MIT CSAIL Alliances | Daniel_Weitzner_Project_6

Welcome to MIT's Computer Science and Artificial Intelligence Labs Alliances podcast. I'm Kara Miller.

[MUSIC PLAYING]

On today's show, for many companies, data is gold. But under the surface, there's a problem brewing.

70% of users surveyed recently are concerned or very concerned that their data is being misused in some way. And only 7% of users feel that they actually understand how their data is being used.

Daniel Weitzner leads the MIT Future of Data Initiative, and he served in the White House as the Deputy Chief Technology Officer of the US.

There are a lot of businesses that are really unsure how to use data in more innovative, aggressive, useful ways without harming their long-term relationships with their customers. So this is a kind of a sticky point that we've gotten ourselves into huge analytic capacity, huge amount of data collection, but without the corresponding trust to really move forward. So today, a peek into Weitzner conversations with businesses and why he thinks a new, safer, and potentially even more profitable day for our data is around the corner.

Our goal with respect to personal data is to have the same level of trust that we have in personal data that we have for the flow of money in our banking system. And that's the way that we're going to really enable new businesses, new research activities, and all kinds of new services to develop is going to be based on that kind of trust.

That's all coming up in just a minute. But first, in so many of our conversations here, we talk about generative AI. It's one of the game changing technologies of the last few years. So how can you understand it better? How can it help your business? MIT CSAIL and MIT xPRO are offering an online course this fall designed to answer those questions. If you're interested in getting more info. Email us podcast@csail.mit.edu. Listeners to the podcast get 10% off the course. So that email is podcast@csail.mit.edu.

Your data, your health data, your financial data, it's without question precious and not just to you.

We are really at the beginning of the personal data revolution of the data analytics revolution, not at the end.

Danny Weitzner spends his time thinking about data, how firms capitalize on it, how you can be sure it's not being used in a way you wouldn't want, and how regulators seek to protect it. He says, look, if discussions about your data seemed ubiquitous in the era of social media, you ain't seen nothing yet.

The most important and valuable uses of data are still out there on the horizon. We have a huge amount of data on our personal health. We have huge capability now because of all the sensors that give us indicators about just what our bodies are doing. We have the ability to live in more healthy ways. We have the ability to discover more about how to cure diseases, to discover more about what the most effective drugs and treatments are.

And it goes on and on. You can learn more about your driving habits and get financial incentives to make safer choices on the road. But here's the thing.

There are some real trust gaps that people are nervous about, what their Apple Watch knows about them or what their car knows about them or what their banks know about them. And I think, unfortunately, gotten ourselves into a situation where on the one hand, people are really excited about these new services and are adopting them quite rapidly but are not doing it from a position of trust.

Weitzner, who's also a Senior Research Scientist at CSAIL and Founding Director of the MIT internet policy research initiative, believes companies have to deal with data differently than they did in the past if they want access to the data that is most important to them. He has a solution to the problem, which he's already working with companies to implement. We will get to that. But first, he says the reason that trust has broken down is pretty simple. Consumers have not always had a good handle on what's happening to their data. Consider one of the most famous examples of this. In 2014, a small company called Cambridge Analytica began to get crucial data from Facebook.

And what was then done with that data was the possible extreme political advertising profiling and maybe even manipulation of users against their will and without their knowledge.

Weitzner says that it's important to understand what went so wrong almost a decade ago if we want things to be any different now.

Facebook, for a long time, had a pretty open environment where they wanted to be able to share lots of data about their users with lots of third-party services. One of those services was an innocent little academic research outfit out of Cambridge University in England, and they collected data on about 250,000 Facebook users with their consent for the purpose of doing research, psychological research.

But the way that data was shared from Facebook to Cambridge University was such that not only did they share the data on the 250,000 users, they also shared the data on the friends and friends of friends of all those users. And then, the Cambridge University research had all this data, turned it over to a commercial entity called Cambridge Analytica, which ended up using that data for commercial political advertising purposes.

The campaign of Texas Senator Ted Cruz, who was then seeking to get the 2016 Republican nomination for president was using the data. Later, the campaign of Donald Trump, who, of course, did get the nomination, also used the data, which by that time, covered something like 50 million Facebook users.

Now, why did this happen? It happened because Facebook shared the data without any ability to exercise control once it had passed to these third parties. Then, they actually had contractual agreements with those parties that says there are only certain things you can do with the data, but they had no way of enforcing those agreements. And certainly, in retrospect, they probably wish they had because they were ultimately, by the way, fined $5 billion for these violations. Facebook was. And Mark Zuckerberg was almost personally sent to jail. So serious consequences, but only after all the harm was done.

Ironically, Weitzner says Facebook was being supervised by the Federal Trade Commission at the time because of previous questionable handling of data. That's striking. And it offers one big takeaway.

The basic technical infrastructure that we're building to collect and analyze all this personal data does not have the ability to provide users basic traceability and accountability of what's happening to their personal data, and it doesn't provide companies who actually want to be responsible with the ability to do that. And the icing on the cake is that even our most powerful regulators don't actually have the ability in a reliable, consistent, and scalable way to really figure out whether privacy rules are being followed or not. So that's how we end up with these large-scale abuses, and that's why users correctly don't trust what's happening out there.

Going forward, Weitzner says consumers and governments are going to hold the private sector to a higher standard. And as AI comes into the picture, companies are going to have to buckle up.

If we're going to move into new environments where we're using personal data even more aggressively, there are going to be higher expectations for privacy. And I think the real question is, who's going to be in the lead in setting those expectations? I think we could wait and see what regulators decide to do, or I think we can-- and I think we are seeing a number of leading financial services firms trying to take a hard look at this question themselves and with us and say, what do we want these relationships to be about fundamentally? And how do we build an infrastructure, technical infrastructure that facilitates this kind of data sharing in a way that also promotes trust?

Right now, lots of consumers-- and I can speak for myself here-- likely have very little sense of where their data goes, who has it, clothing companies, car companies, social media companies, who's selling it to whom. Weitzner says that confusion is pretty widespread.

So I think a lot of the digital world today is really driven by one business model. It's an advertising and marketing-driven business model that relies on what is really covert and unconsented collection of a huge amount of personal data.

We're used to being in a world where there's lots of advertising around. That's the nature of a market economy. I don't think there's anything wrong with that. But what's happened is that we've allowed in the social media and e-commerce environments collection and really surveillance of people's day-to-day activities in a way that they really aren't able to control. So, Kara, you're not alone. 70% of users who were surveyed recently are concerned or very concerned that their data is being misused in some way, and only 7% of users feel that they actually understand how their data is being used.

And now, of course, we recognize that people continue to participate in a lot of these services, but they don't do it happily. And beyond that, there are a lot of businesses that are really unsure how to use data in more innovative, aggressive, useful ways without harming their long-term relationships with their customers. So this is a very a sticky point that we've gotten ourselves into, huge analytic capacity, huge amount of data collection, but without the corresponding trust to really move forward.

OK, so you've been in government. You talk with a ton of companies all the time. You're in academia now. So you understand these different players and the people who would be in the room. And as you said before, we're at this moment where data could not be more important in terms of maybe really helping us to be better drivers, to have better health. So if you had to have a plan, which I know you do, for how we move forward in a way that people do not feel taken advantage of for their data, but data can be used in a useful way, what's that way forward?

So the plan is, at a high level, actually pretty simple, at a technical and operational level, has a certain amount of both computer science and legal complexity to it. But here's the really simple goal that we are aiming for. People who share their data out in these increasingly complicated ecosystems, whether it's financial data, personal health data, driving and location information, whatever it is, they should be able to number one, have control over who has that data at any moment, and number two, they should be able to answer two questions-- Who actually has my data? And what's being done with it?

And here's the punch line of this-- think about the fact that we're trying to restore trust here in this data ecosystem specifically not so that people can hide, not so that people can hoard all their data under their mattresses, but to the contrary, so that people can feel comfortable sharing data with a whole variety of services that can offer them new things, offer new benefits, increase scientific knowledge, increase medical knowledge, help them save money, help them do new things.

And the kind of trust that we are aiming for is actually the kind of trust that people have at a practical level with their banks. We mostly feel that if you deposit money in a bank, it's going to be there the next day and that if you tell your bank to pay someone some money, they're going to pay it. And one of the key reasons that people trust that is because at any moment, we can look at our bank account and say, well, what happened? What did I charge to my credit card? Who did I pay? Did I write a check? Did my payroll check get deposited? Et cetera.

We have basic transparency and traceability of the information that is key to what we expect from our banks, and our goal with respect to personal data is to have the same level of trust that we have in personal data that we have for the flow of money in our banking system. And that's the way that we're going to really enable new businesses, new research activities, and all kinds of new services to develop is going to be based on that kind of trust.

So as a consumer, what would that mean for what kinds of questions I would be asked or what I'd be able to see or know? How does that manifest for the average person out there?

Right, so let's take a case study. We're right now at the beginning of a extraordinary transformation in the financial services industry in the United States and really all around the world. It goes under the name "open banking." And many of us have experienced these services already. In the good old days, we all had maybe one bank and one credit card, and if we were lucky to have investments, we had one brokerage service. Today, people have all financial relationships with all kinds of organizations.

And because of that, we need to move personal information around in that ecosystem much more aggressively than we used to. Anyone who has used a service like TurboTax to prepare their tax returns has gone through the exercise of instructing their bank and maybe their employer and maybe their credit card company or their brokerage service to send all of the data from your last year's financial life to TurboTax so that they can put that together into the beginning of what will be your tax return.

So that's a classic open banking interaction where you as a user actually want to share data with a number-- with a particular entity, and you want them to do one thing with it, but nothing else. Now, in order to facilitate that, there are actually a number of intermediaries of third-party services that have to be involved to make sure that data moves correctly. And the financial services industry is getting better and better at actually moving that personal information around in a secure and reliable way.

The question that has not yet been answered in that set of interactions is, how am I going to be assured of my privacy? We're pretty confident now that the data is secure against malicious theft or hacking because there's good security on all these data flows. What we want to add to that is the ability to have clear privacy controls and traceability on all those flows of data.

So if I tell my bank and my brokerage and my employer to please share data with TurboTax, for example, I want it to be very clear that that's a one-time sharing event, and TurboTax is only to use it in order to help me do my taxes, and then, they should get rid of the data and do nothing else with it.

I don't want to discover six months later that TurboTax looked at my tax returns and decided that, well, maybe my credit worthiness was not quite as good, or maybe I have a lot of extra money and advertisers might want to reach me. I didn't share that data for the purpose of having lots of other unrelated parties get a peek at it or learn things about me. I shared it for one purpose.

What we've been working with our research team here at MIT, as well as colleagues from the industrial members of the Future of Data Initiative is working on a protocol that we call OTrace, which allows users to number one, be able to specify exactly the purposes for which they want to share data in that kind of scenario and, number two, a mechanism by which users can see exactly who has that data at any given point in time and what they're doing with it.

And particularly, we want to be able to assure users that if I share data for, say, a three month period, that when I look at three months plus one day, I'm going to discover that the services I shared the data with have promised me that they've deleted the data because they no longer need it, because they're no longer authorized to use it.

So what we're aiming to do is to create an environment where, at any moment, users can exercise control over the data that they're sharing, specify how it's used, and be able to have a look into the actual uses to feel confident that the rules and the agreements are actually being followed.

It sounds like what you're describing, what you see coming hope is coming in terms of privacy between companies and individuals is like a credit card statement where at any time you could check in on how your data is being used, and, I mean, I don't know if this has happened to you, but it has happened to me, where I normally don't keep like-- I do not check my credit card statement every day, but I will get a call if something very unusual is happening. And they say like, hey, are you buying something very expensive in Canada right now, right? And then, you have to say, actually, yes, I am. Or, no, I'm not. Because as you said, it's very outside the bounds of what I would normally do.

That's right. And I think at the core, what people want is respectful use of their personal data. That does not mean that we want to be asked every single time, is this OK? Is that OK? Is the other thing OK? If it's part of what a person would normally expect to happen and which would benefit them rather than harm them, I think we should be able to err on the side of encouraging those new uses and enabling those new uses.

But in order for that to happen, we have to have exactly the traceability and accountability that you're talking about with credit card statements where I know that either I can look at my credit card statement and say, well, I think my bill was just too high this month. Why did that happen? I didn't think I spent that much money. You can go back and check.

Or, your credit card company or whatever firm is involved in facilitating whatever these transactions are with personal data can analyze all the different interactions with your data and flag the statistical anomalies. We want to be able to apply that same kind of data analytic, machine-learning-driven scrutiny to the use of personal data to help people recognize when something that's happening with their personal data probably isn't quite right.

So if I run a company, how is the sort of new way of approaching data that you foresee in the future? How is this going to change my life? What should I be doing to get ready for it?

That's a great question. So we're working right now with a group of financial services firms to add a layer of privacy protection on top of their existing open banking protocols. And what that means is that the firms will commit to report to users, to customers or to their traceability services exactly what they're doing with the data, and they will commit to making sure that when data gets shared from one party to another, along with that sharing event, comes the consent that users provided along with that data, the terms that users agreed to about what can and cannot be done with the data.

So as a participant, for example, in that financial services ecosystem, it means that you have to be ready to be up front with your customers, with your users about what you're doing with the data and report to them. In the same way as you tell them where their money is, tell them where their data is, tell them what you've done with the data. We're developing the protocols right now in order to facilitate that communication amongst data sharing partners and with consumers. What companies will want to think about is how to use that to build a kind of a constructive transparency and openness and respectful relationship with their consumers in order to be able to win their trust.

Is that scary for firms to-- I don't know-- embark on a whole different discussion than they're used to?

Well, I think that a lot of firms in the financial services industry are looking hard at the privacy status quo on the internet today. And what they're seeing is a kind of a corrosive, distrustful, and disrespectful relationship between internet users and the big services out there. And financial services firms, I think, in many cases, take a longer view of their relationship with their customers. They're not really out to advertise against their customers.

They're out to build long-term relationships, provide services over the long run where they certainly hope to make money and do make money, but they're not trying to make money, we believe, behind their consumers backs. They're trying to make money buying by being upfront about the clear terms of the relationship.

Now, financial services firms don't always go along with this completely willingly. We have regulators, obviously, that have to keep track of the behavior of firms. We've certainly seen financial regulators in the US and all around the world work to make sure that, for example, if someone signs up for a credit card, they get a clear indication of what the interest rate is actually going to be.

We've had regulations that require transparency on the terms of those kinds of credit relationships, and I think everyone's agreed that's good for the credit marketplace, generally good for consumers, and creates a clear set of expectations for firms. We want to have that same transparency and dynamic traceability for the data relationships that firms get into with just as they have learned that they have to be transparent and open on the basic financial relationship.

So basically, companies should get out ahead of this and not wait for regulators to be like, this is what you have to do. Do something really great in advance of the regulator saying that,

Yes, yes.

Weitzner says, the rollout of this new data traceability standard will likely happen over the next several years, and this happens to be a very good time for it.

What we're excited about is that right now, the financial services industry is building a whole bunch of new services, new infrastructure, and new standards to facilitate the flow of personal data. So it's the ideal time to integrate privacy along with it rather than having to layer it on top once the whole system is already put together. He thinks this approach will cascade out of finance and into other sectors. We've started with finance because it's a very concrete and high-need application area because as I said, there's a lot of flux in the technical infrastructure today that we think we can leverage.

But certainly, this OTrace protocol is really designed to provide traceability and accountability and control all for any environment where data is being shared amongst multiple parties, and that covers the range of applications that we've been talking about, finance applications, personal health data, location, and automotive telemetry services. You name it.

But isn't there going to be pushback from companies? The idea of more hoops to jump through is very rarely met with excitement.

What we're hearing from companies today is that as they look forward to more aggressive deployments of various kinds of AI systems of large language models, what they want to be able to do more than anything else is to train those models on a wide range of data that they have inside their enterprises and data that they get from third parties. In order to do that, they have to know where that data comes from. They have to know what they're allowed to do with the data, what they're not allowed to do with the data.

So beyond the individual customers' interests in personal privacy and traceability, firms have told us that they have a very clear need to just have a set of standards and a structure by which they can get their own internal data in order so that they can train on it, so that they can integrate it into the development of new AI models and techniques. So beyond just personal privacy, this capability of traceability, we think, is really the foundation for the data governance that firms need for lots and lots of reasons. So it's going to be an essential piece of infrastructure one way or the other.

Daniel Weitzner is Founding Director of the MIT Internet Policy Research Initiative. He's a Senior Research Scientist at CSAIL and Head of the MID Future of data initiative. Danny, thank you so much. This is great.

Thanks, Kara. Great talking with you.

[MUSIC PLAYING]

For more info on the plan that Danny talked about, just head to FutureofData.MIT.edu. And before we go here, a reminder that CSAIL's Gen AI online course is starting soon. It's a technology that's reinventing jobs, so learn how to use. Apply and strategize with this new course, driving innovation with generative AI created by MIT CSAIL and MIT xPRO. If you're interested in more info, email us podcast@csail.mit.edu. Listeners to the podcast get 10% off the course. Again, the email, podcast@csail.mit.edu.

And if you're looking for another great podcast, check out*Me, Myself and AI.* In each episode expert hosts and researchers talk with AI leaders from organizations like NASA, Upwork, GitHub, and Meta to explore how organizations achieve success with generative AI and what challenges and ethical considerations they face along the way. Listen to *Me, Myself and AI* wherever you stream podcasts. I'm Kara Miller. This podcast is produced by Matt Purdy with help from Audrey Woods. Join us again next time and stay ahead of the curve.