How to accelerate digital transformation with ethical AI

Noelle Bloomfield:

Welcome to this Gloat webinar on how to harness ethical AI to fast track digital transformation. We're really excited to bring this audience together and to have some incredible guests with us.

I'm Noelle Bloomfield. I'm the Director of Product Marketing at Gloat. We have a fantastic talk here in store today, and as we go through the event, please drop in your questions. We'll handle some chat discussions as we go. We wanna keep this interactive and lively and make sure you get your questions answered. Before we kick things off, a quick reminder that all of our previous Gloat Learn webinar sessions are available at gloat.com/resources. This webinar will be recorded and available there, so if you have to hop off or wanna relive the magic, you'll be able to find it and you'll be able to find other stories with incredible guests and some of our other partners like Novartis and Schneider Electric. If you haven't already, please subscribe to our newsletter, Gloat's Humble Opinion, for all of the industry's latest and greatest.

So now let's kick things off. I'm excited to have some incredible guests with us today. We have Paula Goldman, the Chief Ethical and Humane Use Officer at Salesforce; Richard Benjamins, the Chief AI and Data Strategist at Telefonica; and Odia Kagan, Partner and Chair of GDPR, Compliance, and International Privacy at Fox Rothchild. Thank you for joining us today. I'll give you each a quick minute to say hello.

Paula Goldman:

Hi, I'm excited to be here today. There are so many important things to talk about and obviously so many important things going on at Gloat as well. I don't know if there's anything else you want us to talk about or background?

Noelle Bloomfield:

Where you're joining from?

Paula Goldman:

I'm in the Bay Area and I am deeply immersed in AI ethics, so really excited for today's conversation.

Noelle Bloomfield:

Awesome. Richard. Hello from you.

Richard Benjamins:

Yes. Hello, everybody. I'm originally from the Netherlands. I saw there was somebody from the Netherlands as well, so hello. But I'm working from Madrid for many years. And actually it's nice and sunny and hot, it's 24 degrees Celsius. So that's good. I work a lot on AI ethics as well, but also on AI for good. So it's not only avoiding the bad, but also creating a positive impact.

Noelle Bloomfield:

Awesome. And Odia, hello to you.

Odia Kagan:

Hi. I'm Odia Kagan. I'm a lawyer, the lawyer in the room here. So I come at it from, how can you incorporate AI, you know, systems and use them in your company in a way that is compliant. So my big focus is data privacy, both on the US side and the EU side. So, the interplay between AI and data protection, which is a big one, but also increasingly AI compliance with other laws, both new and existing. Those of you that tuned into the statements by the FTC and Lina Khan, the Chair, over the past few days, there are a lot of existing laws that the FTC is ready to apply to AI. So that's where I come in. And I also do a lot

of these updates. I actually just updated a White House press release from this morning. So those of you that are interested in these updates, feel free to connect with me on LinkedIn,

Noelle Bloomfield:

I would be happy to hear your feedback also. Love it. Well, welcome Paula, Richard, and Odia. We're so excited to have you here. So let's kick things off with just a little bit of a starting overview. Obviously, AI is all over the headlines. Ethical AI has been really top of mind as we see things like Snapchat's new AI feature, ChatGPT, all of the news with AI and music. So what do you think about the evolution and recent advancements in AI and kind of what's your overarching perspective on it? Richard, maybe let's start with you.

Richard Benjamins:

I think the most recent things are, of course, around large language models. ChatGP T and the fact that we have now a few machines that are able to hold a conversation that seems to be meaningful in many different aspects. I see the concerns that it creates, and also a lot of what I would say myths that people are discussing everywhere and that have no fundamental understanding of those things. It is an impressive achievement. Yeah. That's the first thing I have to say. I mean I don't know if you heard about the Turing test. The Turing test, which was designed to decide whether a machine can think or not. This machine is almost passing the test. Yeah. So it's a great achievement. That's the positive part. But there are many things that people are not aware about. So they use it as a search engine, as an information retrieval tool, whether or not it’s actually predicting things. So it says a lot of things that are perfect, could be perfectly true, but they are not true. That's what they call hallucinations. So if I comment on the recent developments, what I would say is a great, impressive, achievement of artificial intelligence would do a lot of things that we don't understand, that we don't control. And very importantly, this is in the hands of very few big companies. It's not something that is in the hands of general research and universities and the academic world. So it's very niche, and I think that's also something that is attracting attention.

Noelle Bloomfield:

Absolutely. Paula, what's your take?

Paula Goldman:

Well, I would say similarly, I know the move from AI models that give predictions to those that create content, whether that's text or image or multimodal models. It's not that these are new per se, but the power and the speed with which they're advancing has surprised many, many people, including myself. It feels like every week is a year or a decade in terms of news and advancements and keeping up with the implications of those advancements is really important. I obviously come to it with the perspective of enterprise applications of AI, and while I think those are in some ways more contained there's still a number of really important issues that need to be addressed, whether that's accuracy, as Richard was just saying, or thinking, and some of these issues are not new. They're just accelerated, like thinking about bias in models or thinking about toxicity and the content that's created. These are not new but it's all the more urgent that we talk about it, that we are transparent about solutions. And the thing that makes me happy is that these questions are really part of the public dialogue as people are increasingly fascinated with these technologies. And that's really what we need, is we need kind of all hands on deck with solutions and with guardrails.

Odia Kagan:

Yeah. So I'm gonna tee off at the end of what Paula was saying, I was immediately thinking, it's like that song, “Let's get down to business.” I think that's where we are, right? Like, this has been sort of percolating and now there's a lot of public traction and excitement and buzz. But it depends, right? Like who you listen to on the doom and gloom previews, right? But in any event this is the time now when it's sort of not really early stage, but kind of early stage, right? Early stage in the sense, not in the sense of impact to life, which I think that we're feeling, but in the potential harm, right?

We're sort of at the point where we can still pivot this redirect, and recalculate like GPS systems, right? So I think now, it's time for the call to action, right? Like what Paula was saying, it's really important that there's a buzz now because of the popular opinion, especially in the US right? The sense from the public really is driving regulation. In the post Dobbs, for example, we've seen a huge uptick in regulation, self-regulation, and enforcement, things like that with respect to protection of reproductive rights and reproductive health because of the public opinion. So I think it's really a time to kind of double down on regulation that's

effective with the support of the public, with the support of the industry, which we've now seen with the meeting at the White House, with the heads of big tech companies saying, okay, you know, willing to have these things audited.

And we have standards in place or getting formulated. We have GDPR, we have the US privacy laws, we have the FTC, we have the EEOC, we have the DOJ right? We almost have the EU AI Act. We have the OECD principles. Like we have stuff, we just need to make sure that now it actually gets respected, because that will prevent a lot of the potential harm

Noelle Bloomfield:

Yeah it really is an all hands on deck situation. And I saw Rob say practices are moving faster than policy, and it's true. We've seen this eruption of AI and we're trying to catch up and make sure the guardrails are there and there is no harm done. So I think to that end, let's delve into the risks. There's a lot of concern on what the implications are and how to make sure that this is promoting ethical use and isn't creating bias or inducing harm to others. So I guess, Paula, from your perspective, what, from a corporate perspective, what do executives need to understand to deliver these kinds of AI initiatives successfully?

Paula Goldman:

There's not one answer to that question. A lot of it depends on what you are using AI for and to what end. But I think there are some unifying principles about questions to ask and, and then things to think about in terms of internal governance, which I know everyone on this call could talk about. One thing I just wanted to flag was there was a question about closed captioning, which is near and dear to my heart as I also lead Product Accessibility at Salesforce. Do we have closed caption here? But so back to your question. I think there's a lot of questions here that are not new, right? Questions that one would ask as you're procuring AI and as you're thinking about how you implement AI within your own company, you would ask is there a model card or is there a data sheet? You know what's going into these models? What are the standards that are associated with these models? That's why the news today that Odia was just talking about is really important, as we think about this trend towards independent auditing. And thinking about like, these are things that we've been measuring for a long time, whether that is bias or accuracy, you can go on and on about the different types of measures.

But what's exciting is there's a sort of momentum and convergence towards thinking about not quite standardization, but convergence towards sets of metrics that we can compare across different products.

I think there's a lot to think about for executives in terms of the internal implementation, the signals that are set across the company about how to implement things responsibly. What safeguards are set up in terms of testing. You know, AI is only as good as the data that goes into it. So making sure that when AI is implemented with one's own data, first, testing the data that's going into it, continue tests around the fairness of the models themselves, and then obviously the controls about what data goes in and goes out. I think we could probably spend all day going into the nuances of each of these types of things, but the point being, these questions are by and large, not new, but they're all the more urgent because of the speed and kind of the new implications of generative models.

Noelle Bloomfield:

I think you're touching on a lot of really important questions there on like the data quality and how we're making sure the applications of this are successful and, and that companies are actually taking action for the internal practices that they're doing before putting this out into the market. So, great points. Richard, what about you? What would you flag as the hidden risks of AI that we should be watching out for?

Richard Benjamins:

I think the hidden risks are if you have a plan or you want to launch a product in the market, that you do not do in advance before you launch a certain critical thinking about the potential ethical or social impact or even societal impact of your product. Yeah. So in general product owners or business people, they think about a product, how it can get passionate users the best product, and they think about those things. They don't think so much about dual use or indirect consequences. Just a very simple example. Look at social networks. Yeah, it's a great tool to connect people and to have them better in contact. But then using this technology as a very strong recommendation engine in the end, it formats, it stimulates polarization in societies.

It can even influence democratic elections, and other processes. Not all of those things you can think of in advance. Because these are really indirect consequences, but they have a big impact. So if you're an organization and you see those things, I think you have to act. Yeah. You just can say, 'Hey, my product is legitimate. It complies with all the rules, and this is just collateral damage. 10% have damage, but 90% are enjoying it. Yeah. So if you take a human-centered approach, not a utilitarian approach, then well, 10% of a lot of people is still a lot of people. If you cause harm, then you have to take that into account.

And so this is a big thing that is happening on impact, but you also have bias. Yeah. So giving a different credit depending on the gender, having the same situations of the customers, that's a direct result of not thinking through those things in advance. Yeah. Now, regulation will force that, and self-regulation is doing that on a voluntary basis. And the process is actually, we've done pilots for a few years on the government's model, and we are about to install an internal regulation ahead of the European regulation that makes those things compulsory in all our business processes. Not only if we develop for our developers, but also if we procure, if we use internally, and even if we sell it to our enterprise customers.

Noelle Bloomfield:

Odia how would you build on these two risks that risk profiles that Paula and Richard have laid out?

Odia Kagan:

So I think from what risks are there? I think there is a universe. I mean, there's a universe of risks right? Besides the bias piece and the misinformation, right? Like, those are big things that everybody's talking about. There are also risks, like corporate risks that we've seen. We've actually seen this in the news and there's also trade secrets, misappropriation, and sharing proprietary competitive information that causes damage to your competitive advantage. Like trademark things and trademark copyright. There's a lot of now generative AI being used for photos and generating them right? And of course, data privacy issues when you are involving personal data, either in your, like now going forward in your use of the product. That's a much bigger issue right now in Europe because of the legal basis of GRDPR. But it can go into the US piece, not in the providence of the data part, which we have less of, but in the,

what is the impact of the scraping and does that build in accuracy and bias? And so on the other question that you said, like, what do executives need to do? I think if you wanna break it down into, Paula and Richard said it, but I'll kind of paraphrase it. I think the first thing that you need to do, and this is something it's more important that you do granularly, I think, you know, it seems like a no-brainer.

And I'm sure for a lot of companies it's a no-brainer, but for a lot of companies, it is to really understand your needs. We're coming to onboard something, right? What do I need this for? Like, really not the way my kids take a goodie bag from a party and it's like, oh, do I need this? No, it'd be nice. Like, yeah, whatever. That's a nice to have kind of thing. You can't treat your AI like the goodie bag from the party, right? Like, do you actually, because then the OCD moms are like, no, I don't have room for this in my drawers, right? And if you don't, you just stick it in some junk drawer. But you can't do that with AI because it's not a separate drawer that does nothing inside, it’s immersive in the company, right?

So the first thing is really understand the need, understand whether or not this product actually meets your needs. And then obviously, understand and drill down into the potential risks and see what you can do about them, right? Because we're gonna run into similar issues to what we're running into with other software because it is a piece of software. But now, you're running into an issue of there's gonna be off the shelf AI products that you buy, do you know what the risks are? Are you able to do anything about it other than not incorporate it? Because we've seen the FTC talk about, not in the context of AI, but in the context of data sharing with third parties, they’re like, we understand these are really big tech companies and you don't have any control over what the agreement says, but at minimum you need to know what it says and how it impacts you, and then you can maybe decide not to use this provider.

Noelle Bloomfield:

Yeou covered a lot of ground there, Odia. I think you touched on something that also John Pierre started touching on in the chat. How you're collecting this data and the way you're scraping data and what the inputs are is so important to the ethical use of these models and making sure that they're applied in a way that doesn't create harm. So I think to build on that, how should you know our viewers and companies are thinking about aligning AI strategy, data

strategy, and then ultimately their business strategy as they try and evolve and keep up with the digital age here.

Richard Benjamins:

I would say if you are not digital, you don't have any data, so you can't do anything. So once you have the data only then you can do things with the data and, and the later part you can use AI. So they are definitely completely interrelated. A digital transformation process in any organization needs to have a large data component apart from other IT aspects. Once you are mature with your data, you can also think about using artificial intelligence to use your data to train with the data, but also use all kinds of analytics and applications. They are intimately related to each other. And this has all kinds of consequences. If you want to do a transformation journey, then make very sure that where your digital transformation sits in your organization is not too far away from your data transformation because any problem they have, they have to escalate all to the top and then it goes down again. So bring them very close together such that they are peers. And if there is an issue, if you start with data and AI, it's all about agility. Yeah, I need to give me some data. Here it is. Try something new, it doesn't work. Give me something else. So you need to be very agile, and that means from an organizational perspective, you need to take that into account. That's one of the lessons. But apart from organizations you also have to think about the business case. What is the business case? What is the use case? You can have use cases that are huge, like transforming your business, but hey, maybe it takes you four or five years to put it into place.
So don't start there. Because organizations don't have such patience. Start with the quick wins, show the value, and then continue with the next one. So business impact is important. How much does it cost you, your journey? How do you invest it? Who is paying for it? All questions that are very important in this strategy. Do you hire people? Do you upskill people? Do you promote people towards this strategy?. And how do you use this technology across the organization? So it's not in a small department of, let's say very specialized data scientists or AI engineers, but in the end, the whole organization needs to be able, so you need to scale that out. And then finally, very important: you have to have a responsible business. More and more things like ESG are becoming important. It's not only business anymore, it's business, but it's also responsibility. So as an organization with this powerful technology, you have to think about what are the potential negative ethical impacts. But you also have to think about

how you can use this technology and this data for good. Yeah. To solve big societal and planetary problems.

Odia Kagan:

So I can weigh in a little bit on this generally, but I also wanna answer a couple of questions that are privacy specific. I learned a Chinese saying earlier this week that says “You pull one hair and the whole body moves.” And I think it applies here because all three of these components, right? The AI strategy, data strategy, digital transformation, right? I mean, they're all kind of either subsets or like intertwined with each other, and you mess with one, it impacts the rest, right? And sometimes it actually impacts in a way that you can't really tell, right? Because it's not like, oh, somebody hits you, you break a bone. It's more like, I don't know, endocrine function, right? It takes a month until something goes off and then like you're in a bad mood and retaining water, right?

So I think that my answer is that the really intertwined and the inter-effects are not obvious, they could manifest later. And that's why understanding AI before it goes in is really important. Cuz with AI now we've seen like, it's the epitome of this problem where the AI starts doing stuff by itself that you don't understand. So like the unintended consequences I think are interesting. There was a privacy specific question that asked how do you protect employee privacy as generative AI enables increasing sensing and predictiveness of positive and negative employee behavior? And what do you do?

So you have this situation where you have data privacy issues, right? Some of them are obvious data privacy issues that we're already using. Some of them can go into thought privacy or neuro privacy. If you're talking about predicting based on brainwaves, like god forbid, that actually gets incorporated into the wild, because that will be really interesting. But there are, you know, use cases for security, et cetera, but even mood, right? And we have case law and enforcement in Europe where companies were using this to predict likelihood to resign or effectiveness of customer support. So this is a situation where, first of all, how do you address this? First of all, you need to be mindful that this is sensitive information, right? The White House Bill of Rights called it like sensitive domains. GDPR calls it special categories of data. We call it sensitive data or sensitive information. First of all, this is sensitive because of its health, because it's mood number two, it's a vulnerable population because employees,

by definition, have an unequal bargaining power, right? Between the employer and the employee. And the third piece is more on AI, is fitness for purpose, even if this is a good thing to do, do we know that this thing that you're doing actually does it, is it gonna get false positives? Is it gonna get false negatives? Does it have built-in bias? Right? And so now what you need to do is you need to kind of look at all of these together and say, one, do I need it?

Can I do it better? Can I accomplish this in a better way? Can I do it through surveys? Can I do it through anonymous surveys? Can I do it through, I don't know, like whistleblower things? I don't know, like survey the customers, is there a better way, less invasive way to do it if, if there isn't? Rather, can I limit the use cases? Can I limit their attention? Can I put in human intervention in the end, right? Like, you have to look at it from both, like the AI piece, the labor law piece, and the privacy piece.

Paula Goldman:

There's so much here to comment on but maybe just to pick up on a couple of points that were made. First with Richard and understanding the link between data and AI. You know, we often have folks, customers coming to us very excited about AI, and we're obviously very excited about AI as well. But the first question is like, How's your data ? How's your data governance? How clean is your data? And there's a long conversation and I think it's a mistake to think of these as two different things. They need to be kind of one and the same. There's getting the data and the data governance processes really, really smooth.

Oftentimes, you're thinking about not just getting all the data, but what's the right data? What's the data you really need? And then aligning it with AI products and processes. So I think that's very important. I just wanna plus one, the notion of bringing these processes together within a company. I also think it merges very smoothly with what Odia was just talking about. Like what data is really necessary, and then how can you be, when you're talking about sensitive data, making sure that it's consented and it's transparent and people know what data are you collecting from your customers? Do they know that you're collecting this data?
How easy is it for them to understand it? And there was a comment about responsible AI and how ethics could be seen as guardrails, but could also be seen as a business enabler. And I firmly believe it has to be both. That is to say, I think we've already seen this, but it's becoming

more and more true that people really use products that they trust and that they understand. The more trust that one enables with one's customers, the more adoption we'll see, the more scale we'll see. So it's, it's very, very important to think about trust and the design and uptake of these technologies.

Noelle Bloomfield:

Absolutely. So stay unified, be transparent, and capitalize on the immense opportunity that is there if you are taking the appropriate actions. Now, I think to double click into what we got into a little bit with Odia, and I think Clayton commented on this in the chat, I wanna focus a little bit more on generative AI and talent management. Obviously, you know, this is near and dear to Gloat's heart and what we do, and as we are experimenting with and using generative AI it's important to us that we figure out ways to use this responsibly and to make sure we're not passing on legacy biases using good data. So, I'd love to get a take from this group on how you think generative AI is going to impact skills and talent management moving forward

Richard Benjamins:

Well, I can give one quick guess which is if you speak about generative AI and text, that means it can predict text. So if you say, as an employee, I would like to apply for this job, write me the best skills description that I need to be considered by the matching algorithm. And then it writes something and submits it. So I think we should not mix up generative AI, which is creating things, from traditional AI which is discrimination of deep, deep, deep learning networks or classification trees that is another kind of technology. I think it actually might make it more difficult for talent management because it's easier to look better.

Paula Goldman:

There's a lot of conversation and rightly so, about the impact of generative AI in the workplace. What I will say is that right now generative AI is like a really good assistant. It has a lot of limitations, there are problems around accuracy, there's sometimes problems around toxicity, et cetera. But I think there's no doubt that, if you read any number of headlines, that there are really serious workplace transformation issues that we're gonna need to grapple with as a society. And you see that playing out already, almost even faster than anyone might

have anticipated. I think it's really important that when we think about AI either as designers and developers of it, as consumers of it, as purchasers of it, that we think about AI that empowers people and augments people. That's the version of the vision of AI that I personally want to see in society. And I think it requires really zeroing in on those use cases that do empower individuals and not vice versa.

Odia Kagan:

I agree with all of that. I think I don't really have any additions on the upside, even though I think that there were comments in the chat on how it is gonna affect people and are people gonna be replaced by AI. I think there's been a lot of talk about it. As a lawyer, I always think about the impact and the upsides seem obvious as people start using it for their homework, et cetera. I saw this on the news today that educational websites are worried that they're not gonna get traffic because kids are going on ChatGPT. But the obvious one of thought, obviously for me is the liability side and trying to button down use of generative AI by the workforce and all the issues that have already arisen with that.

Noelle Bloomfield:

Yeah, absolutely. A lot of implications. So thankfully, it sounds like AI's not stealing our jobs right now. It is an assistant, it is a tool, it's something that has to be used carefully for sure. And I think as we see applications in the skill and talent management space, there's a lot to be figured out to make sure that doesn't introduce bias. But right now it's a generative tool. So if you put something in, you can see historically what the trends have been and get that synopsis. As Paula said there's a lot of opportunity to augment and support people moving forward and hopefully empower folks.

Richard Benjamins:

But, I must stress that this is a very debated topic. So will it create more jobs or will it kill jobs, the opinions and the studies are all over the place. So you can't say it will have no impact. You can say it will have an impact. There are studies for everything. So I haven't seen a topic that has so many conflicting opinions from great experts. Yeah.

Odia Kagan:

No wait. The part that says that it won't have an impact was that written by generative AI itself?

Richard Benjamins:

It doesn't say no impact, but it will create much more jobs than it will destroy jobs. So it's not about impact. Everybody agrees that it will have an impact, but it's the type of impact and how good or bad that is. I think that might also have an impact on talent management and workforce planning because the type of jobs shift over a short period of time, well, you might have to plan for that. But then nobody knows in what sense it will shift.

Noelle Bloomfield:

It remains to be seen, for sure. I don't think anyone has the answer just yet, but you're right. You know, that's the center of the debate. To double down on that, as an HR workforce planner, someone who's thinking about the applications here, what do you all see as the role of an executive in ensuring success with AI? We've talked about organizationally how you need to think about this, but what does an executive need to do in order to lead an AI initiative successfully?

Paula Goldman:

One thing I would add, I mean we had a pretty extensive conversation about different checks and balances and governance within a company. And the importance of data, you know, that there's no AI, there's no good AI without good data, all that kind of stuff. But the one thing I would add is that, executives have a really important role in terms of setting the culture and sort of setting the cues about how things get done responsibly and making sure that people know that it's their job, that it's everyone's job to mind kind of the ethical implications of these products. And that it's everyone's job to understand what AI is and isn't good at. As they use these tools in their day-to-day, to understand what are the implications? So I would say in addition to everything we've already said about effective governance, about the role of data, about sort of being smart and thinking about which tools are gonna be most appropriate to

different business goals, and getting really high quality tools, I would not underestimate the cultural tone setting piece of it, because I think that's extremely important.

Odia Kagan:

I can pitch in as outside counsel, but obviously I work a lot with people like chief privacy officers and data officers and people in charge of this kind of governance piece. But I think the key issue here, if we wanna drill it down to basics, is that you need to make sure that there is a system that is fit for purpose, right? And that could be whatever we're working on now, right? The NIST AI framework, the other frameworks for audits and risk management of AI. But you really need to do two things. One is what are the risks and the things that could happen so that you need to have, you know, somebody that you know is on top of developments and can inform you, right?

The other is you need to make sure that it actually happens in practice. And that, I think is probably the most important, is that you, as the AI lead champion and governance person, need to have a voice that is heard in the company. So you can know all of the things to do really well, but if you're not listened to then it won't happen. And so that you guys, Paula and Richard know probably the magic of actually making this happen in the company better than I, but I find that, you need to know all the stuff and that you have good advisors, but you need to actually make your recommendation and have it actually happen.

Richard Benjamins:

I would also say that it depends on what you call success. So if success is maximizing profits, that is a different mindset, then success is profits, but no harm. So those are, and then you can also have a mindset, like let's do good or let's, let's try to get no fines. So those are all mindsets for success. And of course, depending on the executive you are in business or marketing and probably province's profit or in compliance, then it's no fines. I think the important thing is related to what is a kind of open-mindedness. You need good advisors, but you need hard skills. Yeah, of course. You need to know enough that you cannot be fooled.

The skills you need, I think are the typical 4Cs. It's communication, it's critical thinking, it's collaboration, and coordination. This is what you need to do. But not only to be successful

with AI, you know, be successful in any kind of position or even life. So success, the definition of success is more important. And actually, I hope the definition will shift more and more from only profits to a wider perspective.

Noelle Bloomfield:

Well, Richard, I hope so too. I think profits alone without thinking about how to do that without inducing harm, is gonna be really critical as we think about this. I wanna turn to then

Odia Kagan:

The fines, and then the FTC discourages all your profits, and then the profit's gone

Noelle Bloomfield:

Yeah. I think long term, to your point, if you're doing it and inducing harm, it probably will hurt your profits long term.

Richard Benjamins:

Click. But so far there are not many examples of fines that are so high that they make the business change.

Noelle Bloomfield:

Fair enough. Well, we'll see as folks continue to experiment, hopefully they don't induce enough harm to, to change that.

Odia Kagan:

I mean, you know, those of us that have been doing it while are thinking like, what's new here? But you know, when it's like, it feels like, you know, because of the generative AI explosion, like it feels new. But to be fair, and we've seen this one, we've seen it with AI, with the guarantee injunction right on ChatGPT in Italy. We've seen it with data protection. So it's true that you could be too big to fail and too big to find. But the other pieces like injunctions

or like the requirement to delete data that was collected and processed illegally, which we have seen on both sides of the ocean, those are a big deal because you need the data, right? And then now what? And so that we've seen in other realms, we've seen a little bit of it with a guarantee of injunction that was lifted after changes were made. So I think that's also something to think about. It's not just the fines, but like the other tools in the regulatory arsenal.

Paula Goldman:

And also to remember that regulation is proceeding pretty quickly and there's lots of stuff going on all over the world that is coming into force in the next couple of years. So I think that's also something for all businesses to keep in mind, that it's not a static picture.

Noelle Bloomfield:

Yeah. I think Rochelle in the chat mentioned something on this. Curious, if you all have a take, is there a move to actually incorporate privacy and employment laws into these models to address that proactively?

Richard Benjamins:

I don't get the question. Sorry. Is there a move? Uh, here it is, yeah.

Noelle Bloomfield:

Okay. Is anyone actively working to incorporate these privacy and employment laws into the data models to make sure that they're improving data quality to align to those regulations moving forward?

Odia Kagan:

Anyone, companies or anyone? Regulators,

Noelle Bloomfield:

I don't know anyone

Odia Kagan:

I mean companies, I hope so because they need to and regulators semi yes, I think so. In the sense that the EU has been working on an AI act for a long time and it's close, and it's gonna be very impactful in the nature of the obligations that it will impose on companies. And so then of course, the magic is in the enforcement and the kind of uptake by the community. In the US we've seen something a little bit different. I mean, we've seen some legislative initiatives, there was one in California recently. There is a lot of stuff going on with the White House that's kind of trying to ramp up AI, and the promotion of ethical AI regulation.
The other US approach on this, on the general side, right? Um, not on, there are specific laws with respect to specifically in the workplace. But generally speaking, the approach until now has been, okay, like, let's first of all just explain that the laws that we already have, they also applied AI. And we've seen, there was a joint statement by the FTC, EEOC for employment, CEB, for financial institutions, and the DOJ. And there was a joint statement that basically says, Hey, like all the stuff that we are enforcing anyway in impact of data processing, impact of actions, right? Unfair, deceptive, or prohibited acts under the financial or discrimination in the employment sector, like all that, we're already enforcing it, right?

And Commissioner Bedoya said that at the IAPP conference in April and Chairman Kahn just said that basically AI is a type of activity. And so if the company doing it is interfering with competition, then we're gonna do the competition piece. And if it's information and interfering with consumer protection because the data isn't accurate, because there is bias, because there's no transparency because it’s impacting rights without giving notice or without giving choice or consent, we're coming after you with the tools that we already have. And so I think that it is really important to understand that even if we don't have new laws, and maybe there are pieces, not maybe, I think there are in the White House, Bill of Rights actually said there is a gap that needs to be filled by new laws.

But let's say in order to figure out if you are in the gap or you are in whatever, the X percent, that's already regulated, you've gotta make the analysis. And I hope that if you come to the

conclusion that you are in the other percent that can sneak by, please do the right thing anyway. Because as Paula said, one, you may get it, you may make the wrong determination, but two, that number is quickly shrinking. So I hope companies are doing it, and if they aren't, they should be.

Noelle Bloomfield:

With our final five minutes here, I wanna give a chance to hear some of your predictions and hopes for the future of ethical AI. How do we see this being used? What can we do with this? What do you predict given your experience in this space? So maybe Paula will start with you.

Paula Goldman:

Let me first start with, typically we see the pace of innovation and the pace of democratic deliberation as being on very different timetables. And, while that is definitely the case, I think this uptick in urgency from governments around the world to wrap their heads around the situation with AI right now means that I do think we will have regulatory frameworks to depend on sooner rather than later. Obviously Europe has draft language that's almost there. But I think we're seeing similar things in other jurisdictions. So I think that's one thing. The second thing is I'm very encouraged by this question of measurability and standards. And I think we'll see convergence, like any metric about performance or fairness or whatever, they're all complex.

There's trade-offs in how you measure things. But to have some sort of convergence around measurements we see that happening. I think that will get better and better. It will be a moving target. But my prediction is that there'll be a set of standard questions that people ask about systems to compare them. I also think people are becoming smart about what these new tools are good for and what they're not good for. Will there be abuse? Absolutely. Do we need to take that very, very seriously? Absolutely. But I think we're also seeing the emergence of a set of norms around how do we interact with these tools responsibly and then finally, I think this is a discipline that's been evolving for several decades, but I think we're gonna see in every company an emphasis around it because there's gonna be no way to successfully operate in the digital sphere without deep operationalization of AI ethics across the board.

Noelle Bloomfield:

Well said. Richard, what are, what are your parting words here?

Richard Benjamins:

I have three bold predictions. The first one is whether artificial general intelligence will happen. So we now have generative AI. It is not artificial general AI. This is the singularity of AI that is as good as people are. I think if it happens, it is not necessarily bad. Then, I refer to a book that I recently read, which is called A Thousand Brains by Jeff Hawkins. Very original view on how the brains work and the implications for artificial intelligence, especially artificial general intelligence, for me was a mind shifter. Second prediction is, I think the direct consequences of the use of AI like bias, opacity, explainability, autonomy, those things will be solved, as an engineering part. We will understand how it works. It will take some time, but we can do many things.

But the indirect consequences of AI that are beyond direct ones, like technology addiction or addiction to social media, anorexia, because of these strong algorithmic recommendations, and fake news, polarization of societies, those are the consequences that we need to work on because those are not on the control and all the regulations in the world that are currently happening, not the EU act, not the UNESCO's recommendations are looking at those things. So these are high level political statements, and if we don't work at that now, then we have a problem. I also include the concentration of data power and wealth in very few players around the world, which will increase the gaps in all respects between countries and within countries. So that's really a very big concern. And then the last prediction is within 15 years, we don't need to go to the physical gym, but we need to go to the mental gym to keep our minds active.

Odia Kagan:

I'll finish in like 30 seconds. I think I read there was an article in Time Magazine that said that basically compared this to the movie Don't Look Up and like, what's going on with AI. And I think the question is are people too jaded to look up? Right. Which was the problem in the movie, right? Are we like, too jaded? A lot of people are saying, oh, this is the end of the world tomorrow, and there's a lot of people or bots that are saying, no, no, you're fine. Keep using

us, we're good. Right? And as a person who deals with regulation for a living, I think it's obvious that this needs to be reigned in.

And there was a letter with like a thousand CEOs that actually agrees with the lawyers, interestingly. So that's fun. But I think this is the point. I think we need to kind of, it's not too late, but it's time to do something. And I think hopefully if we do that, then I can kind of sign up for Richard's optimistic predictions and not optimistic, right? The people are always gonna look for a magical solution and I see how my kids are doing their schoolwork now, and I see that I put in a lot more effort anyway. So, we will have to kind of find replacements for that. So I think maybe the prediction is that we need both kinds of regulation as well as self-regulation, as well as better parenting to deal with this.

Noelle Bloomfield:

Amazing. Well, Paula, Odia, and Richard, thank you so much for participating today. This was a fantastic discussion and we really appreciate you sharing your perspectives and your expertise with us. To the audience, thank you for joining us. The transcript and the recording will be available at gloat.com/resources. If you wanna revisit it, share it. Thank you everybody, and have a great rest of your day.