Virtual Keynote | Trustworthy Computing (*45 min*)

Jerry Cuomo, Global Industry Leader, Retired IBM Fellow and Distinguished Research Professor, North Carolina State University

Keynote Description

This talk introduces Trustworthy Computing as a shift toward systems built on verifiable, transparent, and privacy-respecting foundations. Using blockchain and AI, it explores real-world impacts—from food safety and identity protection to AI transparency and machine unlearning—emphasizing the transition from reputational trust to algorithmic trust in everyday digital life. Health care and Life Sciences professionals will find particular relevance in applications that enhance drug authenticity, patient data privacy, and AI reliability in decision making.

AI Disclosure

This document was transcribed entirely from the keynote video using automated AI transcription tools

TRANSCRIPT

Hi, Jerry here. Well, actually, I'm a clone of Jerry's voice, starting off by saying, thanks for the invite to speak at Converge to accelerate.

My talk today introduces trustworthy computing as a shift towards systems built on verifiable, transparent, and privacy-respecting foundations.

Using blockchain and AI, it explores real-world impacts, from food safety and identity protection to AI transparency and machine unlearning, emphasizing the transition from reputational trust to algorithmic trust in everyday digital life.

I'm hoping that healthcare and life sciences professionals will find particular relevance in applications that enhance drug authenticity, patient data privacy, and AI reliability in decision-making.

Enjoy the talk, and now over to Jerry. Hey, the real Jerry here. A little bit about myself. I'm a retired IBM fellow, and I recently founded a company called Wild Ducks LLC, where it allows me to do a little writing and consulting stuff.

I also teach at North Carolina State University on this very topic of thinking about trustworthy computing.

And I have a couple of books, Think AI and Think Blockchain, to kind of complement that, plus a podcast, of course, called Wild Ducks.

So I hope you enjoy this presentation. When you think about trust, we always fall back on what has served us humans for quite some time, and that's reputational trust.

And as we all know, reputational trust is building trust through established relationships and long-term and predictable behavior or reputation.

And with that comes at least a perceived reduction in risk based on social proofs. You know, I've been around this person long enough. I can kind of predict their emotions.

They've consistently acted this way. you know, that gives me a sense of trust. And, you know, we've, as people, lived by that trust for quite some time.

And I'm not suggesting that we're going to stop doing that. That's a critical part of trustworthy computing. Companies we trust, people within those companies, et cetera, your social impression, and all of these other ways these days that we can verify a person's relationship and trust between I would like to say there's a little bit more that we can do today, and that's pair this or complement this with algorithmic trust.

So what this does is it quantifies and provides an additional level of trust through mathematics and mathematical certainty.

In short, minimizing risk, as we're going to see through cryptographic proofs. And a lot of this presentation is focused on this latter part, algorithmic trust.

We know about reputational trust, that's not going away, but again, to really build trustworthy computing, let's try to pair reputation with a little algorithmic trust, and we get a beautiful thing that'll occur when you bring these two things together.

So, the first technological ingredient is this thing we all call blockchain, and many associate blockchain with Bitcoin.

And you wouldn't be wrong if you do that. Obviously, the Bitcoin network is powered by a type of blockchain. But blockchain is more than that. I tend to look at blockchain as much as a technology than a social movement or anything like that.

I see it as some really cool cryptography and distributed networking and things of that nature. So I see it as an example as a platform for building trustworthy computing.

And I want to kind of put a little bit more behind that and maybe get a little dramatic and say, would you believe me if I said blockchain is poised to change everyday life for good?

I mean, forever and for the benefit of good. And while we may not always need every element of blockchain, what I'm saying is the ingredients that go into blockchain, are the right ingredients for building this thing that I like to call algorithmic trust.

So I'm a bit of a storyteller. So let me tell you three stories that to this day, I've been doing talks on this topic for about a decade.

And these three examples continue to resonate and inspire me. And so much has changed in this decade. But these examples and also the networks that live behind these examples continue to endure.

So I'll start by saying, has this ever happened to you? And this is an impressionistic view of a trustworthy sandwich, or maybe not so trustworthy.

Have you ever been maybe traveling through an airport and picked up a sandwich before jumping on an airplane?

And then halfway through the flight, you look like this guy here. And maybe there's some kind of foodborne illness that you just got.

I don't know about you, but back in the day, In 2006, I remember spinach taking a bad rap, no pun intended, but bagged spinach was found to be the cause of an E.

coli breakout. And it took regulators two weeks to conduct the trace back and determine the exact source of the outbreak.

Now, in two weeks, a lot of bad stuff happened to people who really got sick and some of them pretty seriously sick.

And then also to the spinach itself, because we didn't know where the good and the bad spinach was.

So we just had to eradicate spinach. So the market for some years later was kind of spinach-less, if you want to say that.

But I like to report that everyday life has been changing now for the better. Around this notion of the food trust network, it started way back when with a kind of consortium with the likes that you see up here.

and I'll name two in particular, IBM and Walmart. And they put together this food trust network with the idea to quickly pinpoint the sources of contamination and to reduce the impact of food recalls, and of course, limit the number of people who get sick or die from such foodborne illnesses.

Let me tell you a little bit about how it works. Using blockchain, we're gonna trace the provenance of the ingredients as they travel from farm to fork.

So think about the sandwich. It has chicken in it, maybe tomatoes, mayonnaise, and lettuce, and maybe we're all fixed that it's the chicken that caused the illness, but maybe it's the tomatoes or the lettuce, right?

So Walmart and IBM started looking at this problem, and they did an A-B test. And the mango, packaged mangoes were picked by Walmart as the source of this A-B test.

They went into a Walmart store, they went into the produce section, looked at a bag of mangoes and said, okay, let's figure out how long it's going to take us manually to trace the origin of this mango.

You know, with the thought that if there's something wrong with the mangoes, they can go back and only remove the mangoes from this farm or supply chain.

And it took seven days to do that trace back. But with the IBM blockchain platform at the time, Frank Giannis, who now I believe works for the Food and Drug Administration, but at the time was in charge of food safety at Walmart, said going from seven days to 2.2 seconds is a really big deal.

That's food traceability at the speed of thought. And I think that you will see is the first example as potentially this cryptography changing everyday life for the better.

Carrefour and Nestle, I guess, is something more concrete. And I used to carry this box of musseline, which is dried potatoes. And on it, there's like a QR code and you can see a mobile app.

And what you can do with this is kind of take a picture of the QR code with your phone and you can see kind of the trace back from the shelf that you picked that food off of to the farm.

Ultimately it came from. And you can look through from tracing back to what would that be? That would be like number four, the Carrefour manufacturers, the product arrived at 25 Carrefour warehouses, it was kept in storage, et cetera.

And then let's see, Number three, going backwards, the product is then stored in two Nestle warehouses located in this region of France.

And then, you know, going back in time even more, it went to the food processing plant, number two, and it looked at doing the packaging and, you know, kind of harvesting it from the crop to, you know, its powdered form, and then ultimately, number one, to the farm.

So tracing it all the way back to the 165 growers. So think about this traceback and having all of these disparate entities from the farmer to the warehouse folks kind of recording on this immutable ledger the farm-to-fork traceback of it.

So that's number one. Number two is my Aunt Tessie. Has this ever happened to you? Aunt Tessie recently had to fill out a application to rent an apartment.

Now, think about what year it is right now. And in this kind of simple act, she tells me she had to fill out a stack of papers.

Yes, papers, as in like things you write on. And in that, it felt like she was giving every piece of information about herself in doing so.

So, of course, the apartment office, but she had to go and get verification records for the bank. They even wanted a cell phone contract to show her address and that she had an active account there.

And then the Department of Motor Vehicles, which had your license, which was your kind of state record of where you lived, right?

So they wanted all of this information. I know I give out geez, I've been giving out information for generations now. Ever since the web was introduced, probably thousands of websites over the years, I've registered that.

I don't know where that information is. And neither did Aunt Tessie. And then we all get some of these on a bad day. Oops, your data has been breached. I hate when that happens.

And, you know, dear valued customer, we're very, very sorry. It won't happen again. And your information. So the dark web, while you may have forgotten where you put your digital information, the dark web hasn't.

So this is really the bane of digital existence, and that's privacy. And every year, Javelin Strategy and Research in 2023 was the last time I checked, and it was up to 15.4 million customers were hit with some kind of identity theft.

So that continues to be a big deal and just keeps growing up. So the next big problem that cryptography and these ingredients that go into blockchain can help solve is around digital identity.

And I had the pleasure a few years ago to work with a company in Canada called SecureKey. And they really did something interesting with cryptography in the form of proofs of identity and using trusted parties like your bank to verify, and the system is called Verified Me, to verify that you are indeed who you say you are.

So rather than sending information all over the interweb, you're keeping information with your state for your driver's license, maybe your bank, your primary bank, and then they become like your friends in a social network and they can vouch for you.

So you don't have to give the millennial partners your information, but all you do is you give them permission to run this proof with your bank and maybe with the state to verify that it's you.

And it gives them enough trust and confidence that those institutions are trustworthy by reputation that they will accept, right?

So that's a pretty big deal. And I think that is taking a run at protecting our identity. The way it works just a little bit is it avoids the honeypot, the big database with all of your identity in it with a big bullseye on it that says, hack me.

So there's no honeypots here because again, it's distributed through these trusted sources, these verifiers.

No tracking. It uses something that the National Standards Group calls triple blind data exchange, where the The person requesting the identity of the apartment doesn't have your information on exactly where you live, what your eye color is, and all those things that you might find on a driver's license.

So you're protected, you as the requester, to be verified. The people vouching for you, they are trusted parties, but The millennial apartments don't know who you're banking with, what state necessarily you have your driver's licenses, but it trusts that those sources are authentic.

And then third, the bank doesn't know that you're trying to rent an apartment. So triple blind, but somehow this works. So no unnecessary information. When I go in, I don't have to worry that now someone has my address that shouldn't have it.

that I didn't approve to have it. So I think you would agree that's example two of everyday life changing for the better.

And this last example does not involve me giving pills to little kids. This happened to me because I lived it. So I know it's a real thing. And that's supposed to be my son with a Taekwondo uniform on.

And the story goes is I used to be one of several parents that would cart the kids back and forth to their tournaments, their martial arts tournaments.

And yes, the kids would get banged up and inevitably someone would say, can I have an aspirin? And I would carry around this, you know, pill jar with aspirins and, and, and stuff like that.

And sure enough, you know, I would, I would give one or two of the kids aspirin to help their ailments.

And one time as I gave a child who asked, um, What I thought was a Tylenol, the parent grabbed it before the kid took it and looked at the pill and said, Jerry, what are you giving my kid?

And my heart dropped. I'm like, what can it be? You know, I have Tums and Advil and stuff like that in there. So I don't know what it was, but we didn't recognize the pill.

And we kind of Googled a number on the pill. We found out it was a generic version of Tylenol. From this point on, I always buy the real thing just because I want to see the real label on it.

But, you know, it scared the daylights out of me, you know, perhaps giving them a, you know, substandard piece of medicine or something I didn't think was really what it was.

But it's a problem. One in 10 medical products circulating in low and middle income countries, for example, are either substandard or falsified.

So said the World Health Organization and Things like cough syrups for children containing powerful opioids, fake antimalarial pills made of cornstarch and potatoes, right?

So, you know, this is a real big deal. But again, I'd like to say that, you know, life, digital life as we know it is moving toward a more trusted world thanks to this thing we call Crypto Anchor Verifier, something created by IBM.

And there's examples like this. Where again, this is not necessarily directly blockchain, but it's using identity proof technology in a very compelling way.

So for example, let's say I had two aspirins, one real and one falsified. What we use is a standard cell phone with a lens on it, a plastic lens and some software with AI, basic machine learning on it.

that creates a light spectrum analysis of this particular material, aspirin, versus known correct versions of this that may have been put on a ledger someplace and say, here's the real one, the spectrum.

And it creates a digital fingerprint of that spectrum. And you create a digital fingerprint of what you think is real. And we'll compare the two. And then if we do that to something that's false against something real, maybe the real one was taken at the factory, at the pharmaceutical company.

And then maybe at the drugstore, every once in a while, the pharmacist does a check on the shelf. They may see this, like, look, the spectrum is different.

And then digitally, you will be able to do that similarity comparison and see that, no, this is probably fake.

It's not matching up. And that was an abstraction, but here is an example. This was in our lab and a test case. So I'm taking a picture of the first aspirin.

There's the color distribution light spectrum. We do it again. And if we have the fake one, apparently, and then if we show the two, you can see one has one main peak.

That's the real one. And then this fake one has kind of two peaks in its color distribution. So it's a little bit off. So the spectrum analysis, you know, kind of flags that and says, hmm, suspicious here.

This might be a fake. So preventing counterfeiting. And again, this works for pharmaceuticals, things in healthcare, other life science materials, you know, plants, wine, olive oil.

We've tested this with all. And light spectrum is a form of fingerprinting. So if you can actually take that into a cryptographic fingerprint, You're well on your way now to being able to trust, to create digital trust to go along with reputational trust.

I trust the pharmacy. I trust the pharmaceutical company. I trust that sometimes malicious things can happen along the way, and that's where the digital trust is built, using these types of proofs.

So what is blockchain? Let me just break it down because I have a more I'd say computer science view of it. You know, I see blockchain right up there with the linked list as a very effective data structure and way to create a pipeline where ultimately you're gaining trust, more trust in the data through proofs and cryptography.

So, you know, databases are, you know, if you were thinking blockchain and database, you would, Wouldn't be right, but you wouldn't be that wrong either, right?

So it's a type of store that may be in a kind of Darwin genus thing. It's kind of in that same camp as NoSQL, SQL, like the type of database.

But unlike those databases, what really stands out to me is that most databases have a single administrator who sets up the rules for the ledger.

Blockchain has multiple administrators. that each have an exact copy. So it's one thing to compromise one single administrator who is setting up the database rules.

Now you got to kind of deal with multiple administrators. So you got to kind of somehow bias them and biasing a group is harder than biasing a single, right?

So that becomes the law of larger numbers. So sharing a ledger is one way to build more trust. consensus. So it's not just about distributing it. It's about, you know, looking at the transaction logs and looking at the consistency of doing things like proofs and looking at that and seeing that first, before we commit something, we have some level of agreement with the group that this is a fit transaction it's proposed.

And then let's mentally thinking about it as voted on, although that's not exactly the way it works.

consented on by the group. And then if some rule is met, let's say majority rules or everyone agrees, consents, then and only then is the new block added, right?

So there is that level of consent to go along with distribution of the shared ledger. Immutability, again, cryptography is used to create this append-only data structure and compromising the last element requires cryptography that would necessitate you going through the whole chain and decrypting the entire chain, which then starts to get into energy.

How much energy does any individual have to apply to this compute problem to go back and kind of reverse or decrypt the entire ledger?

So cryptography is used as a way to foil it, make a little bit more difficult, and in some cases, extremely difficult.

difficult to compromise the data. And it's append only. So with a database, an administrator could go in and delete a record or change a record.

This isn't append only ledger. So again, in order to rewrite the ledger, in order to change a value in the ledger, you basically have to rewrite and re-encrypt the entire ledger.

Cryptography is a very big deal. Again, techniques, to hide and scramble data. And we're gonna go over some of those in a second. That is also a really important ingredient.

So this kind of, when you look at these ingredients, consensus, immutability, distributed network, network proofs, like Merkle trees and Merkle proofs and basic cryptography, creating hashes and all of that, that starts to drive this data integrity and trust paired with reputational trust And now you have something that can help the digital economy operate more efficiently.

And with artificial intelligence being the thing right now, can we use these learnings and apply them, jump the tracks to enhance certain types of AI and allow AI to benefit from not just reputational trust, for the companies created it or lack thereof and bringing some algorithmic trust to AI.

So that's what the second half of the presentation is about. We just laid the foundation that blockchain has these ingredients. Now, can we apply these ingredients similar to how they were applied for the food safety and the digital identity and the, let's say, creating fingerprints out of physical goods?

Can we take that? And can we apply them to things that are prevailing in, let's say, untrustworthy AI to make it more trusty?

De-thinks, worrying about your privacy, like if you're prompting an AI model, is it training on the data that you're sending it?

Patent infringement. Hey, I have some really good things on GitHub that have license laws. licensing to it. How do I know that that license isn't being infringed by AI kind of homing the internet for information?

How can I make sure that that information I just got is real? And then if AI learns something that I prefer it not to learn, could it unlearn it?

So can we apply the techniques we just saw and targeted at some of these problems? So that's what we're going to do now. We're going to go through these basically one by one.

Let's start with deepfakes. I guess the best way to study a deep fake is through a deep fake. So I present you my deep fake. I'm sure you've heard that AI-generated deep fakes pose a great risk by creating convincing but fabricated audio and video content.

But with a little cryptography, deepfakes are being exposed. You see, digital fingerprints, derived from audio and video frequency spectra, verify content authenticity by comparing signatures against known genuine data, detecting deepfakes.

Pretty nice. So this is something I've been playing quite a bit with lately. Creating deepfakes to study deepfakes in a way to really create fingerprints of the real and proposed deepfakes.

samples. In this case, you're going to see my real audio juxtaposed against my fake audio. So let's take a listen in. And we're going to create a fingerprint based on these initially three.

In some of my examples, I have a dozen attributes, you know, fundamental frequency, frequency variation, frequency harmonics.

In this example, we're going to look at a fingerprint based just on these three things. And they've also been investing in guardrails that include a synthetic speech detector and audio watermarking.

Pretty cool stuff. That's the real me from one of my podcasts. So when you go through that, we do an analysis and hear these violin graphs that show how the real voice plots out.

So now let's now run one of my deep fake voices and see if we can visualize the difference. While our ears may not be able to Exactly. Although I kind of know my voice and I can hear little changes, but can we see those differences?

And they've also been investing in guardrails that include a synthetic speech detector and audio watermarking.

Pretty cool stuff. So now let's plot that out. And you can see immediately that the violins are different, right? For the fundamental frequency, it's a little shorter and fatter.

The other one doesn't have as many data points. The variation, it's kind of more around the norm, less spread out. And then the harmonicity is also not as distributed as my real voice is.

So from that, we can derive a set of finger or values. And then together with cryptography, we can kind of map these into our fingerprint.

All right, so that's the first one. Now let's look at privacy. The example of privacy is, you know, looking at user inputs to AI models, and they may have sensitive information, risking privacy breaches.

I think there was a bank or insurance company, or I forget, but they're putting information in, and I think one of the big large language models trained on it, or there's been some many of these in the press especially in the early days of Gen AI coming on board.

But the model here that I've been playing with is encrypted data and if we trained a AI neural net to be able to decrypt, take encrypted data, decrypt the data, come up with an answer and re-encrypt it.

So all that's ever traveling over the wire is encryption. So can we train an AI model, a generative model to think in terms of encryption and decryption.

So this is a very rudimentary example of doing that. And we're not going to use real encryption. We'll start off by using decoding, right? So this is Base64 decoding, define AI in 30 words.

So I'm going to take this output, which was the Base64 encrypted. And this is a tool that allows us to use large language models. I'm going to use Llama. I'm going to send this right to Llama.

a bunch of encrypted data. And look at this. I get back without even saying anything. I sent it an encrypted string. I got back an encrypted string. And if I decrypt it, it's my answer.

Artificial intelligence is a subfield of computer science. That's seamless. So I didn't even have to tell Lama what to do. It knew how to do it. It just saw it, recognized it, It knew it had to decrypt it, and it knew I wanted it decrypted back.

Pretty amazing, actually. So now I'm going to do this to open AIs. And you respond again in English. So I tried to make it say its response in English, and it said, no, I can't.

All responses are to be encoded in Base64. So in the system prompt, I kind of gave it this hint that I want to have a conversation, all in Base64.

So this starts to give you an idea. Imagine if the model was able to communicate almost like HTTPS. Everything is encrypted and decrypted. It's only inside the neural net is it being decrypted.

So if we can start to work with this, and I'm showing examples that show it's plausible. Privacy in healthcare, privacy in banking, privacy can be much more efficiently handled like we're used to handling it on the internet today.

So I think you would see that borrowing from some of the hashing and stuff that we saw from blockchain.

Machine-owned learnings. So AI models have been trained on unauthorized data, and that's an ethical and legal risk under regulations like GDPR, which mandate proper data usage and consent.

So if any of you are familiar with the Men in Black movie and the neuralyzer, I got the idea, can we zap a neural network?

Can we make it forget certain key phrases? And that's what this is. little tool that I built called a Neuralizer does. It's a machine unlearning process by removing data from a model, reversing its impact.

So let me tell you a little bit about the magic. And it all has to do with this notion of the carrot in the blender. If you listen to my podcast, I give this analogy some period of time.

So like a carrot blended into a smoothie, training data in AI becomes untraceable. you know, making it indistinguishable. So like, let's say you went to a party that they were making smoothies and you brought carrots and, you know, halfway through the party, someone says something that offends you and like, I'm going to take my carrot and leave, you know, pardon the silly example.

And then you go and then the host says, no, it's too late, Jerry. You know, I've already used your carrot in the smoothies, right? It's already in there. And I'm like, but I want it back.

And it's like, you can't have it back. It's in there. I was like, well, I can smell it still. I know it's in there somewhere. I'm like, yeah, but I can never give you a carrot.

That's like data to some degree going into AI training. It's there. You don't know where it is anymore, but you can kind of sense it and smell it, but it's not quite there.

So you can't get it back, even if you ask nicely. So this is an attempt to try to get it back, and it's all about tokens. And, you know, in In AI, AI doesn't speak English or a language, a human language.

It speaks digital language. It speaks tokens. So if I say, which animal jumped over the moon? A cow, right? So let's look at the tokens. I'm going to turn on the neuralyzer now, and I'm going to neuralize cow.

And what it did is it neuralized it. It changed the tokens. Let me pause. It changed the weights on those tokens for cow. So normalizing it didn't completely zap it, but it said anytime there's a token related to cow, I'm going to deprioritize that token so that when you're predicting something that ultimately is predicting the output of cow, like in that nursery rhyme, what animal jumped over the moon, Cal jumped over the moon, according to the nursing room.

So the model wants to say Cal. However, Cal has been, that token has been biased, is the word. It's been reduced in its effectiveness. So it's less likely to predict Cal.

So what is it predicting instead? The next best thing, which apparently is Cap. All right. So, but you get, hopefully you get the idea a little bit.

So, let's try it again, list ingredients in cookies. All right, so if I do that, common ingredients, butter, eggs, et cetera. But what happens if I wanted to look for a vegan recipe?

And I don't want mentioning of anything to do with non-vegan things like eggs, right? So let me now neuralize egg. And you could see all of the patterns of those tokens.

It's not just egg, but it's space egg, capital E, egg, taking all those tokens, And now it gives me, as you can see, well, it still has butter, but it doesn't have eggs.

Who's Jerry? Could I have the model forget me? Forget about me. It knows about me. It says I'm an IBM executive known for work with innovation and cloud computing and blockchain.

But now let's have it neuralize Jerry. Jerry Cuomo, Jerry Cuomo. And there's a bunch of variations of me that it has tokens on. And now when you send it to the model, it says, I don't know who Jerry is.

Look at this. I've been forgotten. I've been neuralized. So again, dealing with digital tokens, dealing with cryptography, dealing with bias and weights and understanding that we can make AI more trustworthy.

Hallucinations. That's, if not the last topic, one of the last topics. Language models can produce hallucinations. They're They sure look confident, but incorrect due to lack of user-specific context or directives leading to unreliable output.

So reduce hallucinations by providing specific context using retrieval augmented generation, adding verification steps and structured prompts.

All right, so one of the ways, and you can see this built into some of the more latest chatbots, things like chain of thought.

So prompting guidelines that allow models to, instead of jump to the answer, first reason about the steps to getting to the answer and, you know, magic phrases, things like let's think step by step will really go a long way to help a model, say, pause a little bit, jumping right to the prediction, doing some intermediary steps to have it break down how it's going to get to a better prediction, like thinking about it in phases.

So here's a tool that we built in IBM a few years ago. I hear it's still going strong. And I'm going to pick a very early version of Watson X. And this version, people would say, wasn't very effective as a model.

In fact, you give it like a eighth grade problem, like, you know, I have five tennis balls. Roger buys two more tennis balls. Each can has three tennis balls, how many tennis balls, et cetera.

So it gives the wrong answer, right? So everyone say, well, the model is not very good. Well, let's give it some history here. So let's say what the question is, let's give it a breakdown of the answer like we did.

So instead of saying, before we said the answer was seven. Now we show the math on how we got the answer to be whatever the answer was.

When the model sees that, it emulates it. It says, I'm gonna break down the math. I'm not just gonna guess at an answer. And then lo and behold, it got the answer right.

So the model wasn't that bad. It just wasn't prompted correctly, right? And chain of thought or showing the math through examples is a great way to set up context that will help the model not hallucinate.

Copyright infringement, another area that we've been focusing on. It's AI models that generate code, may inadvertently reproduce copyrighted content.

So how can we do better? And the answer is retrieval augmented generation. We've built this prototype called SimCode It uses the stack with six terabytes of data source code in particular across 358 programming languages, structured variations, et cetera.

And what I really want to be able to do is improve code attribution. You know, where did code like this come from? Look at transparency and guidance in when you're generating code.

So I'll move quickly through this, but I'm going to upload a piece of Python code. It works with Java and some other languages. We use the stack Python, which is the data set.

We use a type of retrieval augmented generation for that. We create, we have a sample, cryptosample.python. We want to see if there's any infringement on this.

It tokenizes the code and it allows for things like if, Variables are in different spots. It'll tokenize it in a way where placement doesn't matter.

And what we'll come up with is, yep, the code has implications around MIT license and BSD license, meaning some of this code is attributed.

It has licenses. But not only can we look at the licensing, now that we kind of have a digital fingerprint of the code, We can look at code review.

We can say, you know, the comment level of the code is low, the complexity of the code, the maintainability of the code.

So once we have it tokenized, we can also compare it to known good references and, you know, kind of make further assertions.

Again, making the code use both trustworthy, because now I kind of know what the licenses are for this code.

I know how it stands against prior art. but also I know the quality of it, right? So think about how this can be applied to gaining trust in other things, medical records and, you know, kind of transaction logs and things of this nature, right?

So in a nutshell back, we talked about trustworthy computing and how the foundation of trustworthy computing is built on good old reputational trusts, trust in people, companies, governments, et cetera.

But we also showed Stemming back from blockchain algorithms and some of the cryptography and proof and fingerprinting and techniques like that, we can kind of see how problems like food safety and digital identity and counterfeiting could be mitigated, but also how algorithmic trust could also be brought to AI, helping with privacy and hallucination and things like fruit you know, forgetting and also deep fakes.

All right. I hope you enjoyed this. And again, Jerry Cuomo, always thinking about trustworthy computing. Got two books out here, Think Blockchain and Think AI, both featured in my classes at North Carolina State University and the topics of my Wild Ducks podcast.

Take care, enjoy, and talk to you soon. Bye.