EPISODE 1130

[INTRODUCTION]

[00:00:00] JM: Biometric authentication uses signals from a human's unique biology to verify identity. Forms of biometric authentication include fingerprints, eye patterns and the way a person walks, otherwise known as gait. UnifyID is a company that builds systems for biometric authentication. John Whaley is the CEO of UnifyID and he joins the show to talk through these techniques for biometrics as well as the implementation details that UnifyID has built to turn these into a reality.

I want to mention that we are looking for writers as well as podcasters. We are considering opening up a number of positions for people to podcast for Software Engineering Daily to broaden the voices and the ideas and the concepts that are explored on Software Engineering Daily. You can send an email to erica@softwareengineeringdaily.com if you're interested in being a podcaster.

[INTERVIEW]

[00:01:01] JM: John, welcome to the show.

[00:01:02] JW: Hey, great to be here.

[00:01:05] JM: You work at UnifyID, a company you cofounded. And you do biometric authentication. What does that mean?

[00:01:12] JW: So, there are all these things that are very unique about each of us. I mean, the most common type of biometric authentication you know about is things like fingerprint. Sensibly, everyone's fingerprint is unique, and then you can record that and then know that whether it's you or not. I mean, what we do is something called behavioral biometric, which basically instead of looking at something static like your fingerprint, we'll look at your behavior. Like this dynamic behavior that you do, things like the way that you walk or the way that you

© 2020 Software Engineering Daily

hold your phone, or there are millions of other little habits and idiosyncrasies that you have and then use that for authentication.

And by doing that, we are able to get something that's very seamless that just doesn't require you to require any conscious user action. You can just be yourself and there's enough that's unique about you. We can actually authenticate you based on that, and it's continuous. So you don't need to – You can walk away from a computer. You put your phone down. Somebody else picks it up, things like that. We're able to detect that change and then know whether it's you in a very seamless way.

[00:02:16] JM: I feel like biometric authentication has been a thing for a while. There have been a lot of companies that have tried to do this. Why is UnifyID different?

[00:02:25] JW: Yeah. So the biggest difference is the type of factors that we use. We focus entirely on passive factors. So, things that don't require any conscious user action to do. I think that's always been the challenge with any form of authentication is the users, the human beings. It's just human nature. I mean, people are naturally lazy and they don't want to go through these additional authentication steps, because it's adds friction to the user experience.

I mean, if you look at the way that people choose passwords – I mean, there's been plenty of information about with password breaches and how to choose a good password, etc. If you look at every year, the list of breached passwords, it's always the same passwords is 123456, or the word passwords, or those type of things. So, changing human behavior is really hard.

Now, I mean, if you look at just historically across technology, like there's always been this shift where initially human beings have to adapt to the limitations of technology. And then eventually technology reaches a point where humans – Where it can adapt to humans and the way that they are. And so authentication is not a new problem. I mean, it goes back to prehistoric times. Any social creature needs the ability to then identify and authenticate like are you who you say you are. Are you part of my group or not? This goes back a long time before technology existed.

And the way that people had always did in the past was you look at their face, you hear their voice. You see the context under which you see them. Maybe their possessions, things like that.

But it was always a very natural thing. And then passwords came along around 450 years ago. And so this was very much adapting to the limitations of technology. Let me enter this sequence of symbols and numbers and letters. And then that's how you know that it will be me, right? But that's not the way that people had traditionally identify themselves in the past. And I think it's about time technology has now reached a point where individuals can be themselves and then there's enough sensors in people's lives as well as like the technology has reached a point that you can then identify people just by their natural behavior rather than having them do something explicit.

So, the difference is because we're doing this completely passively, it opens up a lot more different use cases where you don't really have to change the user experience. You can just have a much more seamless user experience, like opening a door, for example. With our technology, when you make an SDK, you can link it to any iOS or android app. We also have components for web, JavaScript as well. But you walk up to a door and you have the app installed, and then the door already knows that it's you and then the doors unlock. You don't need to take out your phone or do something extra. And so for use cases where you have security concerns but you also care about the user experience, that's where our technology really comes to play.

[00:05:24] JM: So you're talking about like I've got my phone in my pocket and I'm walking up to a door. As I'm walking up to the door, my phone is measuring my gait through its accelerometer or something like that, and then the like Bluetooth talks to the door and tells the door I have my phone? Am I understanding that correctly?

[00:05:40] JW: Yeah. Yeah. I mean, that's right. And we're able to use gait. Our gait analysis is very accurate. We're able to get accuracy on par with a physical fingerprint. So about 1 in 50,000 false-positive rate just based purely on somebody's gait. There's a lot of science that goes behind it, and we have a lot of peer-reviewed complications, etc., on the accuracy and how exactly the technology works. But yeah, it turns out that these little aspects of your emotion and your behavior are unique enough that you can actually use them for authentication.

And so because you're able to do this passively, you can do this in cases where – I mean, it's not just doors and physical world, like a car, or an ATM, or travel, or etc., but also I walk up to

my computer and I sit down, or I'm hailing a ride on a rideshare app or like getting a food delivery or these types of things. There're a lot of these little aspects of people's behavior that are unique that you can use to passively authenticate them. And not only for authentication, but also de-authentication. Everyone always talks about how do I authenticate the user? Nobody talks about how do I de-authenticate them? And the way that people do that today is they'll just set a timeout. They'll say, "Well. Okay, after an hour or a day or a month, we're going to force them to re-authenticate." Because we don't really know when somebody leaves and it's not the same person anymore.

When you're able to passively monitor their behavior – This is why it's called implicit authentication, because you're able to authenticate without making any explicit action. You can just implicitly be authenticated just by your natural activity. Because it's just passive and it doesn't require the user to do anything, then you can also handle de-authentication cases as well in a much smarter way. So I get up and I walk away from my computer. It automatically locks. It logs me out. I walk to my computer, it automatically logs me in. That's the type of user cases that this technology enables.

[00:07:31] JM: Is there anything other than gait analysis that is worth mentioning?

[00:07:35] JW: Yeah. We look at a number of other environmental factors and other type of behavioral factors as well. So environmental factors, these are things like around your device. What are the set of Bluetooth beacons and Wi-Fi access points? And what are the exact RSSI values that they have when you are in a particular position or a location?

This type of things and location and location history, not just where you are right now, but where you came from. I mean, these all provide little hints and these type of data is not necessarily useful for identification of a particular individual, because there are maybe other people who are in the same building as you, who have access to your physical space and would be able to replicate those. But, again, this is just kind of a noisy signal. You can combine multiple of these noisy signals together and they get something that is more – Where you say, "Well. Okay, yes. Look, the phone is in the person's house. I see somebody walks up and like the gait matches. I see them sit down. The way they sit down is the way that they normally do. The Wi-Fi access

points are all consistent, what I've seen before. Maybe the person has a wearable device or other type of Bluetooth headset, etc." These are also consistent.

At that point, when you see all of that data, asking the user for a password at that point is really – You're not really adding much security at that point, because you already know fairly conclusively that this is going to be the correct user when you combine the biometric aspect along with some of the environmental aspects. And so this is what people have always thought of authentication as a binary, like yes or no. It's like, "Yes, it is definitely you." Or, "No, it's definitely not you." And the truth is that these are – Even when using passwords or knowledge-based factors, there's no absolutes. I mean, all you know – When you get a password and a password matches, all you really know is that somebody provided he same sequence of characters, right? And there are plenty of other ways to get somebody's passwords. I mean, people could be phished. They may have just chosen like a bad password. They may have reused the passwords that they use somewhere else.

So, the notion of like using a password as a form of identity, I mean, it has its pros and cons. But certainly, it's not perfect. I mean, there's also the other ways that people do authentication like based on knowledge-based factors like what's your mother's maiden name? Or what's your social security number? This type of data, increasingly, at one point in time it was supposed to be private data. It's increasingly not private data in this era of data breaches and everything else. So these are not reliable factors to be able to authenticate someone.

I mean, even just the whole notion of I'm going to validate your identity by using information from public records to determine whether it's you or not. I mean, there's an inherent contradiction in that. But that's a status quo. I mean, that's the way that people do it today. They'll ask you, "Do you have a mortgage at this address? Or where did you grow up?" Or these types of questions that maybe with a little bit of research, a hacker can very easily figure out what the answer to these questions are. Not to mention, the fact you'll reuse the answers everywhere. One data breach anywhere will then compromise you everywhere. This is why, again, there's no silver bullet in authentication or security. But this capability of being able to passively authenticate someone without requiring them to do any user action was something that didn't exist five years ago. It exists now. And I think and we think that it's going to be an important capability within the toolset of authentication going forward.

[00:11:16] JM: So, this was as simple as collecting the Bluetooth signals that my phone is been around, and collecting the gait of me and knowing that this is my fingerprint. That will be pretty easy. You could just like hash those together or something and get an exact match, and that's me. But I assume that there has to be some kind of fuzziness. There has to be some kind of wiggle room, because I'm not always going to be around the same Bluetooth devices. My fingerprint is going to change over time. Maybe I sprained my ankle and my gait changes. How do you deal with those kinds of variabilities?

[00:11:50] JW: I mean, that's where the machine learning aspect comes in, because you're absolutely right. This cannot be an exact match. It has to be a fuzzy match, and it has to be – Again, it's a probability distribution, and like you have these multiple probability distributions about different factors of like, "Well, what –" If you say, "Well, the gait was a strong match, but like the person is at an entirely new environment. Do we say that that's the user or not?" I mean, it depends on the use case and the threat models that you're worried about. And so this is why you're having that machine learning aspect there to be able to then have these types of fuzzy matches is important to make something which works well in the real-world.

[00:12:33] JM: And keep tell me more about like the ensemble playing of different signals and machine learning and how you use it?

[00:12:39] JW: Yeah. We use machine learning at different levels, and this is not just deep learning like let's take all of this data and like throw it at this magic thing called machine learning, and then you'll magically get the answer. There are a lot of different aspects to machine learning, and some cases it would do something very simple. In other cases it's going to be much more complex and sophisticated.

At the basic level, if you look at each individual factor, like for example somebody's gait or somebody's – The way they walk or when they pick up the phone, that motion, or their typing speed and cadence and the way that they move their phone around as they type around a keyboard. This thing – At the individual level, there's a lot of noise, inherent noise within that signal. I mean, much the same way when you're doing facial recognition. You can't just do like a pixel-by-pixel comparison. In some cases, the lighting will be slightly different. The person's

head will tilted in a different way, things like that. So this is what machine learning comes in where you provide enough examples of that for that user. And then the algorithm basically adapts to find the things, “Okay. What are the things that are very consistent about this that make this person unique? Versus, “This is somebody else. Or this is somebody else who’s trying to attack the system.”

And so much the same way we do the same thing for each of this individual factor. So there's a big machine learning component when we’re talking about individual biometric factors. Then on top of that, there's a meta-level of, “Well, when I’m combining multiple of these things together, how do I actually make a decision about like whether this is the correct user or not?”

So that's the place where there's kind of a – It's not just that pure like unsupervised learning problem. We build in some prior knowledge from our customers and from having human beings being in the loop there. So the way that that works is that there is – Basically, you can imagine that each of these – Imagine that you have a group of experts that they’re all trying to argue about whether this is the correct user or not. And then each of those experts is a specialist in different areas. You may have a specialist, like I am the gate specialist. I am looking at just like the person’s gait, right? Another one will be I am the location expert. I know everything about location patterns and like the way that people behave, right?

And then there're other ones that are specialists on, “Okay. I'm a specialist in understanding attacks and sophisticated attacks when people are trying to spoof data.” So you can imagine that all of these, they all get together in some conference and they're basically trying to kind of battle it out and then say, “Okay. What should be our final recommendation based on kind of all of these different data points?”

And so the truth is, to make a system that works well on that, you need something that's going to be able to then reliably identify which experts are reliable in different contexts and for different users. And so that can be an adaptive. It's not like the gait one always wins and the location is always second, these type of things. You want that to be adaptive and then vary depending on the user and the use case. The other key point there is that there are correlations between these factors.

So, like one of the experts being wrong, like often – I mean, because there are some shared underlying process that's happening that they're both measuring, that like when one of them is wrong, then the other one may often be like wrong in a correlated way, right? And there're also anti-correlations there as well. So understanding which of these experts are correlated so you can then weight them appropriately. There's additional like meta-machine learning problem that's on top of that that is able to combine these in an intelligent way to ultimately provide a score.

Now, what we provide is we provide a set of APIs so you can get access to that low- level data. For example, here's what the user was doing. We saw the user walk. They took 23 steps. Here's how the gate score, the match score varied over time. We then saw them sit down. We saw them put their phone down on the table, and the phone has been still since this amount of time. Those are the raw ingredients that come into the decision algorithms.

We expose those as well, because those are useful for – You can tailor the data they're based on the particular use cases. But then we also provide this meta-level for passive authentication on top that say like, "Okay. Given all of this data, how likely is it that the user is just sat down at their computer and is trying to login? Or how likely is it that they are trying to make a phone call, like call-in to a call center and say like – And then is this the correct? Or how likely is it that they're trying to approach this door?" And that's the type of data that we can use.

The last part I mentioned around the machine learning side is that it's not all about your data versus not. The truth is that there are billions of people in the world and they all act in a similar way. And so much in the same way in the machine learning side like for things like image recognition and stuff, it's not like if you want to train something for image recognition, the right answer is not to just go and collect all of your data, your own data just to do this. That right way to do it is to leverage some existing data, like ImageNet and other things where they have many, like huge, huge datasets, and then use transfer learning to then transfer that into your domain.

Basically, already within the machine learning, the neural network that they've trained. It already has a lot of the knowledge about, "Okay. Well, this is the general shape of things, and here's how to distinguish between different individuals, etc." And then you can then specialize that problem just to work really, really well for like a particular individual or a very particular use case,

hotdog versus not hotdog, for example. That way to do that is not like have a million images and that. It's to basically use a user transfer learning.

In a very similar way, we have a huge dataset like over 30 million devices of people who have volunteered to basically donate their data to this. And then we use that to train these very accurate neural networks. For an individual, we don't need that much enrollment data. We just need a small amount of enrollment data, because we're leveraging all of that data just from humanity and like the way that people walk. So we're going to have to start from scratch every time. We kind of start from 90% of the way there, and then the question then becomes, "How are you unique? What makes you unique compared to other people?" And then just focusing on those aspects instead of just trying to train the whole system from scratch.

[00:19:19] JM: The SDK component of this is definitely appealing, and because I think it's easy for people listening not to quite understand what we're talking about here. But I can have like a – Let's say I run Monzo bank, for example. I have no idea if Monzo is a customer of yours. But Monzo bank. I have a banking application on my phone, and you could make it so that UnifyID pairs, or UnifyID is implemented in the Monzo app. And they might want to use it to use biometric authentication to know that the person who's holding your phone should have access to your Monzo banking app. Am I understanding the use case correctly?

[00:19:58] JW: That's right. Yeah. We don't we don't make our own apps. We make some demonstration apps. But mostly we just integrate within existing apps. And then this is an SDK you can just sit in the background. We leverage the permissions that whatever the app may have. We don't require any special permissions because we just – Fundamentally, all we really need is accelerometer data, which doesn't require any special permissions to access.

If you do have additional permissions, we can leverage those as well. But then yeah – And then you can then use that to then – In the example of the banking app. And this may be depends on what your use cases. It maybe when somebody calls in – So the call center from their phone, we can tell the fact, "Oh, okay. We saw this behavior that happens. Somebody picked up the phone. We saw the motion associated with that. We saw them bring it up to their ear, etc." Then we're able to correlate that on the backend. When you dial into the IVR system to say, "this incoming call is actually coming from this physical device." And not only that, but this physical

device is currently being held by this person with 99.9% accuracy because we know they're calling from their house. And like the motion associated with the call and the fact that they're previously walking a minute or two ago. You're able to combine all these together, and then that provides a signals to the call center agent to say, "Hey, this is highly likely to be the correct person. You can give them a much more personalized experience." Maybe you don't need to go through all of the security.

Whereas on the other side, then you can – If a lot of these signals don't match, then you can then flag this as, "Hey, this may be a potentially fraudulent call that's coming in." And so you want to force – Basically require the user to kind of go through some additional verification steps. And so it is not just about dialing. It's also about like I want to do some type of thing that requires a step-up authentication. It's also not just on the phone, but also on the computer as well. Leveraging sensor data from the phone when you're going to login on the computer on the website.

For example, in your example, imagine that the user sits down and they want to do their online banking like on their – You want to do online banking on the website, right? And so they sit down and they go into the website. Well, okay. In fact, the last 10 times we saw you go into the website, you're in this location. Your phone was like this type of orientation, because you're, again, creatures of habit. Do you put the phone down on the table? Do you keep it in your pocket? What orientation do you do that? What are the devices that are around you? All of these things." And then we're able to build a normative model that say, "Okay. Well, the previous times we've seen this user try to do this action, this is approximately what it looked like."

And so we can then give you kind of some match score that says, "Oh, how consistent in this with what we've seen before?" And not just the visibility of what you see from a server-side, like in terms of what IP address they're coming from. But just even like the behavior, the more detailed behavior information they have there. So hitting extra level of visibility, then it allows you to just be much more adaptive around your authentication and decisions. So you can then make intelligent decisions about, "Okay. Do I need to send this person a push notification to verify to do this additional action? Or do they need to – You know what? I am already pretty confident. I can give them a more streamlined experience." You have the data and the ability to be able to make those types of intelligent and adaptive decisions.

[00:23:24] JM: Is this kind of work data-intensive? I'm trying to understand if it's just kind of lightweight data and then this is not something that's really that intense. Or if this is like you have to collect tons and tons and tons of data and you have to do munching and stuff on the fly? Can you tell me about like how data intensive the UnifyID SDK would be?

[00:23:46] JW: So, number one, the data stays on the local device. We don't send that data off the device. So the processing happens locally. We've done a lot of work to get this to be highly optimized. I mean, in terms of on the machine learning aspect, for example, we have very highly optimized kernels that we use for many of the algorithms that we use. Because this is intended to run even on the lowest end devices and to do so without having an impact on battery life or data usage. And so it's able to run in the background. The amount of power usage is very small, because we've done quite a bit to optimize this.

The other thing just to be clear, it's not recording all of the time. It's only – We have intelligent triggers that say, “Well, okay. Is something interesting happening?” Like, “Hey, we’ve noticed the step counter changing.” Or like, “It looks like somebody just got up and started to walk around.” And we haven't seen any gait scores for a while. Maybe let's record for a few seconds and just make sure this is still the same user.” Things like that. Or other type of triggers that we have around location, etc., that make it so that we can be intelligent about when we do processing. You don’t have to be basically just processing and recording all the time.

The other thing is that we’ve done a lot of work to then make sure that this will run on low-end devices as well. That's one of the benefits here, is that if you think about biometric authentication in terms of things like touch ID or face ID – I mean, those are great, except for the fact that they’re really only available on the more high-end devices. And the US market, they're fairly common. As soon as you go outside of the US, there are cases where the advanced sensor that you're talking about for the facial recognition or for the fingerprints. To make those to be hard to spoof, there has to be a certain level of sophistication within those sensors. Those become cost-prohibitive for lower end devices. But every device, every smartphone has accelerometer. Everyone that’s ever made has accelerometer. And the notion of like the type of motion that we’re collecting here, it's human motion. So, it's not – So even the

lowest end accelerometer and motion sensors have plenty of accuracy to be able to then tell the difference between different individuals.

And so part of the benefit here is that this helps to bring this type of strong biometric authentication even to devices that don't necessarily have the hardware or to lower end devices. And even in cases where you don't want to have that additional friction involved and like, “Okay, let me out to take my phone out and make sure that this is –” And then have them require to do a touch ID or a face ID, which by the way like nowadays in terms with COVID and everything, with people wearing masks and gloves, like those type of biometric authentication like facial recognition become a lot less practical and applicable in these type of cases. But gaits, because it’s passive, it’s happening in the background. There are many cases for which when you're out and about and you want to be able to have this type of continuous authentication. Things gait are a strong alternative to things like touch ID or face ID.

[00:26:54] JM: Can you tell me about how your engineering teams are organized?

[00:26:58] JW: Yeah. So we're a pretty engineering-heavy organization. My background, I was an engineer as well. I was a CTO in my previous startup, and I did my PhD at Stanford. I taught at Stanford as well. So I have a very technical background. Within the company, there's a client SDK team, which is focused primarily on the iOS and android and the JavaScript SDKs. We have a backend team that is focused more on the infrastructure. And then we have an ML engineering team, which is really about taking some of those more advanced algorithms and machine learning algorithms and then getting those implemented in a production way, right?

And so, that's how the teams like within the company are organized. Obviously, there is a fair amount of overlap, and like people kind of wear different hats at different times. It's not like the client team is strictly working entirely on client and then there is a backend – The people in the backend are only working on backend. Many times, you’re talking about protocols and other things, and algorithms, and then getting those implemented. Everybody ends up having to work together. But in terms of the specialization, that's where – That’s the structure of the teams that we have within the company right now.

[00:28:14] JM: What are the hardest engineering problems you're dealing with right now?

[00:28:18] JW: Yeah. Gosh! I mean, certainly as you hear the description of like this type of stuff we do. Myself as an engineer, and I imagine your listeners are like just immediately thinking about, “Oh my gosh! What are the things –” All the things I have to think about, worry about. I mean, battery is one thing that come up. Kind of latency, security. Security system, you don't want it to be easy to attack or break. The machine learning aspect, the real-time aspects. Certainly, this is a hard technical problem that we’re trying to solve. And this is something that had not been solved before, right?

But from my point of view, I mean, these are exactly the type of problems that I like to work on and like our team likes to work on. Life is too short to do things that have already been done and that have already been proven out to make the kind of the iteration of like the same thing. If you're going to spend your life and your time on something, you want it to be something that's going to have a big impact and that will matter. It will be something that's new and be something that's intellectually interesting, right?

And so that's something that definitely what we work on has all of those aspects. Huge impact obviously just in terms of like seamless authentication, authentication being such a huge problem for everyone, being a hard technical problem. This allows us to actually create real value. And then if the problem were easy and anybody could solve it, then you didn't really create much value there. But like when we do run into difficult problems and figure out a solution to them, then that means we created real value. Because the next person that comes along, they’ll probably fall in that same pitfall and they may not be able to figure out how to get out. So, we don't shy away from kind of trying to tackle some of these difficult problems.

To be specific, the biggest challenges that we have, number one, is like the security and privacy and machine learning. I mean, obviously, we don't want to become this honeypot for everyone's data. But we want to be able to build a system that works and works well. So, how to achieve this type of machine learning in a distributed fashion? Where the data is staying on the local device, but you’re still taking advantage of the global state? That’s one of the big technical challenges that we have.

Another one is around the efficiency of our SDK. Because, again, it needs to be able to run in the background. It needs to be able to then do this with minimal impacts on battery or resources on the phone. And so how to be able to best leverage that to achieve that? Leveraging things like – We'll leverage things like if your phone as a GPU, like we can leverage that for doing some background computation and everything. Sso we can free up the CPU on the phone so that we don't have as much of a performance impact and it'll will be more efficient and power efficient as well. And then optimizing things for communication, the times where we do need to do communication. Deciding when we need to communicate, because the power usage of powering up that radio to be able to send data is significant. So, those are on the client side. Those are some challenges on the backend side. Things around the scale and availability are always challenges and how to do that in a way that is secure. That's an additional challenge.

And then on the machine learning side, like these are new algorithms. These are capabilities that nobody has developed before. Being able to identify somebody based on their gait based on five seconds of walking data, etc. There's a lot of sophistication that happens there. It's not just – I mean, it has to be this combination of math, and theory, and machine learning algorithms and capabilities there wedded with what is practical and what we can actually implement to run in real-time like on a phone? It's required multiple innovations on that side just to be able to then get something that kind of works well.

Now, I mean, we've been at it for a long time. We've been working on this full-time for the last 4 years. Before that, there's a side project we've been working on for a time before that as well. There's a lot of history behind it as well, but like that is a never ending quest within the company to improve the effectiveness of the solution, to handle more use cases, to get higher accuracy, handle more people. And so that's been a big focus, and will continue to be a big focus of our ML engineering team.

[00:32:43] JM: Can you do domain-specific verification? I don't know. I'm searching on Amazon. I'm looking on Amazon and you're authenticating me as I am behaving on the website?

[00:32:55] JW: Yes. We have capabilities to do that as well. I mean, if you think about the way that somebody clicks around a website, the way that you move your mouse, the way you scroll, the way you type. Each of these are pretty unique to each individual. And although like you don't

– The amount of data you can get from within a web browser is limited, there's quite a bit there. Those type of situations, it's more useful for risk and fraud type of use cases where it's like, "Okay, let me detect if this is – Does this look suspicious in some way? Does it look like this might be an account takeover or this might be a fraudster? Like somebody, it wasn't the same person as it was before?" It's very hard, because the amount of data you have in these cases is very limited. And often by the time you've collected the data, it's kind of too late, because they're already in the system at that point. Then it's more useful kind of for risk and fraud or kind of as supplemental factors, not as a primary authentication factor.

If you want something that's going to be more of a stronger signal, a primary authentication factor, something is not going to go from, "Okay, this is a 1 in 10 probability or 1 in 100 probability." To, "No. This is 1 in 10,000 probability, or higher." Then that's the case where you need to bring in some more historical data and bring in more context from other devices as well.

The other pieces on the security side, you get a lot more security when your data stream is not coming just from one device, right? I mean, if you're just looking purely from the web browser, that's a very thin data stream there of what you're actually able to collect within the sandbox within that browser. And it's more easy to compromise like in terms of if somebody has either some kind of like remote access agent or some other type of malware on that device. It becomes easier compromise.

Once you get into a product point where it's like, "Well, I'm using data from the phone and the computer." That becomes a lot harder to compromise, because basically, at that point, you need to then have malware and control of both what the phone is reporting as well as what the computer is reporting. And so that becomes just a much more sophisticated level of attack.

Yes. We do have capabilities around like, for example, mouse movements. The way that you click through sites, the way you scroll. It turns out to be really unique for each individual. Typing also turns out to be unique. I mean, if you type around three sentences or so, then usually we can we can tell who you are – Authenticate you just by the way that you type. Not necessarily what you're typing, but just by the way that you type. These are some of the signals we can use in the context of a web browser.

Like on mobile web, you have additional signals there as well. You have ability to then get access to motion data, and we can correlate things like that, so that when somebody is typing, for example, is the motion consistent? Not only like the way that they type like in terms of how long they hold down each key and like what part of the key and like the speed of each keystroke, etc. But you can also take into account some of the motion data that's associated with the typing there, and then determine things like, “Oh! If somebody types with one hand or with two hands, how they hold the phone? What angle they hold the phone? How hard they tap?” These types of things, you can pick up via motion sensors there as well.

Again, if you just look at those things in isolation, it’s not going to be strong enough. You're not going to really get a great signal there. But this is why you have to look at like just the whole thing in aggregates and then look at multiple signals. So even if some of them may have a high false-positive rate or there may be downsides or holes where like gaps where you’re missing data. Places where there are false-positives, things like that. By combining multiple of these passive factors together, you can get something that's more highly accurate and avoid some of the downsides of each of these individual factors.

[00:36:37] JM: I'm realizing now, I don't fully understand the onboarding process. So let’s say I’m a new user. I need to develop my fingerprint. Do I just like login to the app and then like walk around for a little bit and like scroll and stuff, and that builds my profile?

[00:36:53] JW: Yeah, precisely. I mean, there's no explicit onboarding. The most typical way that people do it is like, well, just say just have the app installed. Will this monitor your behavior like the way you walk and the way you scroll and like these other things as just going through your normal life, right? Or they used to do anything. After about a week, like one week of calendar time, by then, we usually have a pretty accurate model for that user. And then we then give pretty accurate results on whether this is the correct user or not.

Now, of course, like there’s perennial learning. So that what happens is that as you get more data, that data is then labeled as the correct user. Then the algorithm improves. And so in cases where somebody's behavior may suddenly change or there're different modes of behavior, etc., we’re able to then track that behavior and then adapt to it. But that's a general, what a recommendation is. You can also do explicit training where it's like, “Okay, get up and walk 150

steps." And then like we're going to use that model, and we can generate a model based on that.

Those types of cases, like those does work when you're in the similar context. So, if I do my training where I'm wearing a particular kind of shoes and it's an office and I'm walking, I'm walking on a carpeted floor, etc. I can build a very, very – With only a small amount of data, I can build a very reliable model that say whether it's you or somebody else kind of within that context.

The issue comes as like, well, how resilient is that model? If I change my shoes. I go barefoot. I change my phone location. These things that may naturally happen in other contexts. Then, if we don't have examples of those, then the system is unlikely to be able to then match and say – It will always err on the side of being conservative. It's like, "Well, I haven't seen the user do this before. So I can't really say that this is the user. I'm going to return back something that's inconclusive, or say this doesn't match what I've seen before."

And so in terms of the training process, this is why kind of just one week of a user's natural life begin in that variation in terms of their natural behavior like on weekdays and weekends in different contexts, like when they're walking fast and slow. When they're kind of using the system in different ways to give enough variation there so we can build a reliable model that's going to be resilient to these different types of changes.

And like I mentioned, the training doesn't aimed at for one week. We have the ability to then include additional data. When we get a data where – For example, let's say that you are walking and you hurt your ankle. You twist your ankle or something and your gait suddenly changed, or you moved to a new city. These types of things happen where you have a sudden change. Then what happens is that then our system will then come back and say, "Well, this doesn't look like what we've seen before for this user." They will then go back and fall back to some type of explicit authentication. It could be biometric, like a face ID, or a touch ID. It could be a knowledge-bases factor. It could be them calling in. There're a lot of different ways to do, kind of step-up authentication or fallback authentication.

In those cases they can then verify, “No. Actually, this was the correct user.” They can then tag that data in they say, “Okay. Well, this recent data here, this actually was the user.” And then that feeds back into the machine learning model. Then from the subsequent times where we see similar context, then that's where we can then say, “Oh, yeah. Actually, we’ve seen this before.” In the previous times, this was labeled as the user. We think this is likely to be the user.

Again, the ultimate question that we’re trying to answer is that – It’s called a binary hypothesis test, where basically here is some recent data. Is this more likely to be you? Or based on what we’ve seen from you, is this more likely to be you? Or is this more likely to be like someone in the universe, like someone else, right? That's fundamentally the question that our product answers, and then attempts to answer to say like, “Well, how likely is this to be you? Based on what we’ve seen, how likely is this to be you versus like somebody else in the universe?” And then we’d give a probability distribution basically that says how likely that is based on the different factors that we’ve seen.

[00:41:04] JM: seem are there any ways to goose the system right now? Are there vulnerabilities that you know about that you can tell me about?

[00:41:08] JW: Well, I mean, so we tried a number of different things. I mean, we’ve done things, like we’ve done tests. We've run tests with identical twins where like people are genetically the same, but then you try to mimic each other's behavior. Our system does appear to be resilient against identical twins. And we also done things where we had trained actors trying to mimic other people's behavior. Those ones are – Which appear to be resilient. The type of factors we’re using right now appear to be resilient against those type of things.

The other things that we’ve tried have been different type of machine learning attacks. We have a lot of history in terms of adversarial machine learning and understanding with people trying to do adversarial perturbations and other kind of synthesized signals, etc. We’ve done a lot of work on there. I mean, just to be clear, there's no silver bullet in security. And the only truly secure system is one that is just not turned on at all. I mean, basically, as long as every system is exploitable no matter what it is – The goal here is to make it so costly and inconvenient for somebody to attack that they will then go and try to do – Just go for targets.

I mean, certainly, if you're dealing with something where you have a state-level actor where somebody who has access to the supply chain, who has access to potentially zero-day exploits or other kind of things, or even things like they'll bribe one of our employees to go and kind of insert backdoors or things like that. In any system, like those type of risks exists. Obviously, we try our best to avoid these type of things. This is why, like just architecturally, we keep the data on the local device. We don't want to become this honeypot.

Again, I don't want a situation where it's like somebody will kidnap my kids and then say, "Give us this data." And then I want to be able to legitimately tell them, "Okay, unless you have a person's device, there's nothing I can do. Architecturally, we've architected the system in a way that that's impossible." That's what we strive to do like in terms of the architecture of the system obviously – And things for operational security, etc., we do that. I think there certainly are attacks that people can do that I think are within the realm of possibility that we've been replicated. But I would not be surprised if a kind of well-resourced enough adversary would be able to figure out some ways around some of the things that we do.

Again, because the things are passive, then in essence, like zero cost. I mean, it doesn't cost anything in terms of your experience. So you can layer on and do as many of these things as you want. Because you're not requiring the user to retrain or change their behavior in anyway, then it becomes much easier to deploy these types of things without a lot of friction and then take advantage of multiple of them at the same time.

[00:44:03] JM: Any predictions for the future of biometric authentication?

[00:44:06] JW: Yeah. I mean, I think – Look. In the future, I think authentication is going to change. Even within the next 3 to 4 years, the password alone will no longer be the predominant method of authentication. And I would also venture to say that the current – Even the current forms of two FA will not be the predominant forms of two FA even within a few years if you look – The status quo for today for authentication is a password alone. More and more services are starting to implement two FA, but they implement two FA by doing like I'll send you an SMS and then you type in a six-digit code. And then that's how you know that it's me. And there are a large number of security issues with that approach just in terms of the SMS protocol. The SS7

protocol is just woefully insecure. Like problems with like SIM swapping or kind of cloning, number porting attacks, these types of things.

Basically, the situation is right now, is if you get access to somebody's phone number, you can probably get access to just about every other part of their digital identity, because from the phone number, you can usually get into their Gmail account or their core email account. From the email account, they can basically reset and get access to everything. That's the status quo. There can be more and more attacks against that. And NIST and the FBI and others have come out and said, “Do not use SMS for two factor.”

The truth is like that's the only ones that's really ubiquitous today, which is why it's so commonly used today. And like the argument is better than nothing. Well, okay, sure, it's better than nothing, but we can also do much, much better both for authentication on the device. Many devices can include biometric. I think you're going to see many more things where authentication is considered being I type in a password. I’m going to get a notification in my phone. I’m going to do some type of fingerprint or face ID or other kind of auth on the phone to authenticate my transaction. And then that's going to be the direction that things are going to be going, right?

You’re going to see many more passive biometric as well. The number of sensors in people's lives is set to explode. If there will be billions and tens of billions of sensors in the world just associated with everything. And the way that people are going authenticate them is going to change. It’s not going to be let me type in my four digit pin, or even let me have this piece of plastic where I have a mag stripe. I’m going to swipe it, or wave it, or those types of things. It's going to be like a lot of this type of authentication has become much more passive and much more natural interactions, right?

I mean, I’m going to walk into my house. My house will recognize me based on the fact I'm carrying my phone. Based on potentially voice or facial recognition. There’s a number of other signals that they can use. But I’m going to walk into my house. It’s going to know that it’s me. It’s not like I have to go and type in my four digit code to like unlock my alarm or anything like that, right? Likewise for my car, likewise for brick-and-mortar stores that I interact with. Likewise for travel, and not just like at the airport, but just the entire travel journey. Going across when I kind

of go, if I step into my autonomous vehicle or like a rideshare vehicle. That authentication is going to be much more seamless in those cases. Hotels, rental cars, restaurants, etc. All these things are going to be just much more natural and, honestly, much more human. I mean, the way that people are going to identify each other. And the technology is going to become more of an enabler for this type of use cases.

And just because like the way that people do authentication today is very contrived and it doesn't really serve the purpose that it needs to. It's hit a scalability limit. I mean, the number of accounts that people have to keep track of and authenticate to, the notion of like the private data and like that the fact that private data is increasingly being compromised, and the fact it can't be relied on for authentication. I think in many cases, because technology and sensors and machine learning have reached the point now that you can do many of these things more passively, I think that you're going to see a transformation in the product experiences. Experiences with interacting with services, where many of these things just come much more natural and much more human interactions rather than something like, "Okay, what is your password? Okay, what is your mother's maiden name? Okay, swipe this card," like those type of interactions.

[00:48:21] JM: All right, John. Well, thanks for painting us a picture of the future of authentication. And very exciting company.

[00:48:27] JW: Great! Yeah, thanks for having me on. And yeah, I'm very excited to be here.

[END]