Sadie Witkowski: So we're both of the generation who grew up right with the birth of the internet.

Ian Martin: Yeah, we're both millennials.

SW: So how often did your parents warn you about sharing information online? Or, like, talk about digital privacy

IM: Oh actually I don't remember my parents ever having that talk with me. Is that bad?

SW: I don't really remember the first time I had "the talk" with my parents. I think I've always been a bit leery of strangers on the internet knowing too much about me. But it turns out, that's probably not the aspect of digital privacy that I need to be most concerned with now.

IM: But before we get into that, introductions?

SW: Sure! I'm Sadie Witkowski

IM: and I'm Ian Martin

SW: And you're listening to Carry the Two, a podcast from the Institute for Mathematical and Statistical Innovation, AKA IMSI.

IM: This is the podcast where Sadie and I talk about the real world applications of mathematical and statistical research.

SW: We might seem like an odd couple to tackle these topics. I'm a cognitive neuroscientist and Ian is a high school choir teacher!

IM: But thankfully, you don't need a degree in mathematics or statistics to learn how to apply them to the world around you. Or to use them to help you evade Ultron!

SW: [laughs] Well we're clearly not in the Avengers slash Marvel universe so maybe we're safe?

Heather Zheng [HZ]: 30:25
usually my job is driving people a little bit [crazy] after talking to me, [they] say that I'm gonna turn off all my IoT devices at home, I'm sorry. [laughs]

SW: To clarify, IoT stands for internet of things. Ie, all your wifi enabled devices.

IM: And who was that?

SW: You just heard Heather Zheng, a professor in the computer science department at the University of Chicago! Heather is one of the lead investigators in the SAND lab, that's security algorithms networks and data, and she's going to walk us through how we can reclaim control of our data and our privacy!

IM: Because when it comes to the internet, I think we're all aware how easy it is to find all sorts of information that maybe we don't want public…

HZ: 38:23
if somebody gets targeted, it's it's very easy to actually get this information.

IM: Ok, we haven't even started this episode and I'm already starting to feel a bit paranoid.

SW: I know, right? In the digital age, I feel like everything we do is tracked. We're always trying to protect ourselves from data harvesting, especially of the sort that can reveal personally identifiable information. Thankfully, we have researchers like Heather to help.

HZ: 39:47
I think that one of the misconceptions is possibly individual user cannot do much about it, and we have to wait for the big shots to change the landscape. Our work is always trying to say you can, there are ways to help build up the defense, but you need to start, be aware of it. And you can apply some of these defenses. And it is feasible. And we're working out that for various different kinds of attacks. But firstly, we have to be able to raise or be good people that are aware of potential attack.

IM: So is Heather talking about phishing scams and how to avoid them?

SW: Not quite. Though we are talking about protecting digital privacy. More specifically, Heather's research group focuses on both potential and actual threats to our data and ways that we can, as individuals, protect ourselves.

IM: Doesn't the EU do this already for its citizens? Wouldn't it make more sense or this to come from regulation? Kind of like the way I feel about dealing with climate change.

SW: Neither Heather nor I disagree with that. But with the breakneck pace that technology moves, it's been hard to properly regulate and in the meantime, our data is getting snapped up by all sorts of potentially malicious actors.

IM: So Heather is trying to help prevent that?

SW: Yup, she wants to provide immediate, concrete tools that can help folks now, while we urge regulators to take these issues on at a more global scale. And while no individual solution will be perfect, it's still better than nothing.

HZ: 13:18
defense is always about raising the cost of the attack.

IM: Oooooo, I really like that! Ok, so if she's not looking at protecting from phishing scams, what is some of the work she's doing?

SW: So the research project that first caught my attention was something she presented here at IMSI for a workshop called Private AI, Machine Learning on Encrypted Data. And it's named after a rather famous person…

HZ: 3:15
Fawkes is essentially how users can help prevent or reduce these third party or unauthorized third party who use their uploaded image in the online platform to run facial recognition against them in reality, in physical or, or in even the online without their consent.

IM: Fox? Are we talking about the Wes Anderson animated whimsy that is the Incredible Mr Fox?

SW: [laughs] No, not quite. More like, "Remember remember the 5th of November…"

IM: "the gunpowder, treason and plot!" So we're referring to Guy Fawkes here.

SW: Right, who was a historical figure, a Catholic who tried to blow up the British parliament building, though nowadays we mostly recognize the name from the V for Vendetta masks. And that's more of the reference Heather is going for here.

IM: The masks? Oooooooo, because she mentioned online images!

SW: Exactly. The Fawkes project is about protecting your image, literally, online from being harvested and used for facial recognition software.

HZ: 3:39
So the general thing is that there are multiple companies like clear world.ai, what they do is they build third party facial recognition software for you know, for profits. So, but how do they get your facial templates, because that's exactly like your, like your fingerprint, like your voice, and they have to get your original data. So in general, what these companies do is they go online, and they either go to your social profile, or anywhere you post the or somebody else, post your image. And usually that's labeled with your like, this is how their job or this is. So they basically label you and then gather a bunch of these data. And that becomes the training data they use to identify you, or do the template to identify you

IM: Oh my god what? And they're just like doing this to random people? I just can't believe this!

SW: Yeah, this is a clear example of using what you post for data mining, in this case for image recognition. And I guess we could say that a person could just choose to never post images of themselves or their loved ones online. But c'mon, really? How many people do you know who don't have an instagram or facebook account?

HZ: 4:28
This is very easy, because before they had to take your photo, you had to go to the DMV to take a photo now it's very easy, just view the script, and they basically just scope many things out of the internet. So that's the one and and and so, you know, so when we post our life, although we had to be really cautious, right? But it's very difficult to either post, don't worry about it. Or you just say oh, I'm not gonna post anything at all. I don't post anything about my kids or my family, but not everybody would want to do that, that that basic they ki-, kill the social interaction. So we're trying to do is what users can do in this regard. You can wait for the law to regulate these companies industries, or is there anything that user can do. So what we're trying to do is using kind of AI to or ml to defend against a ML.

IM: So how does this Fawkes thing help protect my beauteous selfies from being turned against me?

SW: You know the podcast format by now, [laughs]. Obviously, with math! But specifically, Heather uses machine learning.

HZ: 5:25
So there's in adversary machine learning, there's a certain behavior, we call it kind of a perturbation, or you add some kind of a perturbation onto image and that perturbation, it's very visually small or not even affecting the visual view. And yet it creating the, such changes in the feature space that alternate how the machine look at the features extract related to identity. So we try we build Fawkes, Fawkes, essentially a kind of, we call the poison or the cloaking methods. The cloaking is essentially, when you upload your image to the internet, you actually add some kind of imperceptible or some minor

cloaking, or perturbation in the pixel level. And that pixel level perturbation, basically moved your identity feature vector in the feature space. So when these companies use these, to build their facial recognition, they will generate a template about your identity, which is drastically different from what you have in reality. So now, I basically go online or I walk on the street, they take a photo of me, which will not have these perturbations, right? Because I'm not wearing them. And then when this is happening, they will recognize me as a different person.

IM: Soo…… Fawkes is adversarial machine learning?

SW: Right, which Heather sometimes abbreviates as ML.

IM: And these perturbations that she's making to the images are, like tiny pixel adjustments? She's not adding beauty marks or straightening noses or something, right?

SW: No, that's what snapchat filters are for. [laughs]. But basically by using Fawkes, you're adding this layer on top of the image that makes tiny changes to the face that our human eyes wouldn't notice, but will totally throw a computer off.

HZ: 9:05
when people have a perfect face or facial, it's much harder to avoid, you know?

IM: So obviously this wouldn't work on Beyonce.

SW: [laughs] or Jensen Ackles.

IM: No one knows who that is. [laughs]

SW: But there are ways to hide these slight alterations in most images.

HZ: 9:05 [part 2]
Ideally, if you have as your aging you know, if you have any wrinkles or anything on your face is a natural one like age spot or whatever. It's actually perfect to actually hide all these perturbations into what.. I usually have folks suggesting, Can you just make that as a makeup? Yes, you can beautify yourself and adding the perturbation along that line, it just needs a little bit more compute, I mean, a lot more computation optimization for say. So that's all doable as long as this is actually match your expectation, and you can choose to optimize that.

IM: Man, this seems like such an obvious, perfect tool that I never realized I needed! I can both beautify myself (not that I need it) and protect my identity online? Yes please!

SW: I mean, I don't want to make any sweeping statements. Fawkes isn't perfect and Heather can't exactly test it against every potential facial recognition software ever, because they're constantly evolving. It's a bit of an arms race that way.

HZ: 11:07
We tested it across multiple systems, including the open source one. Also, the the online version, like from multiple platforms, and it seems to work very well. However, we cannot exactly test exactly, it at every specific platform. And I think that's the key thing. And it's much harder to deal with, because there's so many instances of these testings, so we tested it on multiple commercial platforms, it seems to works very well. But it's the other caveats of this is you cannot really guarantee it's not like you have a sense of the guarantee about these things, and all the models evolve as well.

IM: I mean, better than nothing. I'll take it!

SW: So I've clearly convinced you of the utility of Fawkes, but I think we should talk a bit about how it actually works.

IM: Oh you mean the point of this podcast? The math behind the cool work?

SW: [laughs] Ok smart guy, why don't you explain what a feature vector is?

IM: Uh….. Why don't we let Heather do that?

SW: I thought so.

HZ: 16:09
what we're trying to do is basically move your identity vector feature vector in the feature space to a different place that will not correspond into your true identity. So, but how far to move, it really depends on where we can hide these visual perturbations without affecting the visual person of your, of your whole face. And there can be different levels of kind of different protection and different perturbation around that as well.

IM: Right, vector space… [pause] Nope, still lost.

SW: I know, this is a little non-intuitive at first, unless you have a background in machine learning.

HZ: 16:58
the general aspect of machine learning how data is getting translated from the input space to the feature space, and how that this is a nonlinear effect. And then that's why you can actually generate a lot of these subtle move in one direction, but with a little perturbation at the input. So a lot of like ML basic stats is very useful. The maths is definitely definitely important. In all the research we actually done right now,

IM: So…. feature space isn't a one-to-one of my facial features.

SW: No, it's not. It's not that your nose is one feature and your eyes are another. Feature space in machine learning refers to the collection of variables that you're interested in

HZ: 17:35
feature space is basically a high dimensional space that translates your complex visual of your face. That's to say, into a series of data points where the machine can understand the pattern, and then use this pattern to map back to the input data. So it's kind of a reduced representation about your specific image, or I suppose a piece of input data in the mathematical form, which captured the essence of the input data.

IM: So my nose isn't a feature, but maybe something like the distance between my eyebrows?

SW: Yeah, it's like these much more small pieces that make up the face as a whole. And even more importantly, the features are the points we use to tell the difference between individuals faces.

HZ: 18:30
Mathematical features, its the patterns that separating maximum separating the two of us. Yeah, although we have the same amount of eyes, and nose, and everything, but there are all different kinds of features along there.
18:48
but you have to make it the right combination in order to change the feature space, so the mapping is very nonlinear. That's why you had to do the optimization search, which can be done using mathematical formula, in this case, and leveraging that computer resource you have. it's all doable.

IM: You had me, and then you lost me.

SW: So, once we have this list of features for photos that build the feature space, we have to figure out what tweaks we want to make to that space, right?

IM: These are the changes or tweaks that make my face unrecognizable to a computer program?

SW: Right, but then we still want our human eyes, ie our friends on instagram, to still recognize our face as our own. So we have to use an optimization search to figure out what big changes we can make in feature space that will minimally impact our images when it comes to our eyes.

HZ: 19:47

Because there's a lot of nonlinear relationship between the input and output that gives you this type of big, small manipulation of the input, which generates a significant shift into the output side of it.

SW: So basically, because machine learning for visual scenes is based on this feature space, looking at pixel changes rather than the holistic way that we view images, Fawkes can take advantage of the non-linearity, the non one-to-one comparison between feature space and what we see, to make these changes that we can't detect.

IM: That's so cool! So the tens of dozens of followers on my private account will still recognize me.

SW: Exactly! And I should mention that you can run instances of Fawkes both across networks of multiple people's images, which is actually the stronger way to cloak images, or on your individual machine.

IM: Where can I get my Guy Fawkes mask?

SW: We'll link to Heather's webpage so that you can download your own instance of it!

HZ: 12:19
So what we're trying to do is say, not completely stopping everybody, I mean, third party from, you know, recognize us, but just make it much harder for them. But also give user immediate react.. I mean action item they can do just to put this on them on their photos when they upload, and then just make it harder for those third party to recognize us in reality, I think that's it's delay. It's a years of a delay that matters as well.

SW: It's not perfect, but Fawkes helps make you a harder target to collect data on.

IM: Wait, but at the beginning of this episode, you said this was her first project…

SW: You're right, we actually have two other applications of machine learning to data privacy that we haven't even talked about yet!

HZ: 20:43
So this Fawkes is essentially a defense against a current kind of a privacy threats that everybody's facing. And the other two are basic, exposing the potential attacks that could be launched against us in terms of privacy. And these are basically not intuitive.

SW: Buuuuuuut, we'll get to those after a short break

[music starts]

[Ad read - Capitalisn't]

[music fades]

IM: So now that we've dealt with an example of real-world data privacy and protection, Heather mentioned potential attacks…?

SW: Right, these next two projects her group worked on are more focused on preventing ways that we could potentially lose privacy. And of course, ways that we as individuals can prevent that.

IM: So what's up first?

SW: Well, imagine if someone could track your movements within your own home, whether that be an apartment or house, just by dropping a passive receiver, a WiFi sniffer, within range of your wifi network.

IM: Wait, they don't even need my wifi password. Which obviously is password1234

SW: [laughs] Nope, and you would literally have no way of knowing unless you discovered the bug, because it's totally passive and doesn't affect your network or devices at all.

HZ: 21:23
So we're just trying to look at, you know, if because, oh, these Wi Fi devices, transmit some signals. And they do transmit signals over time, can a external observer, just observe the fluctuation of these signals, and use that to infer your activity or your presence in the home

IM: I would love it if you said this was all the paranoid delusions of my mind and that it's totally not possible.

SW: Wouldn't we all… Sadly, if you have a malevolent actor who wanted to track you this way, they pretty easily could.

HZ: 22:52
It's completely passive, it doesn't need to send data. So that's why if anything is passive, you can drop a wireless LAN, which does not send data you can just the observe right? Anything on the street, anything can put in your your garden, or a sniffer which lasts for a few days. And then they basically record your daily patterns over time. I think that's the that's the and it's very difficult because it doesn't transmit anything. So that's one it just listen, it does not control the signal. Just listen and look use my simple

analysis, we look at you know we did was self, um calibrated, and remove the noise effect everything and then just identify which room you're at a particular time.

IM: Wait, but how is this bad actor able to tell where I am in my apartment using wifi?

SW: Well, the key is having multiple devices on your wifi network across the different rooms. The sniffer collects that information and….

HZ: 22:07
And if they can analyze that data, and then look at the perturbation within the that data over time, and they can know when you're at a particular room, let's say in the morning, you got out, you go to the bathroom, or you go to the dining room, you go to the kitchen, they know exactly which room you you go a particular instance. So that's basically they know when you enter your home which room you're going to. And that's kind of scary effect where and the more Wi Fi devices you putting into the home, the more precise information they will obtain.

IM: I'm still not getting how this works exactly.

SW: Think of it like echolocation, that thing that bats do to navigate the world. Every device connected to your network is constantly sending out a signal, which is deflected or weakened when it runs into stuff.

IM: Like how I might need a second device in my home office if the router didn't reach there?

SW: Exactly! So walls and furniture affect the signal, and so does our physical body. So if you record the signals over time, you can spot the fluctuations that represent you moving through your space.

IM: That is SO creepy.

SW: And it just needs to monitor the signal changes over time.

HZ: 23:58
It doesn't need to know the internal structure. But the floorplan, it helps tremendous as our but a floor plan, you can actually get it from a lot of real estate websites anyway.

IM: Congrats, you have thoroughly freaked me out. I will be unplugging everything as soon as I get home.

SW: Yeah, sadly that's not a great option for folks though, right? We need wifi to log into work if we're working from home. Or to play music or turn on smart lights or whatever else people use the internet of things for.

IM: Buuuuuuuut, I’m assuming based on Heather’s orientation towards protecting our data that she has a solution?

SW: Indeed! Basically her group has proposed a AP-based obfuscation.

IM: Ah yes, the Advanced Placement-based obfuscation test. I hear it's a real rough one.

SW: [laughs] AP in this case means ‘access point’ for your wifi. Basically, it’s the same idea as Fawkes. You can add noise or a cloak over the original signal that hides these patterns.

HZ: 24:19
And we're trying to say, your AP, your access point can actually help you by transmit some kind of obstruction or obfuscating signals. So it's mostly adding perturbation just make it harder for the attacker to extract information easily. That either needs to have a multiple antenna. So you basically just make it much harder for them to deploy without you knowing it's there. Or they just need a much more sophisticated devices. It's much harder for them. So it's all about raising their cost.

IM: So you can add a noise mask over the signal and that makes it harder to detect the true signal?

SW: Bingo.

IM: Well thank you for taking me on that journey of freaking out to feeling like I can at least do *something* about it.

SW: Are you ready for me to do it to you one more time?

IM: [sigh]If you must.

HZ: 24:48
The keyboard project it's more recent, it's kind of a little bit more scary. As because looking at the personal privacy, our voice, our, our face are all private. But also your data that you type on the machine is also very private. But now if you have people work in public all the time work in the, you know, on the train on the airport, everything is basically, you mix these public spaces with work together, right. But now, if you are typing your sensitive information on a keyboard, right, it's like I'm doing this right, and I'm typing on a keyboard. And you by just recording a video of me typing here. And you don't even need to know my keyboard. And I can after you type about 10 minutes or 500 words, I know exactly what you type, what I know 90% of what you typed,

IM: Woah woah woah, back up.

SW: Right, so this project is perhaps a more immediate fear or concern than the wifi one. Heather's research group is trying to figure out how to protect people's private writing that they're doing in public spaces.

IM: This feels very James Bond.

SW: Oh yeah, it feels like classic spy-craft. In this study, Heather showed that if you're typing away on an ipad or a keyboard and someone can see your fingers, they can record 5 to 10 minutes of your finger movements and then run an algorithm for about 40 minutes and recover what you typed with really good accuracy.

So if you're a government employee or someone high up in a corporation, this kind of espionage bypasses any protections you have from an encrypted wifi network or a screen protector to keep people from reading over your shoulder.

HZ: 25:48
It can be any keyboard, I don't need to you need your Keyboard Layout, I don't need to know have your training data about which one key you type. I don't need to know the content your type, I only need to know probably your typing English. That's all we need to know nothing else about. It's a first encounter attack, I don't need any information about you. I just need to have an RGB camera in front of you and your typing.

IM: How the heck are they able to do this? I tend to be a sloppy typist so how could they tell what I'm tapping out?

SW: Well it gets back to Heather's work in machine learning.

HZ: 28:00
we use a language model, the whatever the language model have some kind of frequency and things about that it's a little bit complex. But also use the language model helps you build the first layer of the label. And then they help you get an initial label to at least give you some sense about these labels, but many of them are wrong, filter them and identify and retrain the and then train the these models to actually self correct some of these errors, it actually works pretty well.

IM: So kind of like the frequency predictions we talked about in Allyson's episode on GPT-3?

SW: Right, these are prediction models which combine the mathematics of the program to produce statistical predictions about what, in the case of GPT-3, makes the most

sense linguistically. And then in the case here, statistical predictions taking the patterns you typed and figuring out what words you were writing.

IM: [stage whisper] If you missed Allyson's episode, then you can find a link to it in the show notes [stage whisper]

SW: And just like GPT-3 used the layers in the neural net to build complexity, in this keyboard study Heather used a similar idea.

HZ: 31:12
it's called the hidden Markov model. So it has a little bit more beyond the frequency, it also uses a transition kind of probability to maximize the ones. So you basically build a initial clustering based on how your fingers move where you touch it. But oh, these are basic estimate from my low level low dimensional analysis, and you use this to give a rough guess. And this rough guess, is basically use... And we actually apply the filter. So the filter is a little bit more complex, because we're trying to look at consistency, look at many things across multiple layer, and just to eliminate potential noise, or at least eliminate the labels that could be wrong or uncertainty. So you use less number of labels with reasonable accuracy, and then you apply additional level of noisy training. And then that helps us build a much better dnn model to recover the contents.

IM: Sadie, I thought we were investing in a jargon alarm or buzzer or something?

SW: [laughs] I'll put it in our budget for the next season. What's tripping you up?

IM: Markov model? Clustering? Filtering? [frustrated sigh] What's happening here that allows this hidden markov model thing to figure out what I was typing?

SW: Well when you boil it all down…

HZ: 32:27
it's all about, you know, finding patterns.

SW: But going back to the jargon. A hidden Markov model is used to predict a sequence or series of unknown, aka hidden, variables, based on observable variables. So in this case, we're trying to predict the hidden button presses-slash-words typed based on the observable movements of a person's fingers.

The cool thing about this, is that you don't need any prior knowledge. You don't need to pre-train the model, because it's only learning based on the finger movements you give it.

HZ: 32:34

we had to deal with noise, we had to deal with that you don't have training data, you don't. And the user behavior varies, right. So I typed and you type very differently. And also, our pattern may not be fully consistent all the time. So so that's another level of noise we have to deal with.

SW: So the model has to deal with noise in the system, which is why it makes clustered estimations of where the fingers are pressing as a first pass, and then you start adding in some other filtering to refine those guesses.

IM: So the clustered estimation might be that my finger actually pressed ‘d’, but the model will say it could also be ‘s’ or ‘e’ that I hit. But it’s really unlikely that that finger tapped ‘k.’

SW: Exactly, at least that’s the case on a qwerty keyboard.

IM: So this hidden Markov model algorithm just needs five minutes of my typing and then 40 minutes to chug along before you’ll know what I typed?

SW: I mean, you need at least 500 ish words typed so the amount of observed time varies for different speeds of typists…. But yeah. That’s the idea in a nutshell.

Of course, this algorithm would be really bad at recovering code, for example, if you were writing in python or R.

IM: Because code has different frequencies of letter or word use?

SW: [laughs] That’s kind of true. But actually, it’s more because it’s hard to write in code fluently like we do for human languages. It takes people, on average, much longer to type 1000 characters in a computer script than it does to write an email.

IM: Wow, that was not the answer I was expecting…

SW: So have I thoroughly convinced you that a spy with an iphone camera could record you for under ten minutes and recover all the terrible fanfiction you were writing at your local coffee shop?

IM: Pfft, Sounds like someone is projecting… But yes, point taken. I see the dangers posed here. Now let’s hear the solution!

HZ: 34:52

I think the immediate thing is if you go to an airport or go to a public space People do this all the time. Right? Yeah. Do you see more people typing on iPads and many of these portable keyboards all the time without any blockage

IM: Blockage?

SW: So the key weakness here is using something like an ipad, where the potential spy can clearly record your finger movements.

HZ: 36:01
I think if you use a laptop, maybe you're safe, because your… The screen actually provides pretty good privacy. But if you're not using a type, you're so good at typing on iPads, which does not have a blockage. That's what we're saying that you need to be aware of this. Because somebody can just pretend to watch a video because what we're showing is somebody who's pretend to watch a video and they record you.

SW: So moral of the story for this one is to use a laptop in public instead of an ipad.

IM: Or bring those manila folders they used to make us stand up to prevent cheating by people glancing at your scantron in grade school.

SW: Exactly!

IM: For a mathy podcast, this is quite a low tech solution, but I like it!

[short music intermission - 15 sec or less]

SW: So, did I live up to my promise? Were you made paranoid and then re-assured that there are things you can do to protect your digital privacy?

IM: Yes. It has been quite the rollercoaster.

SW: I just love how mathematics can be used as the solution, as we saw in Fawkes. Or mathematics are used to reveal potential dangers to our privacy, in the adversarial manner, as Heather describes it.

HZ: 42:38
We're not reinventing these maths we're applying or understanding the implication of doing so whether it's from the machine sides, or from the attack side, or it's from the different sides. They're all involved, you know, what you can do to change the pattern. And what you can do to and subject to certain constraint, whether it's perturbation size, whether is where you use it, how you use it, whether it's a physical barrier versus a

digital edits of your photo, all these are basically what you can do to change the pattern of the of what can be used, but for data analysis, and how do you do that.

IM: Math can be used for good or for evil!

SW: But regardless, we have to be aware of these potential risks and dangers before we can know how to counteract them.

HZ: 41:06
I think the user, that individual user has agency to protect themselves, or has the capability to protect themselves, or at least raise the attack cost is the key thing. We're trying to say here.

IM: So while we should really have governmental powers regulating a lot of this and helping protect folks…

SW: In the meantime, we can use math to proactively protect ourselves!

[Music starts]

SW: And don't forget to check out our show notes in the podcast description for more on adversarial machine learning research and further resources on the mathematics and statistics behind data privacy.

IM: And if you like the show, give us a review on apple podcast or spotify or wherever you listen. By rating and reviewing the show, you really help us spread the word about Carry the Two so that other listeners can discover us.

SW: And for more on the math research being shared at IMSI, be sure to check us out online at our homepage: IMSI dot institute. We're also on twitter at IMSI underscore institute, as well as instagram at IMSI dot institute! That's IMSI, spelled I M S I.

IM: And do you have a burning math question? Maybe you have an idea for a story on how mathematics and statistics connect with the world around us. Send us an email with your idea!

SW: You can send your feedback, ideas, and more to sadiewit AT IMSI dot institute. That's S A D I E W I T at I M S I dot institute.

IM: We'd also like to thank our audio engineer, Tyler Damme for his production on the show. And music is from Blue Dot Sessions.

SW: Lastly, Carry the Two is made possible by the Institute for Mathematical and Statistical Innovation, located on the gorgeous campus of the University of Chicago. We are supported by the National Science Foundation and the University of Chicago.