

Center for Humane Technology | *Your Undivided Attention* Podcast
[2023 Ask Us Anything](#)

- Tristan Harris: Welcome to *Your Undivided Attention*. I'm Tristan Harris.
- Aza Raskin: And I'm Aza Raskin. We are so excited to be doing another Ask Us Anything episode, the first one since late last year. And honestly what a year it's been.
- Tristan Harris: Aza and I had no idea when we were preparing the first AI Dilemma talk, we were getting the calls from the Oppenheimers and feeling the precarity of this moment, that we would be meeting that same year with President Biden and that Aza would be there for the signing of the White House's historic executive order on AI safety just a few weeks ago with the president. Or that we'd be with Senator Chuck Schumer for the AI Insight Forum with the most powerful CEOs of all the AI lab leaders.
- Aza Raskin: Basically, we and our whole CHT team has been in non-stop sprint mode for the past year and you've shown up alongside of us. Our original AI dilemma talk has been watched online by more than 3 million people and became the basis for California Governor Gavin Newsom's executive order on AI.
- Tristan Harris: And more people than ever have found this podcast. We've hit more than 22 million downloads and we're close to hitting our hundredth episode. So when we put in a call for questions for this Ask Us Anything episode, you really came through.
- Aza Raskin: Now while we won't be able to answer all of your questions, our team did read every single one. And from us and from our team, we just really wanted to say thank you. So without further ado, let's get started with the first question.
- Loughlin: Hi Tristan. Hi Aza. My name is Loughlin. I'm a software developer based in Dublin, Ireland. I've been using ChatGPT as my peer programmer slash programming buddy over the past six months to the point where I've come to the realization that if I were to stop using it or if OpenAI turn off the tap and ChatGPT is no longer available to me as a tool to assist me in my day-to-day work, I think my employer, my colleagues and my clients would notice the reduction in my output very quickly.
- I'm curious to know what you guys think of how quickly AI, particularly in software development, has become part of the process to the point where people like me can't really remove it now without its absence affecting directly our productivity and the quality of our work. What are your thoughts on this?
- Tristan Harris: Hey, Loughlin. This is related to our three laws of technology. And if you remember law number two, that if the technology confers power, if AI confers power, it starts a race. The programming teams that adopt ChatGPT to accelerate their programming start out-competing the programming teams that don't adopt ChatGPT. And this is what we mean by the problem of

entanglement. That once AI becomes entangled with work processes, with productivity, with businesses, it's really, really hard to untangle.

Aza Raskin: What you are feeling is the very beginning of entanglement of AI into our societies, into our companies, into our GDP. And so the very fact that you are having trouble disentangling yourself is the signal that we are running out of time.

Tristan Harris: It's also important to say that we're not trying to vilify all AI use cases. And AI that accelerates some productivity for programming teams can be a really good thing. And again, this isn't existential at the early stages. Right now programmers are just getting benefits. To cite our friend Max Tegmark at Future of Life Institute will say, "It's a race to the cliff and the view looks better and better right before you lose control, right up until the point that you hit the cliff."

And so we are actually handing over the control of our code, our companies, our decision makers, our managers, our CEOs, our boards to increasingly be decided by automated AI systems. And that's one of the ways that we kind of slow walk our way to losing control.

Aza Raskin: I mean what I hear you saying Loughlin is that you know you don't want to participate in the race and yet you are forced to, and actually it's really nice to use these tools because it makes you more productive and therefore what should you do? This is actually very similar to climate, where it's actually really helpful to be able to get on planes to go see people, to do business. It makes things more efficient and yet we all know we should be flying less. So then the question is what do you do as an individual?

And the realization is of course that all of these things are not individual action. They are collective action. They are coordination problems. And so the solutions we should reach for are coordination solutions. Now, that's not going to feel very empowering, I think to you as an individual sitting inside of a company because who are you supposed to coordinate with? But if we can get enough people to all see where this goes at the same time, if we communicate clearly about the consequences of letting the race continue, that's where agency comes in.

Tristan Harris: I think the real question that Loughlin's asking is just what do I do as an individual? And we're going to get to that in some of the later questions.

Jack: My concern is about AI's impact on our capacity to think, specifically the fact that the act of writing is simultaneously deeply challenging for most of us, and also fundamental expression of our thoughts. By outsourcing the process of ideation, outlining, drafting, etc. to an LLM, should we be concerned about the simultaneous diminishment of our own ability to think? Apart from the issues relating to academic integrity, which is a whole other minefield, what are your

thoughts regarding how Ai should and should not be used within an educational context?

Should we be thinking about drawing a line regarding what cognitive work we're prepared to offload to an algorithm and what work we need to keep to preserve our own humanity? Thanks very much.

Aza Raskin: Thanks so much for this question, Jack. I think what you're pointing at is that technology, if we don't use it appropriately, can end up atrophying our muscles in various ways. The classic example is relying on GPS means you are less able to navigate when you don't have access to a GPS. So when we think about the use of this technology, we have to make the distinction between uses of the technology that strengthen our muscles and uses that atrophy our muscles.

Tristan Harris: Aza, you have an example of this from a car you bought recently.

Aza Raskin: That's right. This is my Crosstrek, Subaru, and it has a really interesting driving assist feature and it's not full self-driving. It notices if you're drifting out of a lane and it will steer you back into your lane, but it'll only do it once. So what it's saying is, "I've helped correct you when you've made a mistake, but I'm handing agency as soon as I can back to the human." And so it's strengthening your muscle while also keeping you safe.

Tristan Harris: This gets to a theme we've actually talked about even through our work on social media and our earlier podcast episode a few years ago with Yuval Harari. We talked about humane technology should be in a relationship with humanity that is of lifelong human development. What that means is that it helps us develop our cognitive, emotional, relational, interior and exterior skills. Imagine a human and it's growing and it's being taught lessons about morality. It's being taught lessons about relationality.

How do humans become more and more empathetic? How do we become more aware of our environment? These are lines of human development and you can think of technology that is flexing our muscles versus atrophying us as being in a relationship that is strengthening and developing ourselves and yet making us less and less reliant on that technology over time. Just like a good teacher doesn't just tell you the answer to the question and make you more reliant on the teacher.

The best kind of teacher or the best kind of parent might help you, but in a way that makes you less and less reliant on the parent or teacher. And this is the ideal relationship between humanity and technology.

Aza Raskin: Another way of saying makes you more and more dependent is to say, becoming addicted.

Center for Humane Technology | *Your Undivided Attention* Podcast
[2023 Ask Us Anything](#)

Tristan Harris: In our earlier work on social media, we used to reference the phrase human downgrading. That while we've been upgrading the machines, we've been downgrading humans. We've been diminishing our cognitive capacities, making us more addicted, distracted, narcissistic, and polarized. And what we really want is as we upgrade the machines, we are upgrading humans. We are making humanity more resilient, more capable, more empathetic, more aware of its shadow. An assistant shadow that enables us to become a more whole species that loves itself more.

Yasmin: Hi, this is Yasmin from India. It's a pretty basic question. But I just read that the BBC like Reuters and some other organizations have implemented something that stops ChatGPT from accessing their data for copyright issues, I believe. My concern here is if more and more formal organizations do this, then won't ChatGPT sources just be conspiracy theorists and all that, and is that not dangerous? And secondly, does ChatGPT actually curate their data sources right now? So that's my question.

Does it make sense for organizations like BBC to do what they've done? Thank you.

Tristan Harris: What I hear you asking Yasmin is if we take out the BBC and Reuters and the New York Times from the training data that ChatGPT is trained on, then are we just going to be left with Reddit and the conspiracy theorists and the moral outrage economy that's floating around Twitter and training the entire internet on basically the worst of human behavior if we take out all "the good content?" And you're right that the BBC and the New York Times and CNN and Reuters are now preventing web crawlers from accessing their copyrighted material.

And you're asking the question, what's going to be left? What are the rest of the data sources that ChatGPT is trained on? And the real answer is we don't know because the real problem is that these big AI labs are not disclosing all the data sets that they are training on. And there are proposals and regulation and law that are saying that the companies should have to disclose all the data that they're trained on. Because if we knew what portion of this material was based on those sources, we would have a sense of what is the remaining material that it will be trained on.

Aza Raskin: I think it's really important to note that we're setting up almost a kind of false dichotomy right here. Where the only options are that all the high quality content providers close down and they don't get added to the big AI models or they just open up and get their copyrights infringed. And there are obviously going to be middle paths. For instance, there could be controls on what kind of data is forced to go into publicly available models, and those sources are then compensated for them.

Center for Humane Technology | *Your Undivided Attention* Podcast
[2023 Ask Us Anything](#)

Tristan Harris: I do appreciate what you're saying is there's a middle way. It's like do we either let them have the data or we say they can't have the data? It's like, well, they do have the data, but the providers of that data are compensated and attributed.

Aza Raskin: That's right. And have some kind of say in how the data gets used or how the model gets used because they now have some kind of skin in the game. And this of course is a policy question, so we wanted to go to our senior policy manager, Camille Carlton, for a response.

Camille Carlton: So part of the work that we do at the Center for Humane Technology is actually about creating a policy that incentivizes the design of these products to be more humane from the beginning. And we actually know that this approach to pre-training data is doable and can significantly change the output of these models, but most companies don't want to take this approach because it's more expensive. So what we look for within policy is how do we incentivize this?

How do we say, "Okay, it is going to be cheaper for you to be more responsible and how you approach data and how you build these models than it is for you to be irresponsible." Which is the status quo at the moment.

Tristan Harris: And one of the ways you make it more expensive to just extract all this data are the lawsuits that are now happening from authors and from news publishers that are saying that if you take this, we are going to sue you. And that makes it more expensive to take the data irresponsibly. Cool. Let's keep going.

Aza Raskin: This next listener is asking a question about the podcast episode we did on Senator Chuck Schumer's Insight Forum on AI.

Joan: Hey, Aza, and Tristan. This is Joan from Cambridge, Ontario, Canada. I just listened to your episode that was about the forum that you participated in with all of the CEOs of all the tech companies. And something that I thought was really interesting was when you were talking about show me your incentive and I'll show you the outcome. And how at the end what I heard was we need cleaner thinking in the room that is without incentive. I am actually a physician, I work in healthcare, I see humans and how they make choices all the time on a personal level.

And we all are going to only operate out of incentive even for yourselves personally, the Humane Tech organization that you've founded. And what would you suggest to these CEOs and folks who are working in the AI tech space of what is an incentive that still ties back to safety, security, belonging, the things that we actually all use to motivate us in our lives? What have you found to be effective incentives that are reaching for the kind of outcome that you're hoping for?

Center for Humane Technology | *Your Undivided Attention* Podcast
[2023 Ask Us Anything](#)

Tristan Harris: Hey, Joan. This is a great question. Something Aza and I talk a lot about is the Upton Sinclair line that you can't get someone to question something that their salary depends on them not understanding. And so when I hear you say we need cleaner thinking in the room that is without incentive. One of the problems is that the truth of the harms of AI are not incentivized by the CEOs who are building it to think about a lot. In fact, Aza, you've had some recent experiences with this.

Aza Raskin: Recently I was having a conversation with someone who's not one of the top CEOs, but is often in conversation with the Bezos's and the Elons. And they said something really interesting and self-aware, which was that they couldn't point out the bad logic of the CEOs because if they really pointed it out, they wouldn't be invited back. They'd be seen as antagonistic. And so they have to go through this whole calculus of what is the maximum thing I can say that still gets me invited back into the room.

But being self-aware that when they do this, they will never fully correct the bad logic and hence the CEOs just keep doing what they're doing. And that means the CEOs actually never really have truth told to them.

Tristan Harris: There's another thing going on here, which is that people who talk about the positive and optimistic use cases of AI often do get invited back because people who speak in those terms make you feel good, and people who talk about the risks make you feel bad. And so if you just think about the incentives of conversation, those of us who focus on how things could go wrong, don't have an incentive to talk about that so publicly all the time in certain contexts. But I do want to give people optimism that it is possible to change the incentives. In history, we did this with drug companies.

Think about the world before the Food and Drug Administration came along. It was a race to who could basically invent the best snake oil. There were real drugs that actually really helped people, but those drugs were sold alongside people who were just trying to make money. And so the incentives were, who's just best at selling snake oil? And then we created the FDA and phase one, phase two and phase three trials, which basically changed the race from a race to profit to a race to safety, to things that would be safe and effective. That had to make it through those three gates of phase one through phase three trials.

And you can think about AI right now as a pre-FDA world where we're just shipping products to beat the narrow game for market dominance. And the FDA we all know has many problems and failure modes with how it operates. But I think it's an example of you live in a world without institutions and there's no incentive to race for safety then you create institutions that change and bend the incentives towards safety.

Center for Humane Technology | *Your Undivided Attention* Podcast
[2023 Ask Us Anything](#)

Aza Raskin: Another way of changing the incentives is liability. And to give a very human scale model of this, if you're a parent and you bring your kid to the supermarket and they break something. You break it, you buy it. You are liable. You have to pay for the thing you break. And you could imagine in the AI space, if you train a model and then someone uses that model to break something big in the world, or if you train it and they break it, you still have to buy it. And that changes how fast companies would deploy their models into the world because when there are real world consequences, there's financial consequences for them.

Tristan Harris: So laws for example, that created strict liability for AI model developers would actually change the incentives.

Aza Raskin: And finally, instead of just incentives that are sticks, what are incentives that are carrots? And I want to refer people back to the Mind the Perception Gap episode with Dan Vallone, where we talk about perception gaps and bridge rank. And the basic idea here is right now social media sort of rewards you for saying the thing that engages people's nervous systems, gets them to react. And that's of course, as we talk about all the time, is the thing that will polarize people.

It's the angry stuff. What happens though, if you could create an incentive that let you in the sort of nonviolent communication sense, accurately model the other side? Do you really understand what the other side believes? And if you could then promote content that lets people accurately see the other side that starts to knit society back together again. Now, this was a fast explanation of a more complex topic, so I really recommend people go back and listen to the episode with Dan Vallone. And by the way, any of the episodes we mention in this episode, we will include in our show notes.

So they'll be one click away. This next question comes to us from a listener named Kess. They say, "I regularly use AI tools such as Midjourney, and ChatGPT. And your insights about the risks of uncontrolled AI growth has really got me thinking. I work as a pre-sales engineer at a cloud service provider that offers various AI-related services. I'm all about doing right by my customers, but I'm worried about contributing to the growth of unregulated AI. So I'm reaching out to you for advice. How can I do my job well without adding to the problem of unregulated AI?"

Any tips for dealing with clients and promoting AI services in an ethical way? My main focus is on our customer's best interests, and I believe your advice can steer me right. Your wisdom could help me advocate for good AI practices within my company too."

Tristan Harris: We get this question a lot, which is from people who work inside of some part of the ecosystem and ask, "Well, how can I get my one company to have better ethical practices?" And what this really is getting back to is the second law of technology, we've mentioned in our AI Dilemma talk. Which is recognizing these

arms race dynamics at play, that if a technology confers power, it's going to start a race. And those who move faster in that race, who offer the most unregulated form of technology, generally outcompete those who are locking it down and trying to do things in a safer, more restricted, more values-driven way.

If I have a value, that means I'm sacrificing something. That means I'm tying one hand behind my back and saying, I'm not going to release things after I do safety checks. Well, that just means the companies that don't do the safety checks and shortcut their way to the market, they start outcompeting you. So this might make you feel powerless.

Aza Raskin:

And so this is where you have to start reaching for a different tool than the one that's typically in your toolbox. Which is to say that you have to coordinate not within your company, but between your companies of the same sort of size and stature. So that is you have to invite all the other competing companies to come up with a norm so that you can bind the race.

We're trying to come up with some terms for this, but this is the feeling that you have when you see this kind of race dynamics is sort of cower and be like, "I'm powerless." But it's really... it's to go from cower to coordinate. And Tristan, I think you should talk about what you did from inside of Google.

Tristan Harris:

Well, I tried to change Google from the inside for several years, and Google is a unique case because in the attention economy, it doesn't just build one app like YouTube. It actually hosts Android and the app store. And if you can change android and the app store, you can coordinate the incentives. You can change the incentives that all apps have to play inside of. But ultimately, I was unable to change the incentives from within just Google because I was trapped in one company and I had to struggle with what does that mean?

Well, I ultimately decided to leave and say, how do we create a public conversation that enables the whole world to coordinate around a different set of incentives by recognizing that we're currently headed towards bad incentives? So one thing you could do Kess is say, who are the other cloud service providers that are offering the same AI-related services? What if you invited all of them to a conversation or to a screening of The AI Dilemma. And you got them all to see the problem at the same time with clarity.

And said, what would be common practices, common rules, common laws that would essentially enable all of us to coordinate to a better shared outcome? And obviously it's going to be easy to defect as one company on that agreement. And ultimately you can say, well, what if we collectively advocated for laws that enable all of us to do the righter, harder thing?

Speaker 9:

Hey, Tristan and Aza, the race dynamic among tech companies is clear, but it also feels like there's this internal race dynamic among those of us trying to keep

abreast of the latest AI developments, research, policy, debates, etc. So my question is how are you both managing this dynamic personally while also leaving time for looking at work across disciplines or setting aside the time that's necessary for the kind of deep thinking required to address these big questions?

Aza Raskin:

First, I really just want to start by validating that feeling. It is very hard and actually just impossible to truly keep up. And it's actually also hard for, I'll just speak personally, my mental health for constantly being on Twitter, seeing the latest developments and seeing how fast the bits are flowing out into the world. And it's not just me and it's not just you. Jack Clark, the co-founder of Anthropic, we quoted him in The AI dilemma. I just want to read that quote here, which is, "Tracking progress is getting increasingly hard because progress is accelerating.

This progress is unlocking things critical to economic and national security. And if you don't skim papers each day, you will miss important trends that your rivals will notice and exploit. And then after hearing that, you have to have the next thought, which is today is the slowest that it will ever be." So let's just start by saying I validate that feeling and that concern.

Tristan Harris:

And I think what this also speaks to is in terms of making change with an exponential curve, we have to put on our Wayne Gretzky goggles and skate to where the puck is going to be. We have to see where the world is going and actually take actions that are not about meeting the world right now as it is with GPT-4, but skate to where GPT-5 is going to be, to where Gemini is going to be, Google's next model. And ask, what are the constraints or the guardrails that we're going to need looking into the future?

So for us, I want to mirror what Aza said, I struggle with this a lot. If I go away for two days, I have a litany of things that I have to catch up and read, and there's only so many hours in a day and we're a small set of human beings that are trying to track all this.

Aza Raskin:

The honest answer is to force ourselves to have deep time to look at the principles and these sorts of long-term trends. Tristan and I often have to force ourselves to give a presentation. We have to give ourselves a deadline where we know we're going to be going to a high-stakes situation in front of some kind of powerful audience to make us take the time to stop and do that deep thinking. And then also, it's not just the two of us. We're very lucky in that we have a team, and so we have a Slack channel.

Actually a couple where multiple team members are all posting in the things they find most important. And so maybe for you it's consider forming a study group and a Slack channel where you're inviting all of your friends in to help do that curation collectively.

Center for Humane Technology | *Your Undivided Attention* Podcast
[2023 Ask Us Anything](#)

Tristan Harris: I also wanted to recommend former podcast guest Marietje Schaake newsletter called Tech Policy Watch, and you can listen to our episode with her and you can sign up for our newsletter called The Catalyst, which we just relaunched and are starting to put some links in there. I just also want to say this is why in our work we focus on systems thinking, looking at the incentives of the system versus looking at what AI does right now or not. Because when you understand the incentives of a system, you understand where it's going.

Aza Raskin: All right, the next question was submitted anonymously. It asks, "If true legislative regulation to govern and/or slow down the deployment of AI is a slim chance or won't move fast enough, you've said before that legal action, judicial is the only real lever we have to pull. What are your thoughts on the legal action that has been taken to date? For example, the class action lawsuits by Clarkson Law Firm against OpenAI, Google and Cigna?"

Tristan Harris: I don't know about the specific case with the Clarkson Law Firm against OpenAI and Cigna, but I can speak to the importance of lawsuits at transforming industries. And I want to give listeners some good news, which is that recently 41 State Attorneys General sued Meta for intentionally addicting children with their products with Facebook and Instagram, and this has taken years to happen. So the listener who wrote this is correct, that legal action does take time, and we do not have that time with AI.

Which is why we always reinforce the quote by E. O. Wilson that the fundamental problem of humanity is we have paleolithic brains, medieval institutions, aka laws and governance, and accelerating, godlike technology. And so what we're going to need is to upgrade those medieval institutions to move as fast as the godlike technology. Think of it like your immune system has to move as fast as the evolutionary rate of the virus, the virus being the mutations of technology.

Now, that's not impossible, but you probably want to slow down the mutation of the virus, which is why things like slowing down AI development to a pace where we can get this right are worth advocating for.

Aza Raskin: If we're sort of looking just a little bit into the future, I think a really interesting idea is to use the power of AI's cognitive labor to strengthen all of our laws. Because what will happen if our institutions don't use the asymmetric power of AI first is that every rational actor that has access to money and compute will use all of the open source models to find every possible loophole in existing law.

And the only way to protect against that is by doing something sort of like Alpha Law, where you have institutions point the AIs first at all of the laws to discover every loophole to patch them up before the bad guys get there.

Center for Humane Technology | *Your Undivided Attention* Podcast
[2023 Ask Us Anything](#)

Tristan Harris: Imagine if GPT-5, which is more powerful than GPT-4, was used to identify all the loopholes in your legal system where people could get away with fraud, murder, crime, etc. And you're using the more advanced AI to patch all the loopholes in the illegal system before the less powerful AIs can catch up and exploit all those loopholes. But you always want to make sure that the more advanced AI is used for defense rather than enabling the more advanced AI to offensively hack the system.

Aza Raskin: That's right, and this is actually one of the many challenges with companies, especially OpenAI, racing to deploy their models to the public as quickly as possible because it doesn't give the institutions the time to create defenses.

Tristan Harris: One of the key challenges to the law is the gap between the spirit of the law and the letter of the law because the technology keeps evolving. We asked Casey Mock, who's CHT's Chief Policy and Public Affairs Manager to comment on this.

Casey Mock: To be clear, it's a false choice between a legislative or a judicial solution on AI. These two branches of government have a symbiotic relationship, particularly on a fast-moving technology like this. There'll be fact patterns that legislative drafters can't imagine in advance, and it's a fool's errand to try to write the perfect comprehensive bill on a complex issue like generative AI. Courts keep the law adaptable, let it grow organically, keep it future-proof. Without that adaptability, legislation becomes like a new car; it'll lose its value right when you drive it off the lot.

Thinking about these products in much the same way as a car or an airplane or even a kid's car seat, where the manufacturer has a clear duty of care to design and build a safe product or else they can be sued, is crucial to changing the financial incentives of the companies building this technology.

Manuel: Given the stage that we are at, with a need for a global conversation about AI, how might I contribute even if I'm not directly involved in tech? What do you believe is required from an individuals not directly involved in tech or legislation to be part of this essential dialogue? Speak soon, Manuel.

Aza Raskin: Manuel, thank you so much for this question. This is another one of those questions that we get time and time again, which is what can I do as an individual, especially if I'm not at one of the companies? And of course we're in early days of this conversation and before there's a big public awakening and there will be a mass movement in the future, but there isn't one yet. It makes it feel like there's not much to do, but these first voices and being brave, it really counts.

So we'll give a couple thoughts on what you can do. The first is contribute to clarity, right? There's a lot of noise out there both on, AI will kill us all, on AI will save us all. And we find it very helpful because everyone wants to know which

way is this going to go? Are we going to get the promise or we're going to get the peril is to focus people's minds back on incentives and the race and to powerfully communicate that in your community and to your friends and to your co-workers.

Tristan Harris: And the more we make that exceptionally clear why we can predict the future if we can look at the incentives just we like we did with social media. That I think helps people move from the cacophony and the noise and the confusion into the clarity that there are major risks ahead. And just to say that unlike social media, AI is really moldable. It's really moldable by what people think about it. I would never have believed that we can influence the conversation as much as Aza and I have been able to. Except for the fact that people's opinions on it just really hadn't been developed because people need to get educated about it.

So if you could help educate communities, if you can bring your friends, family, co-workers together and host a watch party of *The AI Dilemma*. It might sound ridiculous, but we've seen many people do this and it's led to some local community actions. You can also organize or attend a protest. We've seen small groups get together to let AI labs know that we really need guardrails. There's groups like *Pause AI* that are helping citizens create visible media moments and that have drawn some outsized attentions to problems like the race to AGI or too early releasing open source models.

Aza Raskin: The problem is individual action is insufficient. It has to be coordination. You get coordination by doing communication, but communication that's local only has local effect, and this is a global race, so you have to do global communication. That's the solution, but the act of doing it locally, if enough people and many different places all do it locally, that can have a more global effect. There's a fairly famous example of, it's actually one of our friends who got a billboard outside of Facebook's office that was calling out all of the climate misinformation being amplified on Facebook. And it did something.

And so there's something interesting to think about. What are these non-obvious ways to communicate beyond just you and your friends but to your community? And if it's from your community up to your city and beyond. It's more that question that we should be asking ourselves than us being able to give you any specific answers of what to do.

Tristan Harris: I think one of the things that listeners need to think about for 2024 is how AI will be impacting the 2024 elections around the world. I believe it's the case that something like 2 billion people are having elections this coming year, and that is going to be a complete mess with generative AI. So we all need to be vigilant about how AI is going to be affecting elections, and we hope to cover this a lot more in the coming year. It makes these issues all the more urgent. So please focus attention on that and how to educate people around you, your communities to try to inoculate them about the effects.

Center for Humane Technology | *Your Undivided Attention* Podcast
[2023 Ask Us Anything](#)

Aza Raskin: That's right, Tristan. I think it's the United States, India, Brazil. It's all honestly the major democracies. We'll turn to that in future episodes and really just wanted to say how much we enjoyed answering your questions and how grateful we are to be on this journey with you. And just one of the things I wanted to say, Tristan, about what is it like to show up to do this work? Honestly, often it's pretty bleak and it doesn't seem like there are many paths to good outcomes. But for me personally, that means that I have found the time that I spend with friends.

There's a kind of preciousness that I get to feel that makes me more grateful every single day. And I just wanted to offer that to all of our listeners because it's important to find the right conjugation, the way to feel as we look at the hard truths of where we may go.

Tristan Harris: Lastly, I want people to think about what is a stable state between AI and humanity? Do we just keep racing and scaling bigger and bigger models that are capable of more and more magic powers that leave society more and more overwhelmed and unprepared? Or is there some stable state with certain levels and amounts of AI, certain kinds of magic powers that are restricted to certain domains with certain people that have the wisdom to wield those powers? How do we get this right by linking power and wisdom? I think we need a lot more people envisioning positive stable states with AI and humanity.

When we interviewed a lot of the people who are inside the AI labs, they often sort of don't have a positive stable state. Which says, then why are we continuing to race and scale as fast as possible? And I think we urgently need a project of many people thinking about, from many diverse backgrounds and cultural ideas and frameworks, thinking about what this positive vision looks like. Because when you tell people, just don't do something. That's not as easy as saying, well, let's walk over here instead. And I'd like to encourage many more people to be thinking about what world we actually want to live in.

Your Undivided Attention is produced by the Center for Humane Technology. A nonprofit working to catalyze a humane future. Our senior producer is Julia Scott. Kirsten McMurray and Sara McCrea are our associate producers. Sasha Fegan is our executive producer. Mixing on this episode by Jeff Sudekin. Original music and sound design by Ryan and Hays Holladay. And a special thanks to the whole Center for Humane Technology team for making this podcast possible.

You can find show notes, transcripts, and much more at humanetech.com. If you liked the podcast, we'd be grateful if you could rate it on Apple Podcast because it helps other people find the show.

And if you made it all the way here, let me give one more thank you to you for giving us your undivided attention.