SCHULTZ FINANCIAL GROUP®

A Legacy of Trust & Innovation Since 1982

‘The ultimate fraud machine’: Scammers are using AI to target people and businesses with increasingly convincing deepfakes

THE GLOBE AND MAIL

Alexandra Posadzki & Joe Castaldo
Feb. 21, 2025

Dan Kagan was in a meeting at his company’s downtown Toronto office last spring when he received a panicked phone call from his 80-year-old mother.

iStock-2187060130

She told him that she’d just gotten off the phone with a police officer, who had pulled over her grandson – Mr. Kagan’s son – for speeding and found a large stash of cannabis in his car. The only way to prevent her grandson, Jordan, from getting arrested was to fork over $2,500 in cash – right away.

Initially skeptical, Mr. Kagan’s mother had asked the caller if she could speak with her grandson. The caller obliged. The voice she heard on the other end of the line sounded exactly like Jordan, and even called her “Bubbie,” like he always did.

There was just one problem with the story: Jordan Kagan, who worked at the same company as his dad, was sitting at his cubicle, where he’d been all day.

The call, it turned out, was a voice deepfake, a recreation of Jordan Kagan’s voice made with artificial intelligence by someone attempting to con an elderly woman out of $2,500.

The scammers were no match for the Kagans, who, as luck would have it, both work at an identity management company called Okta Inc. – Dan as vice-president and country manager for Canada, his son Jordan as a business development representative.

“My mom swore up and down it was my son,” said Dan Kagan. “They had my son’s voice perfectly, and the inflection and just the way he spoke.”

AI-generated deepfakes, such as the one that targeted the Kagans, are an increasingly popular tool for cybercriminals, according to experts. The deepfakes offer a low-cost, high-tech way to convince victims to part with their cash by posing as someone they know and trust.

Schultz Financial Group, Inc.
Wealth Advisors
Office : (775) 850-5620
info@sfginc.com
sfginc.com

“There’s a massive upswing in the use of deepfakes for fraud campaigns,” says Sam Rubin, the global head of operations at Unit 42, the threat intelligence and cybersecurity consulting arm of California-based Palo Alto Networks Inc.

But consumers, who can be targeted with deepfake videos of celebrities hawking fraudulent investment schemes or by scammers sampling the voices of loved ones, are not the only ones at risk. Generative AI is also being used to peddle deepfake media aimed at defrauding businesses, creating a new financial risk for companies to manage.

Last November, the U.S. Financial Crimes Enforcement Network cautioned that since 2023, it has seen an increase in the use of deepfake media in fraud schemes that target financial institutions. For instance, criminals are using generative AI to create or modify images used on identification documents such as passports and drivers’ licences, and then using those fraudulent identities to open accounts and launder the proceeds of other fraud schemes, FinCEN said.

Entrust Corp., a Minneapolis-based company that processes millions of identity verifications each year, spotted a similar trend. The number of deepfakes increased by 3,000 per cent from 2022 to 2023, Entrust noted in its 2025 identity fraud report.

Meanwhile, more than a quarter of C-suite and other executives polled by consulting giant Deloitte in May, 2024, said their organization had experienced at least one deepfake incident targeting their financial data – with more than half of 2,131 respondents expecting the prevalence of such attacks to increase over the following year.

“We’re only at the tip of the iceberg,” said Robert Blanchard, principal data scientist at data analytics firm SAS Institute Inc.

The good news, said Mr. Blanchard, is that while technology has created the problem, it may also hold the solution. The bad news, he notes, is that the situation is likely to get worse before it gets better.

Deepfake scams are growing in popularity because of how inexpensive they are to create, and how substantial the financial rewards can be. The process of cloning an individual’s voice and appearance has become easier, cheaper and faster. And the clones have become more convincing.

“It’s taking off dramatically in the scam ecosystem,” said Ashley Jess, senior intelligence analyst based in Hamilton with U.S. cybersecurity company Intel 471. “Deepfakes are more readily accessible, you need fewer images to create a convincing one, and you can just buy it as a service.”

Scammers are using deepfakes to create chief executive officer frauds – which involve impersonating a company’s CEO to trick employees into making payments – that are more convincing. Traditionally done via e-mail or text messages, nowadays scams may use fake video or audio of the CEO.

Employees of information technology services provider Kyndryl Holdings Inc. are familiar with this scheme, having been contacted many times over the years by “fake Martin,” an impersonation of their CEO, Martin Schroeter. Some of those attacks utilized deepfakes, though none have succeeded, according to Fortune .

But one high-profile incident in Hong Kong last year demonstrates how effective – and lucrative – such schemes can be. Fraudsters created deepfake recreations of multiple employees of British multinational design and engineering firm Arup Group Ltd., including its chief financial officer, according to CNN. The effort paid off: After attending a video call with the deepfake recreations, one of the company’s Hong Kong staff was duped into sending the scammers more than US$25-million.

In one case that Mr. Rubin worked on, hackers infiltrated a company by calling its IT help desk using what he believes was a deepfake recreation of an employee’s voice.

That sort of thing is still relatively rare, said Mr. Rubin, but experts predict that the use of deepfakes against businesses will become more common.

Fraudsters can choose from an array of options that fit their budget. Cloning someone’s voice is “extraordinarily cheap” says Kyle Wilhoit, a threat researcher at Unit 42. There are free tools available on Microsoft Corp.’s GitHub, as well as inexpensive paid services that typically range from US$10 to US$15 per month, he said.

Voice samples can be obtained from social media posts or even a voicemail greeting.

“It’s something like three to five seconds or less of voice sample that is required in order to be able to replicate someone’s voice convincingly enough that most people can’t discern an AI doing it,” said Jon Ferguson, vice-president of cybersecurity and DNS at the Canadian Internet Registration Authority.

Mr. Wilhoit said his daughter received a series of phone calls at work in the days before scammers unsuccessfully tried to extort her family. They told her that she’d won a radio contest and they needed her to say specific things.

For scammers lacking in technical know-how, they can turn to a thriving and anonymous online marketplace of people offering tips and tutorials, who also hawk their own deepfake creation services. There is, for example, an active Discord server for users of DeepFaceLab, a popular open-source program for building deepfakes, and a spinoff called DeepFaceLive, which enables people to pose as someone else in live video calls. The software was developed by a ragtag group of researchers from around the world in 2020. They even published a paper about their techniques.

The Discord community is actively trying to improve the software, offering guidance to newcomers and sharing completed face models for others to use, often of celebrities. Other models are more generic, with names such as “normal asian woman” and “random asian man.” There are additional datasets for download to help increase the realism of deepfakes and eliminate the technical glitches that can expose the ruse, such as one dataset that has some 6,000 images of human teeth.

People often post in the Discord to ask if anyone can make them a custom deepfake model, and offering to pay for it. A user by the name of Cxsmo, who developed another open-source program called DeepAscension, regularly responds. Contacted by The Globe and Mail, Cxsmo was exceedingly cautious and declined to reveal their name or location while only communicating through encrypted messaging. But Cxsmo’s answers were entirely benign. “I’m all about helping people enjoy the tech,” Cxsmo wrote. “I try to keep things positive, and if someone’s looking to do something sketchy, I steer clear of it.”

The reality of open-source software, however, is the creators cannot control how it’s used. Even proprietary generative AI applications can be coaxed into automating financial scams. In October, researchers from the University of Illinois Urbana-Champaign built a few AI agents (applications that can do things such as navigating to websites and entering login details) on top of a voice-enabled large language model from OpenAI. The researchers instructed the agents to carry out a variety of phone scams, including soliciting banking information and transferring funds. (Doing so required jailbreaking the OpenAI model, circumventing built-in safeguards.)

The researchers posed as the gullible victims on the other end of the phone line with the AI agents. The study found the agents successfully transferred funds through a major U.S. bank 20 per cent of the time, and succeeded 40 per cent of the time on a cryptocurrency platform. When the agent failed, it was due to transcription errors or trouble navigating complex websites.

The study ultimately showed these feats are technically possible. The glaring caveat is that since the researchers themselves played along as victims, it’s not clear how successful an AI bot would be in the real world. “We do not focus on the persuasion aspect of the scam,” they wrote.

Despite the concerns and headlines, some lawyers who respond to cyberbreaches for corporate clients say that deepfakes are not yet playing a major role. “We have seen some of them, but the examples we’ve seen have been caught by clients,” said John Cassell, a partner at Norton Rose Fulbright in Calgary.

But Mr. Kagan believes that deepfake scams are happening more frequently than people realize. “I just think a lot of it is going unreported, because it’s embarrassing,” he said.

One thing seems certain: As the deepfake technology continues to improve, the danger that it poses will grow.

“Prior to about five months ago, if you had a deepfake in front of you on the screen, and the guy moved his face too quickly, it would digitize a little bit on the side of the face,” said Chris Mathers, president of consulting and investigative firm Chris Mathers Inc. “Or if you put your hand in front of your face, sometimes there would be a trail, and you could see the person behind it.”

That’s no longer the case, according to Mr. Mathers, a former undercover RCMP officer who gives talks on the topic of deepfakes.

“The technology is increasing exponentially. It’s going to get better and better. It’s going to improve itself.”

Just as AI has opened up new avenues for scammers, it’s also created opportunities for technology providers to thwart bad actors.

A report published by IBM last summer found that companies that had integrated AI and automation into their cybersecurity operations experienced breaches that, on average, were 54 days shorter and cost US$2.84-million less than those not using the technologies.

“It’s certainly not all doom and gloom,” said Chris Thompson, who’s based in San Diego, Ca., and is the global head of IBM’s team of ethical hackers, X-Force Red. “Just like attackers are trying to leverage AI to move faster, defenders can leverage AI to investigate attacks faster.”

A number of startups have sprung up that have a specific focus on combatting nefarious uses of generative AI, including deepfakes.

Shuman Ghosemajumder, a Canadian who previously served as the global head of product, trust and safety at Google, co-founded a company called Reken in San Francisco last year to defend against security threats posed by generative AI. The limitations of the technology, including the fact that large language models can invent information, might be a problem for businesses, but

that doesn't matter for cybercriminals, he said. "If I want to be able to just generate content, including images, video and audio that sounds believable, now I've got a tool that enables me to do that," he said. "It's actually the ultimate fraud machine."

Merely detecting AI-generated content isn't sufficient, he said. Doing so accurately all of the time is impossible because of the proliferation of open-source AI models that detection software might not have seen before, and therefore not be able to spot. Secondly, AI-generated content is already spreading within companies, as Microsoft and Google give people the ability to write e-mails and produce images with AI. "Just because something is AI-generated doesn't necessarily mean that it's malicious," Mr. Ghosemajumder said. "It's ultimately the maliciousness that you're trying to determine."

To that end, Reken is building an AI-powered platform to launch later this year that, in part, will be able to sniff out malicious content. "The only way you can really do this at scale is by deploying large-scale machine learning systems and analyzing as much data as you can," he said. But he declines to elaborate on how exactly Reken will determine maliciousness.

DeepTrust, another San Francisco-based startup, is already testing its security platform with a small group of companies. Noah Kjos co-founded the startup in 2023 partly due to an experience his grandfather had in which he received a phone call from someone who had cloned Mr. Kjos's voice. (His grandfather didn't fall for it.) Mr. Kjos and his co-founder originally pursued a pure deepfake detection company, but felt there was a bigger opportunity to apply the technology to live interactions and to guard against attempts at social engineering.

The company has built software that integrates with Zoom and other video calling platforms to detect AI-generated audio and also analyze the content of the conversation against a company's internal policies to determine if something is amiss.

"The AI models we've trained have extensive knowledge of social engineering tactics and can frame the current conversation against that," Mr. Kjos said. "If we're going through a password reset process, and you're dodging certain things or trying to circumvent anything, that's what it will be able to flag." The software can provide the employee with additional questions to ask or details to confirm, while also sending a notification to the company's security team.

When Mr. Kagan got the call from his mother, he knew right away that he was dealing with a scam.

The fraudsters had told Mrs. Kagan they'd send someone to her home to pick up the cash. Unfortunately, by the time the real police had been enlisted to help, the impersonators had seemingly realized that their cover had been blown. They never showed up at Mrs. Kagan's condo, where York Regional Police officers were waiting to arrest them.

Mr. Kagan still muses about how things may have unfolded.

"Had I not answered the phone I think things would have gone a little bit differently," he said.