses and turning people into a kind of gambling addict. You might not have been drawn to casinos, but now you are drawn to Instagram. That's a problem. It's a problem that our society basically allowed companies to infest the entire world with casino and slot-machine design, but we are not helpless. There are apps that you can install that will limit your usage of these things; you have those at your disposal, and you can use them to keep you from spending too much time online. You can uninstall apps. You can get a simple phone instead of a smart one, or you can install an app that turns your phone black and white to reduce its addictiveness. There are options.

Hopefully we'll be able to weigh the short-term displeasure of not being able to get on to Instagram whenever you want or having to look at a black-and-white screen against the longer-term benefits to your mental health that will accrue if you don't spend as much time on these things. Those choices are available to you. You can make them, and of course, you can also advocate for stronger regulation so that governments might restrain companies from infesting the world with addictive design.

Andrew Erickson: It occurs to me that the gamification of social media and adjacent technologies is increasingly distracting us from its externalization and takeover of our recorded memory. You've addressed the complete technological mediation of memory in your fiction. "Remem" in "The Truth of Fact, the Truth of Feeling" records everything that happens in the world from the various perspectives of different people who use it. The story at one point predicts that episodic memory—the record of personal and interpersonal experiences we have of the world—will become entirely technologically mediated.

Is that what's currently happening through social media, and if so, what impact will that have on the way that we experience the world and the way that we remember things in the past, when you have this detailed official record that can be called up at any time to look and actually see what the truth of that that moment was? How does that change the way that we interact as people?

Ted Chiang: I don't see a strong similarity between the "Remem" technology and the social media that we see today. There are two very different forms of technological mediation at work. The issues that I was trying to investigate in "The Truth of Fact, the Truth of Feeling" have to do with the question of: to what extent do we rewrite our own personal narrative when we are telling the story of ourselves? Are we doing it in such a way as to flatter ourselves? The technology posited in that story pushes against that

rewriting of personal history. Even if the technology is purely private, as Remem is, it would raise the same issues about the narrativization of the self.

What is happening with social media is something entirely different. The technological mediation happening there is one where people are not so much interacting with other people on a face-to-face level but through a layer of performance. They perform for TikTok, and that becomes the primary way that they interact with each other, instead of socializing in person. That is not to say that people are not also engaged in some degree of performance when they interact face to face, but that's a different kind of performance. As a quick anecdote, once my sister told me about an argument she had with her son, who was a teenager at the time. They were arguing, and at one point her son pulled out his phone and texted her his response. And she said, “Why are you texting me? We are standing right, facing each other.” And her son said, “it's just easier this way.” That is an example of what ubiquitous technology does: it makes people uncomfortable with direct interpersonal interaction. A lot of young people don’t like talking on the phone because it’s so immediate. Exchanging text messages provides a reduction in risk, whereas a direct, face-to-face interaction involves a kind of vulnerability. This is the type of technological mediation that's going on with social media, which is entirely different than what I was trying to get at with “The Truth of Fact, the Truth of Feeling.” I'm not sure that that story has much to say about what is happening with social media right now.

[transition to Q&A with general audience]

Audience Question: Is the use of AI tools the opposite of developing critical thinking skills, where people rely on them and then sort of deactivate their own critical thinking, or is there also a positive use for AI?

Ted Chiang: Generative AI is absolutely an enemy of critical thinking. It’s a way to offload all your thinking onto a machine. This can be politically advantageous for certain factions, and it is very easy to sell people on something that promises to reduce effort. There's this alignment between politics and capitalism when it comes to generative AI. It's an attractive product, and the capitalist motivation was probably what initially drove companies in a certain direction; they're looking for something that they can sell to millions of people. And then certain political actors realized, actually, this is really great, this works well for us, so we're going to get behind this too.

I think machine learning does have a value to researchers, but it can only be effectively used by experts; it’s not something which is of enormous value to non-experts, in the same way that an electron microscope is not useful for most people. There's this tension because of the capitalist motivations behind companies selling AI. You're not going to make a billion dollars selling electron microscopes, because the market simply isn't big enough. To make billions of dollars, or to justify the billions you spent to develop generative AI, you need to market something to a very wide audience. So the really

positive applications of AI are not profitable, because they are ones that require work and effort and critical thinking, and you can't sell that to a lot of people.

Audience Question: Your novella “The Lifecycle of Software Objects” was published in 2010, prior to the to the rise of large language models, and explores the coming of age of AI. But it’s not a dystopian science fiction imaginary, where AI is going to take over the world and humans are going to go extinct. Instead you explore AIs as social companions for humans. What prompted you to back then to explore this aspect of AI—as social entities coming of age with help from human beings rather than as something to fear?

Ted Chiang: I was writing in response to this trope in science fiction, where you basically flip a switch on a robot and it goes from being completely inert to suddenly being a conscious, competent, and obedient servant, a fully-fledged person that came out of nothing. That just seemed really weird to me. I didn't see how that could work, because people don't come out of nothing; to make a competent and trustworthy servant out of a human being, you need at least twenty years of effort, maybe more. At one level the story is an attempt to explore the question of, how do you make a person? I believe there's good reason to think that the process of making a person is irreducibly complex, and the way we currently make people might be close to optimal. Maybe you can't make a person in less than twenty years; maybe you can't make someone who is competent and trustworthy in a shorter amount of time. If that's the case, if you’re dealing with entities which start out as helpless and require twenty years of care in order to become competent, how can you possibly get someone to invest that much effort? I don't know that you can get good results by paying someone to do that. It seems like the only way to get someone to spend twenty years making a person is if they have an emotional involvement in that person. So that might also be irreducible, something you can't get around. It might be that you will always need something analogous to child rearing.

Audience Question: What potential does AI offer when it comes to speaking with others, in terms of animal communication systems and the like?

Ted Chiang: Some people are making very optimistic claims about the power of AI to decipher animal communication. I'm skeptical about these claims. Machine learning could be a useful tool for detecting patterns, either in an acoustic modality or some other modality, that we might otherwise ignore or overlook. That is certainly a possibility. But I don't think that’s the major obstacle in animal communication. The biggest difficulty in animal communication is the matter of the radically different perspective. What are whales trying to communicate to each other? I don't think that a better spectrographic analysis of whale song is going to tell you what they are trying to say. I think you need a better sense of what whale culture is, what it means to be a whale. Linguists will tell you that language is pretty much inseparable from culture. Understanding whale communication, to the extent that it constitutes a language, is also inseparable from understanding what it is to be a whale. Acoustic analysis could be helpful, but it's not going to be helpful unless you first have a sense of what it is that they are trying to

communicate. It's not going to be a simple translation; there's not going to be a clear one-to-one mapping to human concepts.

Ted Chiang's fiction has won four Nebula awards, four Hugo awards, six Locus awards, and the PEN Malamud Award. His first collection *Stories of Your Life and Others* has been translated into twenty-one languages, and the title story was the basis for the Oscar-nominated film *Arrival*. His second collection *Exhalation* was listed as one of the Top Ten Books of 2019 by *The New York Times* and included in President Barack Obama's 2019 reading list. Chiang has written about AI for *The New Yorker* and in 2023 he was named one of *Time Magazine's* 100 Most Influential People in AI.

References

Chiang, Ted. “The Lifecycle of Software Objects.” 2010. *Exhalation*, Picador, 2020, pp. 62-172.

---. “Story of Your Life.” 1998. *Arrival (Stories of Your Life MTI)*, Knopf Doubleday, 2016, pp. 88-141.

---. “The Truth of Fact, the Truth of Feeling.” 2013. *Exhalation*, Picador, 2020, pp. 185-230.

---. “Why A.I. Isn’t Going to Make Art.” *The New Yorker*, 31 Aug. 2024. Condé Nast, https://www.newyorker.com/culture/the-weekend-essay/why-ai-isnt-going-to-make-art. Accessed 23 June 2025.

Mountfort, Paul. “The *I Ching* and Philip K. Dick’s *The Man in the High Castle*.” *Science Fiction Studies*, vol. 43, no. 2, 2016, pp. 287-309. *JSTOR*, https://doi.org/10.5621/sciefictstud.43.2.0287. Accessed 23 June 2025.

Rees, Tobias. Personal correspondence, n. dat.

Williams, Paul. “The Author and the Oracle.” *PKDS Newsletter* 25, Dec. 1990, pp. 1-10.