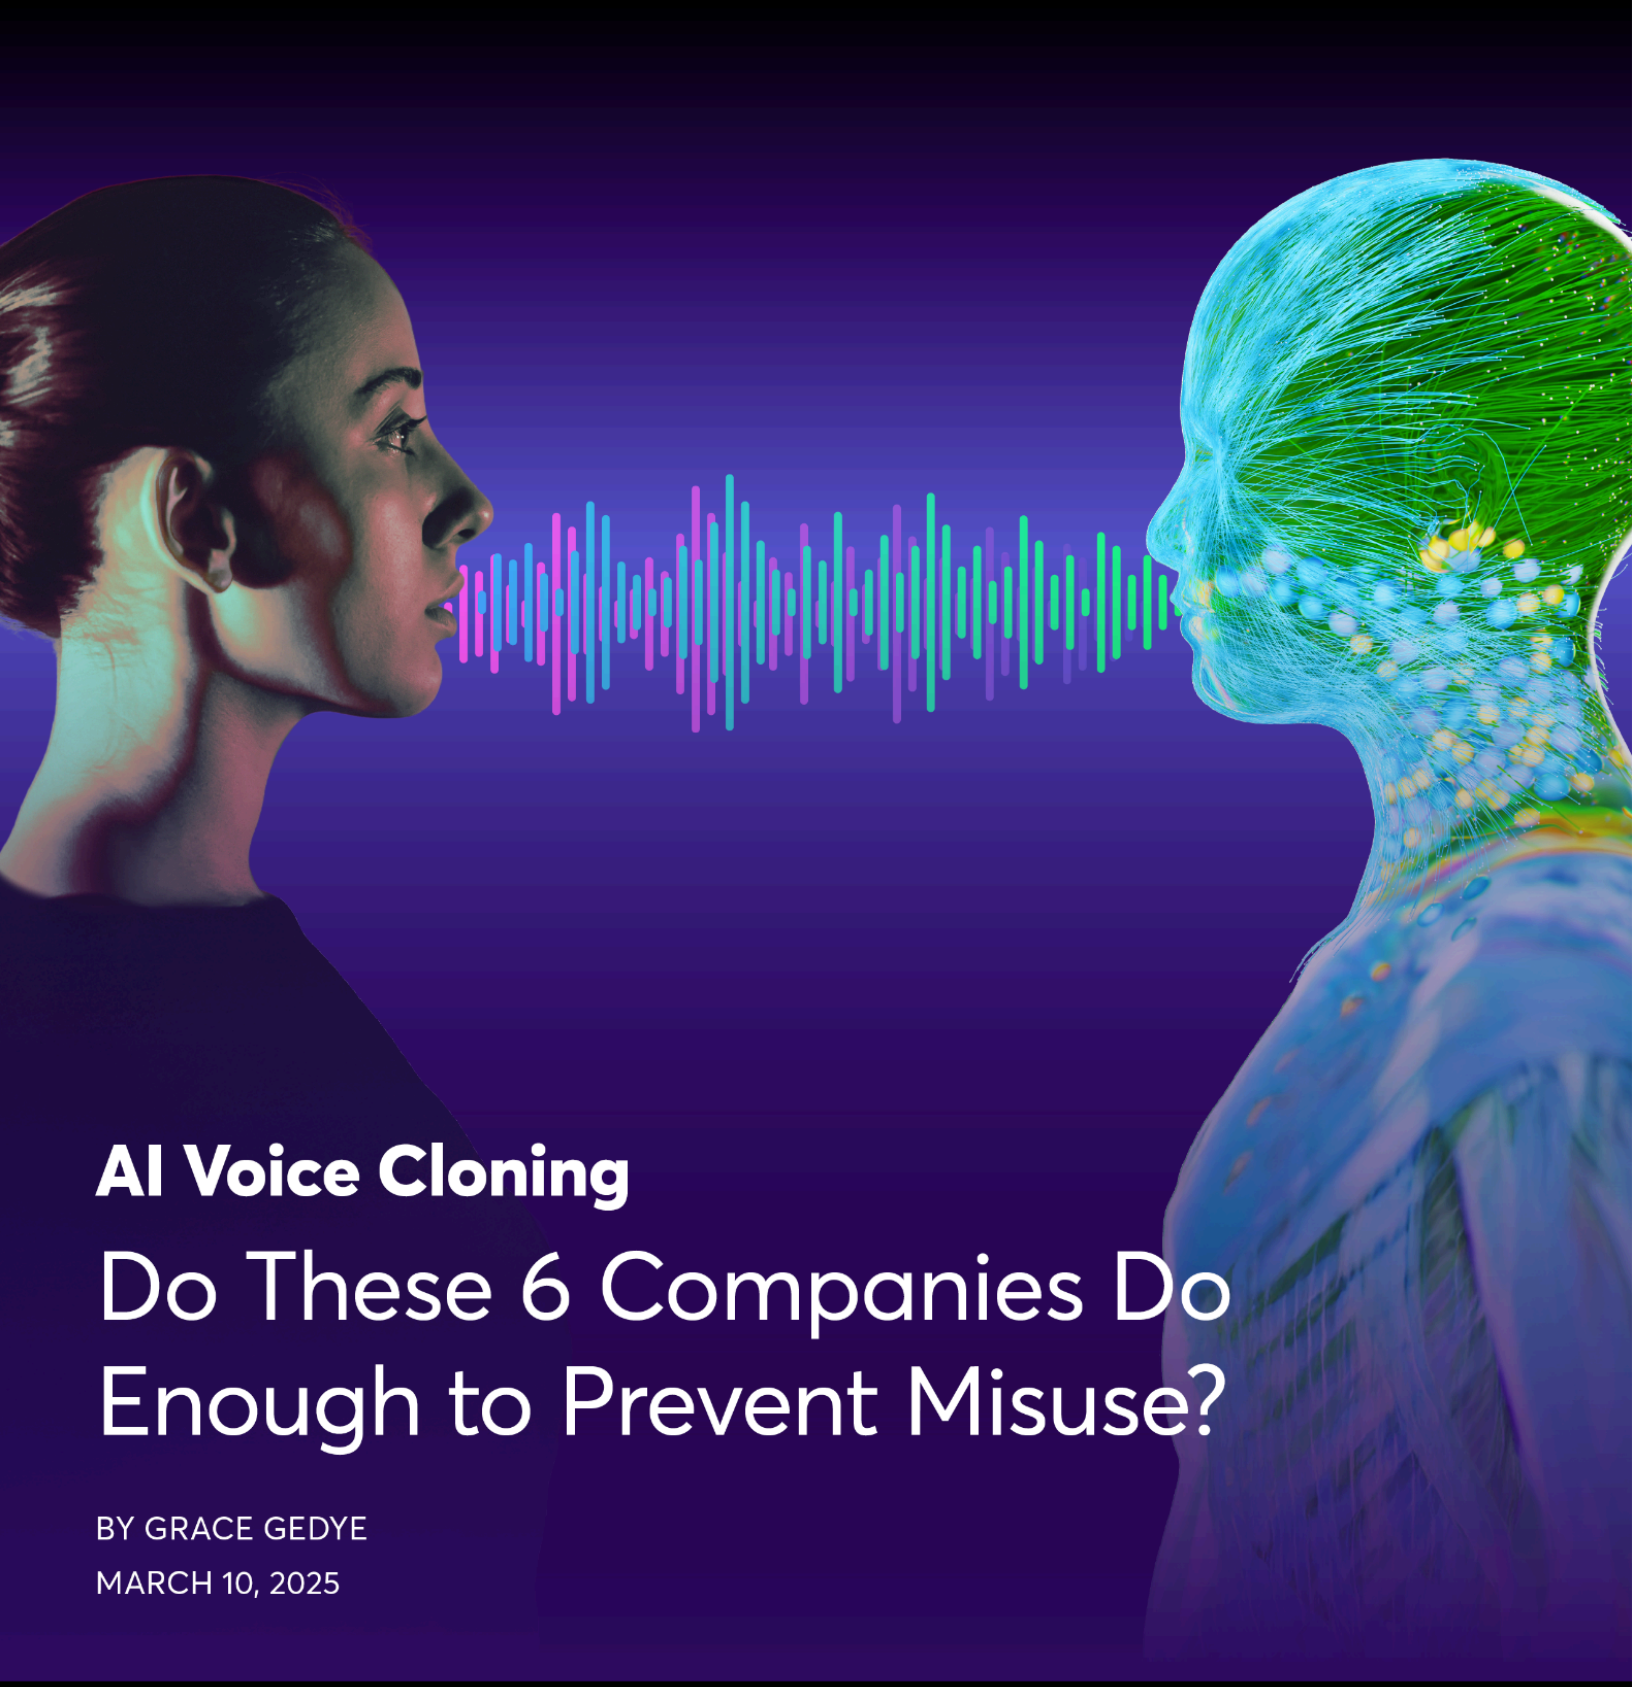

AI Voice Cloning

Do These 6 Companies Do Enough to Prevent Misuse?

BY GRACE GEDYE

MARCH 10, 2025

Consumer Reports

Table of Contents

Introduction

Consumer Reports (CR) is an independent, nonprofit, and nonpartisan testing organization. We use product ratings, journalism, research, and advocacy to create a fairer, safer, and healthier world. We've long worked to protect consumers from fraud, and artificial intelligence tools have unlocked new avenues for scammers.

This report is concerned with the rise of AI voice cloning products that enable consumers to clone—that is, create an artificial copy of—an individual's voice using only a short audio clip of the individual speaking. AI voice cloning products have many legitimate uses, including speeding up audio editing, enhancing movie dubbing, and automating narration. But without proper safeguards, these products also present a clear opportunity for scammers, who have used the technology to impersonate, for example, a consumer's grandchild calling in need of money, and celebrities and political figures endorsing dubious products and bogus investment schemes. These tools are easily found using Google and other search engines; it costs little, or sometimes nothing, to use them.

Consumer Reports assessed voice cloning products from six companies and found that four of those companies erected no meaningful barriers to cloning someone's voice without their consent. The other two provided a mechanism aimed at confirming consent. We make recommendations for companies that offer AI voice cloning, and we argue that the nascent voice cloning industry should adopt norms and standards to mitigate the risk of fraud. We also argue that the Federal Trade Commission already has the authority to regulate these technologies under Section 5 of the FTC Act and that it should take action against companies that offer voice cloning tools without reasonable mechanisms in place to prevent harmful uses.

Background

Legitimate uses of AI voice cloning tools

This report focuses on AI tools that generate speech. Some AI voice tools provide preloaded voices that consumers can use as is or, in some cases, modify. AI voice *cloning* tools are a subset of these products that enable users to create custom voices that sound like a specific person. Users can clone their own voice or, in some cases, other people's voices. AI voice cloning products require users to record or upload recordings of the voice they wish to clone. The AI tools process that audio to emulate the speaker's voice. They can generate novel and sometimes uncannily realistic audio, including speech that is disparaging or fraudulent.[1]

AI voice companies market their products for use in a range of scenarios, including:

- Automating narration of audiobooks, articles, and blog posts
- Creating character voices for video games
- Automating voiceovers for social videos
- Editing audio clips and inserting new content without re-recording
- Translating language (including audio-to-audio translations) and dubbing movies
- Customizing speaker content or accent in marketing campaigns to target different regions or customers
- Providing voice clones to people who are going to lose or have lost the ability to speak
- Generating educational or corporate training content
- Generating voices for answering customer service inquiries

Some voice cloning tools are aimed at specific populations, such as The Voice Keeper,[2] which is focused on serving children with speech or language impairments. Other companies, such as ElevenLabs, market themselves for a wide range of use cases.[3]

---

[1] Kate Knibbs, "Researchers Say the Deepfake Biden Robocall Was Likely Made With Tools From AI Startup ElevenLabs," *Wired,* January 25, 2024, https://www.wired.com/story/biden-robocall-deepfake-elevenlabs;
Ben Finley, "Deepfake of principal's voice is the latest case of AI being used for harm," *AP,* April 30, 2024, https://apnews.com/article/ai-maryland-principal-voice-recording-663d5bc0714a3af221392cc6f1af985e;
Stuart A. Thompson, "How 'Deepfake Elon Musk' Became the Internet's Biggest Scammer," *The New York Times,* August 14, 2024, https://www.nytimes.com/interactive/2024/08/14/technology/elon-musk-ai-deepfake-scam.html;
Drew Harwell, "An artificial-intelligence first: Voice-mimicking software reportedly used in a major theft," *The Washington Post,* September 4, 2019, https://www.washingtonpost.com/technology/2019/09/04/an-artificial-intelligence-first-voice-mimicking-software-reportedly-used-major-theft.

[2] *The VoiceKeeper,* https://www.thevoicekeeper.com/kids-about-us.

[3] "Explore uses of our generative AI audio platform," *ElevenLabs,* https://elevenlabs.io/use-cases.

Problematic uses of AI voice cloning tools

With the ability to simulate a specific person's voice comes acute risks for impersonation and fraud, which we break into three general categories below.

A) Bad actors use AI voice tools to defraud and mislead consumers by impersonating everyday Americans.

AI voice cloning tools have the potential to supercharge impersonation scams, including a phone scam sometimes known as the "Grandparent scam," in which a consumer is told that a loved one is in trouble—for example, they crashed their car or landed in jail—and need money fast.[4] In the past, scammers might try to achieve a rough approximation of a relative's voice to make the ruse compelling. Now, if scammers can obtain audio of the loved one speaking from, for example, a social media video, they can create a potentially compelling AI clone of the voice.

This is already happening. In March 2023, The Washington Post wrote about parents who lost more than $15,000 after purportedly receiving a call from a "lawyer" who told them their son had killed a U.S. diplomat in a car accident, and hearing what they thought was their son's voice on the phone.[5] In March of 2024, The New Yorker highlighted the stories of several parents sent into a state of terror after thinking they heard their panicked child's voice, followed by dark threats.[6] Scammers have targeted companies as well, using AI voice tools to convince employees they are getting a call from an executive who needs them to transfer funds. In one case, the managing director of a British energy company wired $240,000 to Hungary, thinking he was speaking to his boss.[7]

In February 2024, CR asked consumers across the country whether they had ever received a phone call from a scammer mimicking the voice of someone they knew or a well-known figure. We heard from hundreds of consumers who had this experience and said it left them feeling "vulnerable," "shaken by the experience," and "really weirded out." Below are excerpts of their testimonials:

---

[4] Alvaro Puig, "Scammers use AI to enhance their family emergency schemes," *Federal Trade Commission,* March 20, 2023, https://consumer.ftc.gov/consumer-alerts/2023/03/scammers-use-ai-enhance-their-family-emergency-schemes.

[5] Pranshu Verma, "They thought loved ones were calling for help. It was an AI scam," *The Washington Post,* March 5, 2023, https://www.washingtonpost.com/technology/2023/03/05/ai-voice-scam.

[6] Charles Bethea, "The Terrifying A.I. Scam That Uses Your Loved One's Voice," *The New Yorker,* March 7, 2024, https://www.newyorker.com/science/annals-of-artificial-intelligence/the-terrifying-ai-scam-that-uses-your-loved-ones-voice.

[7] Drew Harwell, "An artificial-intelligence first: Voice-mimicking software reportedly used in a major theft," *The Washington Post,* September 4, 2019, https://www.washingtonpost.com/technology/2019/09/04/an-artificial-intelligence-first-voice-mimicking-software-reportedly-used-major-theft.

- “My grandpa got a call from someone claiming to be me. Supposedly, I was traveling, and my car broke down and I needed to have him send money, so I could complete my travels. Grandpa said there was no doubt in his mind that I was the caller and was preparing to do as asked. Luckily, before he went through with the transaction, he reasoned that if I was in trouble and honestly needed money, he would have heard from my mom. … Scary that the tools they use could imitate my voice that closely as to fool a close relative.” —CR member, Minnesota

- “The initial caller’s voice sounded very much like my nephew’s. He knew family details, pleaded with me not to call his father, and promised to pay me back as soon as he got home—all very convincing. I should add that I spent more than 60 years in law enforcement and intelligence work. This scam was so carefully arranged and executed that I fell for it nevertheless.” —CR member, Massachusetts

- “I received a phone call from my grandson explaining that he was in a car accident at college and needed $5,000. He sounded scared and upset and asked that I not tell his parents. So I went to my bank to get the money, and the bank teller told me it was a scam. I did not believe her, as I was sure it was my grandson's voice.” —CR member, New York

- “The voice on the other end sounded just like my grandson, and it said, ‘Gramie, I've been in an accident.’” —CR member, Florida

- “... I was skeptical, and told him I had heard of scams such as this. So he said, ‘I'll let Nate say a few words to you.’ It sounded exactly like my Nate! He has a rather unusual voice, so I was then almost convinced.” —CR member, Indiana

- “The phone rang and a voice said, ‘Hi Gramma, this is Mac. I’m in New Jersey with my friend Chris. We had an accident. I broke my nose.’ I immediately knew it wasn’t my grandson. He calls me Gramma Beth … and he’d have no reason to be in New Jersey. He’s New York, born and bred. … The voice did sound exactly like him, however, and I could easily have been duped.” —CR member, New York

- “I received a call and heard my daughter crying hysterically! She wasn't making sense, so an ‘officer’ took over the call. He stated I needed to come right away but would not answer my questions. Thankfully, I have Life360 and looked to see where my daughter was at, and it showed her at home. … To hear my daughter’s crying voice shook me for a long time!” —CR member, Minnesota

We don’t know whether or not the scam attempts above relied on AI voice cloning. Either way, they are examples of the types of scams to which AI voice cloning naturally lends itself.

Imposter scams are common. In 2023, 853,935 imposter scams were reported to the Federal Trade Commission's Consumer Sentinel Network.[8] Twenty-one percent of those scam reports included monetary losses, which totaled $2.7 billion in 2023. Imposter scams were the second-most-frequent category of fraud reported to the Consumer Sentinel Network in 2023 out of the 29 categories that the network tracks.

Bad actors may also use AI voice cloning tools to bypass voice-based security measures, including voice recognition systems used by some banks. For example, a Motherboard journalist successfully used ElevenLabs' AI voice cloning tool to bypass a voice ID security measure and break into a bank account in 2023.[9] At the time, several large U.S. banks used voice ID measures, including TD Bank, Chase, and Wells Fargo. In 2024, a BBC journalist used voice cloning tools to bypass voice ID mechanisms for her accounts with the banks Santander and Halifax.[10] The extent to which scammers are using this technique is unknown, but if reporters have thought to test their bank security in this manner, fraudsters likely have too.

AI voice cloning tools can also be used for deception with different ends—including harming someone's reputation. In 2024, a Maryland high school athletic director reportedly used AI voice cloning tools to mimic the voice of a school principal.[11] The recording came after the athletic director and the principal had discussed the athletic director's poor work performance. The audio clip, which sounded like the principal, reportedly contained racist remarks about Black students' test-taking abilities, as well as antisemitic remarks.

Some AI voice companies explicitly market the fact that their products can be used for deception. PlayHT, a voice cloning company, lists "pranks" as a use case for its AI voice tools in a company blog post.[12] Speechify, another AI voice company, also suggests prank phone calls as a use case for its tools. "There's no better way to prank your friends than by pretending you're someone else. Voice changer apps allow you to make real-time changes to your voice and trick your friends into thinking you're a different person when you make a fake call online," reads one page of Speechify's company website.[13]

---

[8] Federal Trade Commission, "Consumer Sentinel Network Data Book 2023," February 2024. Available at: https://www.ftc.gov/system/files/ftc_gov/pdf/CSN-Annual-Data-Book-2023.pdf.
[9] Joseph Cox, "How I Broke Into a Bank Account With an AI-Generated Voice," *Vice,* February 23, 2023, https://www.vice.com/en/article/how-i-broke-into-a-bank-account-with-an-ai-generated-voice.
[10] Shari Vahl, "Cloned customer voice beats bank security checks," *BBC,* November 28, 2024, https://www.bbc.com/news/articles/c1lg3ded6j9o.
[11] Ben Finley, "Athletic director used AI to frame principal with racist remarks in a fake audio clip, police say," *AP,* April 25, 2024, https://apnews.com/article/ai-artificial-intelligence-principal-audio-maryland-baltimore-county-pikesville-853ed171369bcbb888eb54f55195cb9c, and Ben Finley, "Deepfake of principal's voice is the latest case of AI being used for harm," *AP,* April 30, 2024, https://apnews.com/article/ai-maryland-principal-voice-recording-663d5bc0714a3af221392cc6f1af985e.
[12] "Boy Voice Changer: Everything You Need to Know," *PlayHT,* September 20, 2024, https://play.ht/blog/boy-voice-changer.
[13] "Girl voice changer online for a call," *Speechify,* https://speechify.com/blog/girl-voice-changer-online-for-a-call.

B) Bad actors use AI voice cloning tools to defraud consumers by impersonating public figures.

AI voice cloning tools have enabled scammers to generate deepfake videos falsely depicting celebrities and political figures endorsing products, recommending investments, and otherwise trying to influence ordinary people. Some AI voice tools, such as Parrot AI, offer a wide range of preloaded voice clones for public figures and celebrities, including President Donald Trump, former President Joe Biden, Joe Rogan, Kim Kardashian, Will Smith, Rep. Alexandria Ocasio-Cortez (D-N.Y.), and more.[14]

AI-powered celeb-bait has proliferated on social media. An investigation by ProPublica identified videos on Meta platforms seemingly depicting Trump and Biden—each with their distinctive tone and cadence—offering cash handouts to people who filled out an online form.[15] 404 Media has reported on the proliferation of AI clones of Rogan, Taylor Swift, Ice Cube, Andrew Tate, Oprah Winfrey, and The Rock pushing Medicare- and Medicaid-related scams on YouTube.[16] Scammers using AI clones of Taylor Swift's likeness and voice have also hawked Le Creuset dishware.[17] Elon Musk's likeness and voice have been frequently repurposed by scammers using AI video and voice tools to push fraudulent investment schemes. One consumer was reportedly scammed out of $690,000 after watching a deepfake Musk endorse an investment opportunity.[18] Recent research suggests that consumers often fail to recognize deepfake videos for what they are and also overestimate their own ability to detect deepfakes.[19]

C) Additional concerns.

Voice cloning tools raise many other concerns beyond the fraud, impersonation, and scams that are the focus of this report. For example, there are ethical and legal implications of some training data collection practices. Voice actors are suing AI voice company Lovo, alleging that

---

[14] "Make a celebrity say anything," *Parrot AI,* https://www.tryparrotai.com.
[15] Craig Silverman and Priyanjana Bengani, "Exploiting Meta's Weaknesses, Deceptive Political Ads Thrived on Facebook and Instagram in Run-Up to Election," *ProPublica* and *Tow Center for Digital Journalism,* October 31, 2024, https://www.propublica.org/article/facebook-instagram-meta-deceptive-political-ads-election.
[16] Jason Koebler, "Deepfaked Celebrity Ads Promoting Medicare Scams Run Rampant on YouTube," *404,* January 9, 2024, https://www.404media.co/joe-rogan-taylor-swift-andrew-tate-ai-deepfake-youtube-medicare-ads.
[17] Tiffany Hsu and Yiwen Lu, "No, That's Not Taylor Swift Peddling Le Creuset Cookware," *The New York Times,* January 9, 2024, https://www.nytimes.com/2024/01/09/technology/taylor-swift-le-creuset-ai-deepfake.html.
[18] Stuart A. Thompson, "How 'Deepfake Elon Musk' Became the Internet's Biggest Scammer," *The New York Times,* August 14, 2024, https://www.nytimes.com/interactive/2024/08/14/technology/elon-musk-ai-deepfake-scam.html.
[19] Nils C. Köbis, Barbora Doležalová, Ivan Soraperra, "Fooled twice: People cannot detect deepfakes but think they can," *iScience* 24, no. 11 (2021): https://www.sciencedirect.com/science/article/pii/S2589004221013353.

the company cloned their voices without permission or proper compensation.[20] One plaintiff claims he was contacted for voiceover work on the freelancer platform Fiverr by someone he later discovered to be a Lovo employee. He was reportedly told the recording would be "used for academic research purposes only." Years later, he heard his voice in news videos and podcasts made by companies with whom he did not have a contract.

Another concern is what happens with a customer's voice data when they upload audio to clone their own voice. Some companies give themselves the right to use that audio to train their underlying model; others provide an opt-out. Companies can also update their privacy policy at any time to give themselves the right to sell or share customer voice data with government entities or companies developing surveillance products and then subsequently begin doing so.

Non-consensual voice clones can also be used to edit audio of influential figures to spread misinformation or disinformation. Ahead of New Hampshire's primary election, a political consultant and a magician used ElevenLabs to create an AI clone of Joe Biden's voice discouraging citizens from voting in the primary and then sent the message out as a robocall to New Hampshire voters.[21] In February of 2024, CNN reported that a fake recording created with AI of a top candidate in a Slovakian election went viral; the recordings sounded like the candidate was bragging about rigging the election and talking about raising the price of beer.[22]

Some companies are aware of the potential for harm

Voice cloning tools clearly present real risks, both to the public and—in the form of liability and bad publicity—to the companies that offer them. In fact, some larger tech companies seem to have responded to those risks by deciding to delay the release of already-developed voice cloning capabilities to the public.

For example, Microsoft developed a text-to-speech model called VALL-E but says it currently has no plans to make it available to the public or to integrate it into a product, writing:

---

[20] Winston Cho, "Actors Hit AI Startup With Class Action Lawsuit Over Voice Theft," *The Hollywood Reporter,* May 16, 2024, https://www.hollywoodreporter.com/business/business-news/actors-hit-ai-startup-with-class-action-lawsuit-over-voice-theft-1235900689.

[21] Holly Ramer and Ali Swenson, "Political consultant behind fake Biden robocalls faces $6 million fine and criminal charges," *AP,* May 23, 2024, https://apnews.com/article/biden-robocalls-ai-new-hampshire-charges-fines-9e9cc63a71eb9c78b9bb0d1ec2aa6e9c; Maggie Astor, "Behind the A.I. Robocall That Impersonated Biden: A Democratic Consultant and a Magician," *The New York Times,* February 27, 2024, https://www.nytimes.com/2024/02/27/us/politics/ai-robocall-biden-new-hampshire.html; Vijay Balasubramaniyan, "Pindrop Reveals TTS Engine Behind Biden AI Robocall," *Pindrop,* January 25, 2024, https://www.pindrop.com/article/pindrop-reveals-tts-engine-behind-biden-ai-robocall.

[22] Curt Devine, Donie O'Sullivan, and Sean Lyngaas, "A fake recording of a candidate saying he'd rigged the election went viral. Experts say it's only the beginning," *CNN,* February 1, 2024, https://www.cnn.com/2024/02/01/politics/election-deepfake-threats-invs/index.html.

> It may carry potential risks in the misuse of the model, such as spoofing voice identification or impersonating a specific speaker. … If the model is generalized to unseen speakers in the real world, it should include a protocol to ensure that the speaker approves the use of their voice and a synthesized speech detection model.[23]

Another example is OpenAI, which developed a voice cloning product called Voice Engine but has thus far held it back from wide release, noting that "generating speech that resembles people's voices has serious risks."[24] In a statement about Voice Engine, the company writes:

> We believe that any broad deployment of synthetic voice technology should be accompanied by voice authentication experiences that verify that the original speaker is knowingly adding their voice to the service and a no-go voice list that detects and prevents the creation of voices that are too similar to prominent figures.

Some companies provide more bespoke voice cloning with some human oversight. WellSaid Labs and Murf AI, for example, indicate on their respective websites that they offer AI voice cloning but that customers must contact a sales manager in order to clone a voice.

Murf AI requires customers to provide 1 to 2 hours of professional studio-level recordings of the voice a client wishes to clone, to record a script provided by Murf, and to then wait one to four weeks for it to develop an AI avatar, according to its website.[25] The company also claims to limit access to the resulting clone exclusively to the client and their team and protect that access with two-factor authentication.[26] WellSaid Labs claims on its website that it never allows users to independently clone voices, which "ensures users cannot create a voice clone of a non-consenting individual by uploading their audio data."[27] WellSaid also claims to have content moderation mechanisms in place that can scan for variations on the kinds of content it prohibits and then bar the content from production until it's able to conduct a review.[28] CR contacted both companies for research interviews to learn more about how they safeguard their products but did not get a response.

---

[23] "VALL-E 2," *Microsoft,* https://www.microsoft.com/en-us/research/project/vall-e-x/vall-e-2; Benj Edwards, "Microsoft's new AI can simulate anyone's voice with 3 seconds of audio," *ArsTechnica,* January 9, 2023, https://arstechnica.com/information-technology/2023/01/microsofts-new-ai-can-simulate-anyones-voice-with-3-seconds-of-audio.

[24] "Navigating the Challenges and Opportunities of Synthetic Voices," *OpenAI,* March 29, 2024, https://openai.com/index/navigating-the-challenges-and-opportunities-of-synthetic-voices.

[25] "Voice Cloning Process," *Murf.AI,* https://help.murf.ai/are-there-any-specifications-for-recording-the-voice-data.

[26] "AI Voice Cloning: Craft Custom Voice Clones for Unique Experiences," *Murf.AI,* https://murf.ai/voice-cloning.

[27] "We believe in Responsible AI," *WellSaid,* https://www.wellsaid.io/resources/ethics.

[28] Martin Ramirez, "AI for Good: Content Moderation at WellSaid Labs," *WellSaid,* February 14, 2022, https://www.wellsaid.io/resources/blog/content-moderation-ai-for-good?utm_source=rss&utm_medium=rss&utm_campaign=content-moderation-ai-for-good.

Other companies, however, have in some cases released voice cloning products to the public without technical mechanisms to ensure the original speaker has consented.

Methodology

Consumer Reports set out to evaluate how a handful of companies offer voice cloning to consumers. We chose six companies that offer voice cloning tools and that represent a range of practices when it comes to safeguarding against the misuse of their products.

For each company selected, we attempted to create a voice clone using publicly available audio of a CR employee. We conducted this part of the study between September 2024 and January 2025. We also reviewed the companies' privacy policies.

Based on the information we gathered, we attempted to answer the following questions designed to address our core concerns surrounding impersonation, fraud, and data privacy.

1. **What customer information is required before you can create a custom voice clone?**
2. **How much does it cost to create a custom voice clone?** We sought to determine the least a user could pay to access custom voice cloning. If the company offered a free trial that enables customers to access custom voice cloning for a period, we counted that as "free." If the initial price was low and then later increased, we recorded the initial price.
3. **Are there technological barriers to frustrate non-consensual cloning?** We were looking for mechanisms to prevent the creation of non-consensual voice clones beyond checkbox-style self-attestation that customers would abide by the terms of service and/or clone only voices to which they had the legal right.
4. **Does the privacy policy give the company permission to use customer voices to train or improve its model?** We sought to determine whether a company's privacy policy explicitly or implicitly permits the company to use customer voices to train or improve its model(s). We also looked to determine whether the policy explicitly mentions an opt-out or opt-in.
5. **Does the privacy policy permit the company to allow other customers or companies to use your voice model or data?** We also noted whether the privacy policy explicitly or implicitly permits the company to sell or share customers' voice models or data with other companies. And we noted whether the company permitted customers to opt in to sharing their voice data or allowed them to consent.
6. **Does the privacy policy grant users the right to delete all their voice data and other personal data?** We also noted whether a company offered a right to delete but nevertheless maintained its right to retain data for "legitimate business interests." If the company offered a right to delete but nevertheless maintained its right to retain some data for specific and more narrowly defined purposes—such as fraud detection or compliance with laws—we recorded that as a right to delete with minor exceptions.

If a company offered multiple voice cloning products or multiple tiers of the same product, we answered the questions based on the least expensive product or the product that required the least information from its users.

In addition, in November 2024, we contacted the six companies in the test set to ask additional questions, including whether they mark or modify the audio their products generate so that it can later be detected as AI-generated content, and whether they attempt to detect and prevent the creation of any kinds of harmful content. Because privacy policies can be vague, we also asked detailed questions about their data practices. A full list of these questions and the companies' responses appear in the appendix. We followed up with companies in January 2025 to confirm our findings. In some cases, companies told us their data practices were more protective of consumers than their privacy policies appeared to allow. Companies' full responses can be found in the appendix.

Findings

For four of the six products in our test set, we were able to easily create a clone based on publicly available audio. These products did not employ any technical mechanism to ensure we had the speaker's consent to generate a clone or to limit the cloning to the user's own voice. These companies—ElevenLabs, Speechify, PlayHT, and Lovo—required only that we check a box confirming that we had the legal right to clone the voice or make a similar self-attestation.

The other two companies, Descript and Resemble AI, took steps to make it more difficult for customers to misuse their products by creating a non-consensual voice clone.

Descript's voice cloning tool required us to read and record a consent statement, then used the resulting audio to generate the clone. This method is not foolproof: With any voice cloning platform that uses this safeguard, it would be theoretically possible to make a non-consensual voice clone on a second platform and then use that clone to generate the consent statement audio required by the first platform. This was demonstrated by journalists for Proof News in their reporting on AI voice cloning tools.[29] But in that instance, the second platform—the one that generated the original non-consensual voice clone—is the larger problem.

Resemble AI has a slightly different setup for preventing misuse. The product required a user's first voice clone to be based on audio recorded in real time, rather than uploaded audio. But our tester was able to hit record, then play a clip of publicly available audio from another device while recording on their laptop. This process could enable a user to clone someone's voice without their consent; though in our test case the resulting clone was not a compelling impersonation, probably because the quality of the audio had been degraded with each successive recording.

Resemble AI confirmed that using an audio recording from a separate device does degrade the quality of the audio, and "that is a known factor we have prepared for." The company added that it has "safeguards in place that would intentionally limit the quality of the unauthorized clone attempt."

After that first voice clone, subsequent voice clones made with Resemble AI required additional safeguards, including requiring the user to record or upload a recording of a consent statement, which would then be matched against other uploaded or recorded audio to confirm the speaker's consent. When we attempted to clone a CR employee's voice by uploading publicly available audio of that employee speaking, paired with a recording of that same employee reading the consent statement, Resemble rejected our attempt, incorrectly asserting the two audio clips were from different people. We asked Resemble AI about this, and it said that its "algorithm is tuned to be more cautious to prevent misuse."

[29] Janus Rose, "AI Tools Make It Easy to Clone Someone's Voice Without Consent," *Proof News,* June 25, 2024, https://www.proofnews.org/ai-tools-make-it-easy-to-clone-someones-voice-without-consent.

Businesses often gather and verify data about their customers' identity; these practices are referred to as "know your customer." Knowing the true identity of customers is a way to reduce misuse and fraud; it allows companies to investigate attempts to misuse their products and take action, and helps ensure that customers who have had accounts shut down in the past due to misuse can't easily make new accounts under a fake name. Four of the six companies (Speechify, Lovo, PlayHT, and Descript) required only a customer's name and/or email address to make an account, making it easy for customers to enter false names and email addresses. In practice, this means those companies know little about the identity of the consumers who can misuse their products to scam and defraud. The other two tested companies, Resemble AI and ElevenLabs, also required credit card information, which is somewhat more challenging to falsify. Other industries, such as the financial services industry, have much more intensive know-your-customer regulations when opening accounts to mitigate the risk of crime.

The table below shows a condensed version of our findings:

AI VOICE CLONING COMPANIES						
	ELEVEN LABS	**SPEECHIFY**	**LOVO**	**PLAYHT**	**DESCRIPT**	**RESEMBLE AI**
What customer information is required before you can create a custom voice clone?	**First name, email address, credit card information**	**Email**	**Name, email**	**Name, email**	**Name, email**	**Name, email, payment information**
How much does it cost to create a custom voice cloning?	**$5**	**$0**	**$0**	**$0**	**$0**	**$1**
Are there technological barriers to frustrate non-consensual cloning?	**No; users presented with checkbox to confirm they will abide by terms of service**	**No; users just enter a full name after a self-certifying statement confirming they are abiding by the company's terms**	**No; users presented with checkbox to confirm they will abide by the company's terms**	**No; users presented with checkbox to confirm they will abide by the company's terms**	**Yes; users must record or upload audio of an authorization statement that also trains the clone**	**For customer's first voice clone, any audio the product can record live works. For each subsequent voice clone, a consent statement must be recorded or uploaded.**

AI VOICE CLONING COMPANIES						
	ELEVEN LABS	**SPEECHIFY**	**LOVO**	**PLAYHT**	**DESCRIPT**	**RESEMBLE AI**
Does the privacy policy give the company permission to use customer voices to train or improve its model?	**Yes, with opt-out**	**Yes, without offering an opt-out**	**No**	**Unclear**	**Yes. No opt-out of having your voice model used for potential model training, but the company does offer an opt-out for having your projects used to improve the service.**	**Yes, without offering an opt-out.**
Does the privacy policy permit the company to allow other customers or companies to use your voice model or data?	**Yes, if you opt in. Those who opt in are compensated.**	**Yes, with user's consent**	**No**	**Unclear. The policy says the company can do anything with your consent and to further legitimate business interests.**	**Yes, with user's consent**	**No**
Does the privacy policy grant users the right to delete all your voice data and other personal data?	**Yes, though possibly only in jurisdictions that offer data deletion rights**	**Yes, though the company does retain some data for "legitimate" business interests**	**Yes, with minor exceptions**	**Yes, with minor exceptions**	**Yes, though the company does retain some data for "legitimate" business interests**	**Yes, though the company does retain some data for "legitimate" business interests**

Additional information from companies

We asked companies questions about how they safeguard their products and how they handle customer data. Their full responses can be found in the appendix. Some key safeguards mentioned included:

- **Preventing the creation of non-consensual voice clones.** A Descript representative said users have to read prepared consent statements; this is intended to prevent anyone from using an existing recording to create an AI clone of someone else's voice. Descript also said that a portion of consent statements are audited by humans. ElevenLabs shared that it has developed a "No-Go Voice" technology, which blocks the voices of hundreds of public figures from being created. ElevenLabs also pointed to a safeguard it implemented for its higher-quality Professional Voice Clones, Voice CAPTCHA technology, which verifies voices by requiring users to pass a voice similarity test before creating a clone.[30]
- **Post-hoc detection.** ElevenLabs said it is implementing standards from the Coalition for Content Provenance and Authenticity by embedding cryptographically signed metadata into the audio generated on its platform. It has also released a publicly available classifier that assesses the probability that a clip of audio was made by ElevenLabs. Resemble AI shared that it embeds a watermark in all of the audio it generates.
- **Moderation.** ElevenLabs shared that it uses classifiers to detect harmful content—such as content related to scams and fraud, sexual content involving minors, and content related to self-harm—and then uses a combination of automated moderation and human review to take action against violating content. Resemble AI shared that it allows users to report harmful or inappropriate content, which is then reviewed by its team for appropriate action.

Lovo, Speechify, and PlayHT did not respond to CR's outreach.

Policy Analysis

Voice cloning companies have a legal responsibility to limit the risks

Section 5 of the Federal Trade Commission Act prohibits companies from engaging in unfair business practices. The FTC uses a three-part test to determine whether a business practice is unfair under the FTC Act. Under the FTC's unfairness test, a practice is considered unfair if (1) it causes a consumer injury that is substantial, (2) the injury is one that consumers cannot reasonably avoid, and (3) the consumer injury is not outweighed by countervailing benefits to consumers or competition that the business practice generates.[31] Below, we submit voice cloning tools that lack safeguards against misuse to the FTC's three-part test.

1. *Do safeguard-free voice cloning tools cause substantial injury?*

Yes, the injuries caused are substantial. Individual consumers can lose thousands of dollars or more to scammers who use voice impersonating tools. In 2023 alone, 853,935 reports from

[30] "ElevenLabs Voices: A comprehensive guide," *Eleven Labs,* https://elevenlabs.io/voice-guide.

[31] Federal Trade Commission Act of 1914, 15 U.S. Code § 45 (n) (1994), https://uscode.house.gov/view.xhtml?req=granuleid%3AUSC-prelim-title15-chapter2-subchapter1&edition=prelim.

consumers about imposter scams were reported to the FTC's Consumer Sentinel Network, with monetary losses totaling nearly $2.7 billion.[32] The network does not track what proportion of those scams involve voice cloning. But losses from scams where fraudsters posed as government or business representatives more than tripled between 2020 and 2023.[33] And anecdotal evidence of scams involving voice cloning has been mounting. As noted previously, a New York Times investigation of the proliferation of "Elon Musk" deepfakes endorsing fraudulent investment schemes identified one consumer who was scammed out of more than $690,000 in retirement funds.[34] CBS identified another consumer who was scammed out of $10,000 after seeing a deepfake "Elon Musk" endorse investment opportunities.[35] In March of 2023, The Washington Post reported on parents who lost more than $15,000 after purportedly receiving a call from a lawyer who told them their son had killed a U.S. diplomat in a car accident, and hearing what they thought was their son's voice on the phone.[36]

2. *Can consumers reasonably avoid being harmed by safeguard-free voice cloning tools?*

No, they cannot. Consumers cannot reasonably avoid the possibility of being surreptitiously recorded or having clips of their voice shared publicly online. Nor can they reasonably avoid being targeted by an AI voice-powered scam or being tricked by AI celebrity deepfakes online. The very point of these impersonation attacks is to deceive people. While some deepfakes may include telltale signs of inauthenticity, as AI becomes more sophisticated it will be harder and harder for even experts to distinguish between real and fake content without technical assistance.

3. *Do the consumer benefits of safeguard-free voice cloning outweigh the harms?*

No, they do not. Many of the pro-social or neutral uses of AI voice tools do not require custom voice clones. For example, a Substack writer can gain most of the benefits of AI voice tools—cheap and fast audio narration of their articles—by using a preloaded AI voice rather than a clone of their own voice. Similarly, consumers could derive much of the value of real-time AI voice-to-voice language translation without generating a voice that sounds like their own.

A stronger case for custom AI voice clones can be made for other pro-social or neutral uses, such as movie dubbing or customizing ads to different markets. But the benefits of most of these

---

[32] Federal Trade Commission, "Consumer Sentinel Network Data Book 2023," February 2024. Available at: https://www.ftc.gov/system/files/ftc_gov/pdf/CSN-Annual-Data-Book-2023.pdf.
[33] "Impersonation scams: not what they used to be," *Federal Trade Commission,* April 1, 2024, https://www.ftc.gov/news-events/data-visualizations/data-spotlight/2024/04/impersonation-scams-not-what-they-used-be#ftn4.
[34] Stuart A. Thompson, "How 'Deepfake Elon Musk' Became the Internet's Biggest Scammer," *The New York Times,* August 14, 2024, https://www.nytimes.com/interactive/2024/08/14/technology/elon-musk-ai-deepfake-scam.html.
[35] Brian New, Lexi Salazar, Mike Lozano, Scott Fralicks, "Deepfakes of Elon Musk are contributing to billions of dollars in fraud losses in the U.S.," *CBS News Texas,* November 24, 2024, https://www.cbsnews.com/texas/news/deepfakes-ai-fraud-elon-musk.
[36] Pranshu Verma, "They thought loved ones were calling for help. It was an AI scam," *The Washington Post,* March 5, 2023, https://www.washingtonpost.com/technology/2023/03/05/ai-voice-scam.

AI voice cloning applications would be enjoyed most directly by businesses rather than consumers. If businesses save money by using custom AI voice clones, rather than hiring voice actors, they may choose to pass those savings along to consumers—or they may not. However, any benefit to consumers will be indirect.

In weighing the consumer costs and benefits of AI voice cloning, it's also important to note that we should not be comparing a world with voice cloning to a world without it. Instead, we should be comparing, on one hand, a scenario in which AI voice cloning tools make it cheap and easy to create non-consensual voice clones and, on the other, one where companies that offer voice cloning tools supervise the use of those tools, know and work directly with their clients, and require clients to agree in advance to use the technology appropriately. The latter approach may be more expensive but may also be more affordable than paying human voice actors to record and re-track long scripts, while posing a significantly lower risk of harm to consumers.

There is an important voice cloning use case for a small share of consumers: providing assistance to people who have lost the ability to speak due to a medical condition. AI companies often point to this use case when promoting AI voice cloning tools.[37] However, serving those consumers does not require that voice cloning be widely available to the public. Organizations serving those consumers can collaborate with patients and medical professionals to tailor the product to their needs.

Consumers and businesses can derive many of the benefits of AI voice tools with a combination of less risky products: AI voice tools that offer only generic voices and supervised, safeguarded AI voice cloning. On the flip side, the costs to consumers of widely available, easy-to-use AI voice cloning tools with limited safeguards are clear: the proliferation of fake celebrity endorsements and misinformation, hyper-persuasive scams, and the possibility that a consumer's own voice will be used to damage their reputation or defraud them.

It's also important to highlight a precedent for the FTC bringing an action against a generative AI company for failure to implement reasonable safeguards against abuse. In 2024, the agency brought a complaint against Rytr, a company whose AI technology could be used to generate thousands of detailed but fake product and service reviews and testimonials based on a few simple inputs.[38] In December 2024, the agency prohibited the company from advertising or selling any service related to generating consumer reviews.[39] Consumer Reports submitted comments in support of the agency's proposed settlement with Rytr, arguing that companies

[37] See, for example, "Navigating the Challenges and Opportunities of Synthetic Voices," *OpenAI,* March 29, 2024, https://openai.com/index/navigating-the-challenges-and-opportunities-of-synthetic-voices, and "Impact Program: On a mission to help 1 million people reclaim their voice," *ElevenLabs,* https://elevenlabs.io/impact-program.
[38] Complaint, Rytr LLC, FTC Docket No. 232-3052 (December 18, 2024), www.ftc.gov/system/files/ftc_gov/pdf/2323052rytrcomplaint.pdf.
[39] "FTC Approves Final Order against Rytr, Seller of an AI 'Testimonial & Review' Service, for Providing Subscribers with Means to Generate False and Deceptive Reviews," *Federal Trade Commission,* December 18, 2024, https://www.ftc.gov/news-events/news/press-releases/2024/12/ftc-approves-final-order-against-rytr-seller-ai-testimonial-review-service-providing-subscribers.

have a legal obligation to protect their products from being used for harm, and pointing to the agency's extensive history of data security cases with similar findings.[40]

Post-hoc detection alone is insufficient

An important strand of AI research focuses on devising methods to "watermark" AI-generated content so that it can be detected later. Effectively watermarked AI audio paired with detection tools could help identify misleading AI-generated audio. Researchers at companies and universities are already working on techniques for watermarking AI audio. ElevenLabs, for instance, offers a tool that allows users to determine the probability that an audio clip was created by its voice cloning product as mentioned above.[41] And a group of major companies has formed the Coalition for Content Provenance and Authenticity (C2PA) and created an "open technical standard providing publishers, creators, and consumers the ability to trace the origin of different types of media."[42]

But post-hoc detection alone is not a sufficient solution to the problem. One reason is that watermarks can be removed or tampered with.[43] More importantly, someone persuaded by a voice on the telephone that a family member is in danger is unlikely to have the presence of mind or the technical sophistication to check the audio against an AI-detection tool in real time. (And while it may be possible to monitor phone calls for AI-generated voices, such technology may come at a steep cost to privacy.)

For these reasons, post-hoc detection techniques must be paired with continued consumer education around common scams, as well as safeguards that reduce the likelihood that fraudulent voice clones are created in the first place.

Imperfect safeguards are much better than none

One objection to the argument that AI companies should put greater safeguards in place is that open-source developers and hobbyists can develop AI voice cloning tools and post the code online. Even if safeguards become the norm, this argument goes, open-source projects could enable sufficiently motivated actors to create safeguard-free tools that could be used to perpetrate fraud or spread misinformation. Foreign governments unconcerned with American laws could also create their own AI voice cloning tools to spread misinformation in the U.S.

---

[40] Justin Brookman, Matt Schwartz, Grace Gedye, "Consumer Reports files comment in support of FTC's settlement with Rytr," *Consumer Reports Advocacy,* November 4, 2024, bit.ly/3X7qwLp.
[41] "Detect whether an audio clip was created using ElevenLabs," *ElevenLabs,* https://elevenlabs.io/ai-speech-classifier.
[42] "An open technical standard providing publishers, creators, and consumers the ability to trace the origin of different types of media," *Coalition for Content Provenance and Authenticity,* https://c2pa.org.
[43] Jacob Hoffman-Andrews, "AI Watermarking Won't Curb Disinformation," *Electronic Frontier Foundation,* January 5, 2024, https://www.eff.org/deeplinks/2024/01/ai-watermarking-wont-curb-disinformation.

Of course, no set of safeguards will prevent all impersonation, scams, and misinformation enabled by AI voice cloning. But increasing the amount of time and money required to conduct such fraud is an important and meaningful way to limit it. Scammers forced to use open-source tools or tools created outside the reach of U.S. laws may not be as effective or efficient at separating consumers from their money. Moreover, companies' legal responsibilities don't necessarily evaporate when they post their code online and allow other people to reproduce it. The same three-part test the FTC uses to decide whether a business practice is unfair likely applies to businesses' decisions to make certain products freely available. Finally, emerging content labeling and provenance standards, like C2PA, could eventually enable platforms to block or warn users about content developed by companies or people who do not use industry-standard identification schemes, making it harder for scammers to use untrusted content generation tools.

Policy Recommendations

Better practices for AI voice cloning companies

CR consulted with external experts[44] and collaborated internally to generate a set of practices that we feel should be adopted by companies developing and making AI voice cloning products available. These recommended practices are meant to be a starting point and are not comprehensive.

Policies that address fraud, deception, and impersonation

- At a bare minimum, companies must:
  - Have mechanisms and protocols in place to confirm the consent of the speaker whose voice is being cloned, such as by requiring users to upload audio of a unique script.
  - Collect customers' credit card information, along with their names and emails, as a basic know-your-customer practice so that fraudulent audio can be traced back to specific users.
- AI voice companies should watermark AI-generated audio for future detection and update their marking technique as research on best practices progresses.
- AI voice companies should provide a tool that detects whether audio was generated by their own products.
- AI voice companies should detect and prevent the unauthorized creation of clones based on the voices of influential figures, including celebrities and political figures.

[44] We are grateful to the experts who shared their time and expertise with us. See the Acknowledgments section for a full list.

- AI voice cloning companies should build so-called semantic guardrails into their cloning tools. These should automatically flag and prohibit the creation of audio containing phrases commonly used in scams and fraud and other forms of content likely to cause harm, such as sexual content.
- Companies should consider supervising AI voice cloning, rather than offering do-it-yourself voice products. Supervised voice cloning might entail the AI company meeting its clients, understanding their use cases, and ensuring that the company, the client, and anyone whose voice will be used consent to how the tool will be used. The AI voice company might also ensure that access to the voice model is limited to necessary actors and enter into a contractual agreement about which entity is liable if the voice model is misused. The AI voice company could delete the model once the project is complete.

Additional policy recommendations relating to privacy and cybersecurity:

- **Limit the use of consumer information according to data minimization principles.** Customer payment information should be used only to facilitate payment and verify identity. Customer information and voice data should be used only to provide the requested AI voice service.
- **Give consumers a non-retention option.** Companies should provide consumers with the option to have their voice data and clone deleted when they exit the app or retained only for a specific number of days.
- **Do not reuse voice models.** Companies should affirmatively promise consumers that they will never sell, share, or allow someone else to use their voice data or clones and that only the customer or someone designated by the customer may use them. Companies should not reuse or resell customer voices or voice models.
- **Do not reuse other customer data.** Companies should not reuse or resell other customer data, such as email addresses, location, occupation, or other information that they might collect.
- **Use reasonable cybersecurity practices.** Companies should use reasonable safeguards to stop attackers from gaining access to personal information and voice models. One example: a vulnerability disclosure program that enables researchers to easily and securely disclose security vulnerabilities they discover.[45]
- **Use consent-based training data.** Companies should train their models using only the voices of people who have given consent to have their voice data used in this way.

[45] For more, see Stacey Higginbotham, "Who Ya' Gonna Call? Why IoT Companies Should Embrace Vulnerability Disclosure Programs," *Consumer Reports Innovation,* July 29, 2024, https://innovation.consumerreports.org/who-ya-gonna-call.

More robust enforcement and new rules

Current policy context

Existing laws apply to the use of AI voice tools for fraud and impersonation, and new proposals are in the works. As discussed above, the Federal Trade Commission's Section 5 unfairness authority appears to apply to several AI voice cloning products that lack robust safeguards against misuse. In addition, in April 2024, a new FTC rule went into effect prohibiting the impersonation of business or government entities and giving the agency stronger tools to combat scammers.[46] The FTC also proposed amending that rule and extending those protections to the impersonation of individuals, calling attention to the threat of AI deepfakes.[47] As part of the proposed amendment, the agency also sought comments on whether the rule should declare it unlawful for a company to provide goods or services "that they know or have reason to know [are] being used to harm consumers through impersonation." The Federal Communications Commission has also proposed rules that would require callers to disclose the use of AI on AI-generated calls, but that rule has not gone into effect.[48]

Many states have "right of publicity" laws that generally entitle individuals to control the commercial use of their image, voice, and likeness.[49] Some states have begun updating their right of publicity laws with generative AI in mind. Tennessee, for example, passed the ELVIS Act, which, among other provisions, clarified that "voice" means "a sound … that is readily identifiable and attributable to a particular individual, regardless of whether the sound contains the actual voice or a simulation of the voice of the individual."[50] The ELVIS Act created liability not only for people who make non-consensual deepfakes but also for those who create any "algorithm, software, tool, or other technology, service, or device, the primary purpose or function of which is the production of an individual's photograph, voice, or likeness without authorization from the individual."

Stronger enforcers, new rules, and new laws

Many of the harms outlined in this report could be addressed under the Federal Trade Commission's Section 5 unfairness authority. But in order to change companies' practices, the

---

[46] Federal Trade Commission, "FTC Announces Impersonation Rule Goes into Effect Today," press release, April 1, 2024, https://www.ftc.gov/news-events/news/press-releases/2024/04/ftc-announces-impersonation-rule-goes-effect-today.
[47] Federal Trade Commission, "FTC Proposes New Protections to Combat AI Impersonation of Individuals," press release, February 15, 2024, https://www.ftc.gov/news-events/news/press-releases/2024/02/ftc-proposes-new-protections-combat-ai-impersonation-individuals.
[48] "FCC Proposes First AI-Generated Robocall & Robotext Rules," *Federal Communications Commission,* August 8, 2024, https://www.fcc.gov/document/fcc-proposes-first-ai-generated-robocall-robotext-rules-0.
[49] "Right of Publicity Statutes & Interactive Map," *Right of Publicity,* https://rightofpublicity.com/statutes.
[50] H.B. 2091 (Tenn. 2024). https://www.capitol.tn.gov/Bills/113/Bill/HB2091.pdf.

FTC would need to take high-impact enforcement actions, which in turn requires a well-resourced, empowered FTC.

Unfortunately, the FTC is currently severely hamstrung by budget constraints. In 2024, it had only 1,292 full-time employees to pursue its competition and consumer protection missions.[51] This represents a decrease from 1,746 full-time employees in 1979. The economy has more than tripled in size during that time, while the agency's capacity has decreased by more than a quarter.[52] As of this writing, it is unclear whether the Trump administration will attempt to enact further staffing or budget cuts at the FTC.

Even when the FTC does manage to bring a case, it often cannot get meaningful relief from the wrongdoer. For most violations of Section 5, the agency cannot get statutory penalties from offenders. Historically, the FTC was at least able to obtain restitution—to get back the money that consumers lost to fraudsters. However, in 2021, the Supreme Court held that the agency's enabling statute does not give it that limited authority in many instances.[53] Despite bipartisan agreement that the FTC should at least be empowered to force wrongdoers to disgorge ill-gotten gains, Congress has failed to enact legislation to restore that power.[54]

To ensure that the U.S. is poised to counter AI-powered scams, Congress should grant the FTC additional resources to hire attorneys and technologists and expand its legal powers.

Despite limited resources, the agency has brought actions against companies that sell AI technology that can be used in deceptive and unfair ways.[55] The behaviors uncovered in this paper constitute Section 5 violations, and we encourage regulators to take actions against companies with limited safeguards. The agency could also provide clarity by promulgating rules for AI voice cloning.

The FTC is not the only agency that should take action; many state attorneys general can and should bring cases under state consumer protection laws as well. More than half of states'

---

[51] "FTC Appropriation and Full-Time Equivalent (FTE) History," *Federal Trade Commission,* https://www.ftc.gov/about-ftc/bureaus-offices/office-executive-director/financial-management-office/ftc-appropriation.
[52] "Real Gross Domestic Product," *Federal Reserve Bank of St. Louis,* January 30, 2025, https://fred.stlouisfed.org/series/GDPC1.
[53] *AMG Capital Management, LLC v. Federal Trade Commission,* 141 S. Ct. 1341 (2021), https://www.supremecourt.gov/opinions/20pdf/19-508_l6gn.pdf.
[54] Testimony of Anna Laitin, Director, Financial Fairness and Legislative Strategy, Consumer Reports, Before the House of Representatives Committee on Energy & Commerce Subcommittee on Consumer Protection and Commerce on "The Consumer Protection and Recovery Act: Returning Money to Defrauded Consumers," (April 27, 2021), https://www.congress.gov/117/meeting/house/112501/witnesses/HHRG-117-IF17-Wstate-LaitinA-20210427.pdf.
[55] Federal Trade Commission, "FTC Announces Crackdown on Deceptive AI Claims and Schemes," press release, September 25, 2024, https://www.ftc.gov/news-events/news/press-releases/2024/09/ftc-announces-crackdown-deceptive-ai-claims-schemes.

Unfair and Deceptive Acts and Practices statutes contain broad prohibitions on unfairness, according to an analysis from the National Consumer Law Center.[56]

While the FTC and state attorneys general likely already have the power to take action against AI voice cloning companies with minimal safeguards, new legislation would clarify the situation for companies. Legislation could prohibit companies from offering, distributing, or making a voice cloning product available unless they follow basic best practices, such as the ones outlined above. Such legislation could be introduced at the federal or state level. In the face of gridlock in Congress, states have moved quickly on AI policy, with a particular focus on consumer protection.

Appendix:

Supplemental information provided by companies

Descript Responses

Do you mark or modify the audio your product generates so that it can later be detected as AI-generated content?
At present, Descript is not aware of any effective watermarking methods that are not easily circumvented but we continue to assess new tools/technologies as they may be developed.

Do you attempt to detect and prevent the creation of any kinds of harmful content? (Sometimes this is referred to as semantic guardrails.) If so, please elaborate.
We believe that our customers should own their content, and that monitoring what they create would violate their privacy. We're not a distribution platform, we're a creation platform—so just like Microsoft doesn't monitor what you write in Word, we don't monitor the videos or podcasts created in Descript.

Do you have a vulnerability disclosure program? If yes, please provide a link where we can find information about it.
We do not.

When people upload audio to make a voice clone, does your company also use that audio to train your underlying model?

[56] Carolyn L. Carter, "Consumer Protection in the States: A 50-State Report on Unfair and Deceptive Acts and Practices Statutes," *National Consumer Law Center,* February 2009, https://www.nclc.org/wp-content/uploads/2022/08/report_50_states.pdf.

We present people with a data sharing opt-in during initial onboarding, and if they do not opt in, their data is not used for AI training. They also have the option to change this preference later on. More information on how here:
https://help.descript.com/hc/en-us/articles/10255866490125-Account-data-and-privacy

If your company uses consumers' audio to train its underlying model, do you offer consumers the option to opt out?
Customers can opt out of data sharing at any time. More information on how here:
https://help.descript.com/hc/en-us/articles/10255866490125-Account-data-and-privacy

Do you share, sell, or make available consumers' voice recordings, uploaded audio, or synthetic media designed to sound like a particular customer?
We do not.

Do you have any safeguards in place designed to prevent the creation of non-consensual or fraudulent voice clones that you would like to tell us about?
First, as with any tech product, all Descript users have to review and sign our terms of service, which prohibit impersonating others or attempting to clone the voice of anyone who hasn't consented to it. If we learn that a customer has violated those terms, or has used Descript in connection with any illegal activity, we will deactivate their account.

Additionally, we've developed software to validate that a customer's voice is their own, and a process for training speech models that depends on verbal consent verification – meaning our customers can only create text-to-speech models that have been authorized by the voice's owner.

1. To use Descript's AI voices, users have to a) record themselves reading a prepared consent statement affirming that they want Descript to create an AI version of their voice and b) agree to authorize Descript to create that voice. The consent statement is intended to prevent anyone from using an existing recording to create an AI clone of someone else's voice; we preserve it for verification.
2. As a further safeguard, when a customer uses an AI voice to make a change to existing content, Descript compares the audio being edited to their recorded consent statement before completing the action.
3. A portion of our consent statements are audited by a human to verify that the system is working.

IIElevenLabs

ElevenLabs' Responses to first round of questions

Do you mark or modify the audio your product generates so that it can later be detected as AI-generated content?
Yes. We are implementing Coalition for Content Provenance and Authenticity (C2PA) standards by embedding cryptographically-signed metadata into the audio generated on our platform. We believe that provenance and transparency are critical to AI safety, and are collaborating with industry, governments, and other stakeholders on these efforts, including through membership in the Content Authenticity Initiative (CAI) and the U.S. Artificial Intelligence Safety Institute. We have also released a publicly available classifier to detect ElevenLabs AI-generated content.

Do you attempt to detect and prevent the creation of any kinds of harmful content? (Sometimes this is referred to as semantic guardrails.) If so, please elaborate.
Yes. We use proprietary and third-party classifiers to detect harmful content, including content related to scams and fraud, sexual content involving minors, and content related to self-harm. We rely both on automated moderation and human reviewers to take action against violating content. Audio generated by ElevenLabs is traceable to individual accounts–when we identify content that violates our Prohibited Use Policy, we take decisive action, including removing voices, placing users on probation, banning users from our platform, and if appropriate or required by law, reporting to authorities.

In addition, we have reporting mechanisms in place to enable the public to flag misuse. We also work with external threat intelligence teams that provide insights on potential misuse, and we have set up information sharing channels with government and industry partners.

Do you have a vulnerability disclosure program? If yes, please provide a link where we can find information about it.
We have a vulnerabilities disclosure program for security vulnerabilities identified by ethical hackers. The program is managed through HackerOne and we intake reports through vulnerability_disclosure_program@elevenlabs.io.

Separately, we also have a publicly available form for reporting abusive or prohibited uses of our services, including unauthorized use of voice, impersonation, and fraud. The link to submit a report is here:
https://help.elevenlabs.io/hc/en-us/requests/new?ticket_form_id=13145996177937#

When people upload audio to make a voice clone, does your company also use that audio to train your underlying model?
We use certain data that some customers provide to us to improve the quality of our audio models for everyone, but all customers have the ability to opt out of such use. Information on our training practices is available here,
https://help.elevenlabs.io/hc/en-us/articles/29952728805393-Is-my-data-used-to-improve-ElevenLabs-AI-models
and in our Privacy Policy here https://elevenlabs.io/privacy-policy.

If your company uses consumers' audio to train its underlying model, do you offer consumers the option to opt out?
Yes. Please see our answer to the previous question.

Do you share, sell, or make available consumers' voice recordings, uploaded audio, or synthetic media designed to sound like a particular customer?
No, ElevenLabs does not share or sell customers' created voices or audio. In fact, we have created a marketplace mechanism that allows users to make their own voice available to others through the ElevenLabs Voice Library, for which they can receive compensation when the voice is used. Please find additional information on Voice Library sharing here.
https://elevenlabs.io/vla

Do you have any safeguards in place designed to prevent the creation of non-consensual or fraudulent voice clones that you would like to tell us about?

Our prevention measures include:

- Customer screening. All customers who obtain voice cloning tools must provide contact and payment information, which helps us block accounts linked to fraudulent activity or high-risk geographies at sign-up.

- Voice CAPTCHA. We require self-serve users who wish to create high-quality professional voice clones, as well as those who we have identified to be in violation of our Prohibited Use Policy, to undergo voice CAPTCHA prior to creating a clone.

- "No-Go" Voices. We have developed a No-Go Voice technology that blocks the voices of hundreds of public figures, including celebrities and politicians, from being generated. We are continuously expanding this safeguard to increase its effectiveness, and we monitor and take enforcement actions against users who attempt to generate blocked voices.

In addition to the measures above, we also:

- Monitor for misuse. We use proprietary and third-party classifiers to detect content, including content related to scams and fraud, sexual content involving minors, and content related to self-harm. In addition to automated means, we use a team of human reviewers to evaluate the content. We also work with external threat intelligence teams that provide insights on potential misuse, and we have set up information sharing channels with government and industry partners.

- Enforce against abusive accounts. When we identify misuse, we take decisive action, including removing voices, placing users on probation, banning users from our platform, and if appropriate or required by law, reporting to authorities.

- Support third-party detection tools. We work with third-party AI safety companies to improve their tools for identifying AI-generated content. In July 2024, we partnered with Reality Defender – a cybersecurity company focused on deepfake detection – providing access to our models and data to strengthen their detection tools. Our partnership helps Reality Defender's clients detect AI-generated threats in real time, shielding people around the world from misinformation and sophisticated fraud. We're also involved in other academic and commercial projects, including research initiatives at UC Berkeley's School of Information, to advance AI content detection.

ElevenLabs' responses to second round of questions

Does ElevenLabs' privacy policy allow the company to use customer voices to train/improve the model? If so, is an opt-in or opt-out offered?
In accordance with our Privacy Policy, we use certain data that some customers provide to us to improve the quality of our audio models for everyone, but all customers have the ability to opt out of such use. Information on our training practices is available here, and in our Privacy Policy.

Does ElevenLabs' privacy policy allow anyone else (other customers or other companies) to use customers' voice models or uploaded voice data? If so, does the customer whose voice model or data is being used have to opt in or give consent?
Voices created by an ElevenLabs customer can only be used by others if the customer has chosen to make a voice available in our "Voice Library" – a marketplace we set up on our platform to allow voice actors and other individuals to monetize their voices. Since launching Voice Library payouts in February 2024, we have paid contributors more than $1 million. Those who contribute voices to the Voice Library can remove them at any time. Please find additional information on Voice Library sharing here.

Does ElevenLabs' privacy policy provide consumers the right to delete all their voice data and other personal data? If so, does the policy provide all consumers with this right, or only consumers in certain jurisdictions? Does the policy provide a right to delete but still enable the company to retain customer data for specific purposes? If so, for what purposes?
We honor the data subject rights of all consumers regardless of jurisdiction. Please find additional information on our Data Subject Request Process, including how to submit a request, Here.

Consumer advocates are concerned about harms that can stem from AI voice cloning tools if they are misused to defraud or mislead consumers—particularly via impersonation of a consumer's contacts or influential figures. Do you have any mechanisms or business practices in place to prevent the creation of non-consensual voice clones or to prevent the creation of harmful content? If potentially harmful content is created, do you have mechanisms in place to mitigate harm from it?

As previously submitted in November in response to your questionnaire, our prevention measures include:

- **Customer screening.** All customers who obtain voice cloning tools must provide contact and payment information, which helps us block accounts linked to fraudulent activity or high-risk geographies at sign-up.
- **Voice CAPTCHA.** We require self-serve users who wish to create high-quality Professional Voice Clones, as well as those who we have identified to be in violation of our Prohibited Use Policy, to undergo voice CAPTCHA prior to creating a clone.
- **"No-Go" Voices.** We developed No-Go Voice technology, which blocks the voices of hundreds of public figures, including celebrities and politicians, from being generated. We are continuously expanding this safeguard to increase its effectiveness, and we monitor and take enforcement actions against users who attempt to generate blocked voices.

In addition to the measures above, we also:

- **Monitor for misuse.** We use proprietary and third-party classifiers to detect content, including content related to scams and fraud, sexual content involving minors, and content related to self-harm. In addition to automated means, we use a team of human reviewers to evaluate the content. We also work with external threat intelligence teams that provide insights on potential misuse, and we have set up information sharing channels with government and industry partners.
- **Enforce against abusive accounts.** When we identify misuse, we take decisive action, including removing voices, placing users on probation, banning users from our platform, and if appropriate or required by law, reporting to authorities.
- **Support third-party detection tools.** We work with third-party AI safety companies to improve their tools for identifying AI-generated content. In July, we partnered with Reality Defender – a cybersecurity company focused on deepfake detection – providing access to our models and data to strengthen their detection tools. Our partnership helps Reality Defender's clients detect AI-generated threats in real time, shielding people around the world from misinformation and sophisticated fraud. We're also involved in other academic and commercial projects, including research initiatives at UC Berkeley's School of Information, to advance AI content detection.

Users who wish to defraud, mislead, or impersonate others, or to otherwise misuse your product, may decide to input false names and emails to mask their identities. How do you deal with that concern?

We block any email account associated with violating users that we are able to identify, and we also block email accounts with known fraudulent email domains. In addition, we use third-party providers to run fraud risk analysis on user-submitted information.

CR was able to create a customer account using an email address, name, and credit card for the first month's payment. Beyond a self-attestation that the user

would adhere to the company's terms of service, CR did not find additional technological barriers, such as requiring the speaker's consent when creating a voice clone. Would you like to comment on this finding? (We may edit responses for space.) Are there additional safeguards that we did not find?

We take a risk-based approach designed to keep bad actors off our platform and to prevent abuse, while allowing our users and their audiences to benefit from our technology. We require self-serve users who wish to create high-quality Professional Voice Clones, as well as users we have placed on probation, to undergo live voice verification in the form of passing a voice similarity test known as Voice CAPTCHA prior to creating a clone. We do not currently require all self-serve users who wish to create lower-quality Instant Voice Clones to pass Voice CAPTCHA. However, those users are subject to the range of safety-related technological barriers – discussed in detail above – that create substantial friction for any bad actor seeking to abuse our tools, and drive them to our many commercial and open-source competitors that allow for cloning with little to no friction.

This approach is designed to prevent misuse while allowing our users to utilize our tools in a remarkably wide range of beneficial, benign, and lawful ways. Importantly, many users seeking to clone their own voice have difficulty passing CAPTCHA – for example, if they don't have a good microphone, if they're dealing with background noise, or if their voice has changed. Users who have lost their voice entirely due to illness or accident – among the individuals who benefit from our tools most – simply cannot pass CAPTCHA. And other users have the authorization or legal rights to clone the voice of someone else – for example, a movie studio that holds contractual rights to clone and use a voice or a creator who is producing a First Amendment-protected satire. We do not want to prevent these users from benefitting from our tools, and the suite of industry-leading safeguards that they are subject to, which we are always scrutinizing and working to enhance, have proven very effective.

LOVO

Lovo did not respond to Consumer Reports' additional questions.

PlayHT did not respond to Consumer Reports' additional questions.

Resemble AI Responses

Do you mark or modify the audio your product generates so that it can later be detected as AI-generated content?
Resemble AI offers its Neural Speech Watermarker (https://www.resemble.ai/watermarker/), an "invisible watermark" that tackles the malicious use of AI generated voices. With a deep neural network watermarker, the data is embedded in an imperceptible and difficult-to-detect way, acting as an "invisible watermark."

Do you attempt to detect and prevent the creation of any kinds of harmful content? (Sometimes this is referred to as semantic guardrails.) If so, please elaborate.
At Resemble AI, we have always been committed to upholding the highest standards of ethics (https://www.resemble.ai/ethics/) in the development and application of our AI-powered voice generation technology. As the technology continues to evolve, we recognize the potential for new forms of misuse and continue to implement sophisticated safeguards to protect against unethical voice impersonation.

At the core of our ethical approach is a robust consent system that ensures voice cloning only occurs with the explicit consent of the individual. This built-in consent mechanism is a critical safeguard against unauthorized use and misuse of our technology.

We employ advanced tools like Resemblyzer for speaker identification, which allows us to verify the identity of individuals providing consent for voice cloning. Additionally, our deepfake detector helps us identify and prevent attempts to create fraudulent or misleading content across audio, images and video. By prioritizing consent and leveraging cutting-edge technologies for speaker identification and deepfake detection, we are setting a new standard for ethical AI generation. Our commitment to consent-based voice cloning ensures that individuals retain control over their voice data and protects against unauthorized use.

We also have prohibited uses of our technology and you cannot use AI Voices built by Resemble for:
-Claiming to be from any person, company, administration, or entity without explicit authorization to make this statement and/or impersonating to gain illegal information or privileges;
-Propagating hate speech;
-Discrimination, libel, terrorism, or violent activities;
-Spreading unattributed content or misrepresenting sources;
-Exploiting or manipulating children;
-Making unsolicited phone calls, vast communications, postings, or messages;
-Deceiving or deliberately misleading people.

For example, when recording, Resemble enforces the user to say an array of particular sentences in their own voice. Misuse of this can be easily detected by its algorithm. Once the voice is created, the user owns all rights to that voice. Resemble does not use that voice data to train other models, nor do they resell the voice data to third party companies.

For customized solutions, they work with companies through a rigorous process to make sure that the voice they are cloning is usable by them and have the proper consents in place with voice actors.

We also have a system in place for users to report harmful or inappropriate content. This is reviewed by our team for appropriate action given our Terms of Service. Feedback is also taken into consideration as we improve our automatic AI content policy system.

Do you have a vulnerability disclosure program? If yes, please provide a link where we can find information about it.
We currently do not have a vulnerability disclosure program.

When people upload audio to make a voice clone, does your company also use that audio to train your underlying model?
We do not use that voice data to train other models, nor do we resell the voice data to third party companies.

If your company uses consumers' audio to train its underlying model, do you offer consumers the option to opt out?
Not applicable

Do you share, sell, or make available consumers' voice recordings, uploaded audio, or synthetic media designed to sound like a particular customer?
Resemble AI offers a marketplace with a selection of preset voices created in collaboration with professional voice actors, allowing users to find the perfect voice for their specific needs.

But in the case of individuals, once the voice is created (with explicit consent), the user owns all rights to that voice. Resemble AI does not use that voice data to train other models, nor do they resell the voice data to third party companies. For customized solutions, we work with companies through a rigorous process to make sure that the voice they are cloning is usable by them and have the proper consents in place with voice actors.

Do you have any safeguards in place designed to prevent the creation of non-consensual or fraudulent voice clones that you would like to tell us about?
Yes, at the core of our ethical approach is a robust consent system that ensures voice cloning only occurs with the explicit consent of the individual. This built-in consent mechanism is a critical safeguard against unauthorized use and misuse of our technology.

As we continue to push the boundaries of what is possible with AI voice generation, we remain dedicated to fostering a culture of responsibility, transparency, and trust. Our goal is to harness the power of this technology for good, enabling new forms of creative expression and accessibility while actively working to prevent its misuse.

Other ways we prevent and protect:

Resemble Protect – Resemblyzer
Generative models are advancing at a rapid pace. It is our duty to ensure that we create the right tools to prevent misuse of technology wherever we can. We open-sourced Resemblyzer – a powerful package that uses modern AI and Deep Learning to analyze and compare voices. Resemblyzer will help tackle Fake Speech Detection, Speaker Verification, and Diarization. While it is difficult to prevent all misuses of generative technology, we urge consumers to be evaluative of everything we hear, see, and even read.

Preserving Authenticity: Resemble Watermark
To further our commitment to ethical AI use, we have developed the Resemble Watermark technology. This cutting-edge technology seamlessly integrates the watermark into AI-generated audio, making it unnoticeable to listeners while providing a reliable method to verify the authenticity of AI generated content. It ensures that any piece of content can be verified for its origin, thereby preventing misuse and enabling traceability. The Resemble Watermark is a crucial step towards maintaining transparency and authenticity in the era of generative content, ensuring that our clients can trust the content they consume and create.

Innovation with Integrity: Resemble Detect
In our ongoing effort to combat the misuse of deepfakes, we've introduced Resemble Detect, a cutting-edge tool designed to identify AI generated voices, images and videos. This initiative represents our commitment to ensuring the responsible use of AI voice technologies. Resemble Detect aids in distinguishing between authentic and AI-generated content in real time, providing an additional layer of security and trust for creators and consumers alike. It's a testament to our dedication to upholding ethical standards in the digital realm, reinforcing the authenticity and integrity of digital communications.

Second round of Resemble AI responses

Does Resemble AI's privacy policy allow the company to use customer voices to train/improve the model? If so, is an opt-in or opt-out offered?
We do not use customer voices to train/improve our models.

Does Resemble AI's privacy policy allow anyone else (other customers or other companies) to use customers' voice models or uploaded voice data? If so, does the customer whose voice model or data is being used have to opt in or give consent?

At Resemble AI, we have always been committed to upholding the highest standards of ethics (https://www.resemble.ai/ethics/) in the development and application of our AI-powered voice generation technology. At the core of our ethical approach is a robust consent system that ensures voice cloning only occurs with the explicit consent of the individual. This built-in consent mechanism is a critical safeguard against unauthorized use and misuse of our technology. We do not allow other customers or companies to use customers' voice models or uploaded voice data without explicit consent.

Does Resemble AI's privacy policy provide consumers the right to delete all their voice data and other personal data? If so, does the policy provide all consumers with this right, or only consumers in certain jurisdictions? Does the policy provide a right to delete but still retain customer data for specific purposes? If so, for what purposes?
Yes, Resemble AI does provide consumers the right to request deletion of their voice data and other personal data, no matter the jurisdiction. According to our Terms of Service, Resemble AI does retain and own all right, title and interest in and to such Aggregated Data (unidentifiable and anonymized data), which may be used for troubleshooting, bug fixes, continuous development and improvement of our services and processes, internal learning and training, delivery of support, among other things.

Consumer advocates are concerned about harms that can stem from AI voice cloning tools if they are misused to defraud or mislead consumers—particularly via impersonation of a consumer's contacts or influential figures. Do you have any mechanisms or business practices in place to prevent the creation of non-consensual voice clones or to prevent the creation of harmful content? If potentially harmful content is created, do you have mechanisms in place to mitigate harm from it?
We have several safeguards in place including

i. Prohibiting uses of our technology and you cannot use AI Voices built by Resemble for:
    1. Claiming to be from any person, company, administration, or entity without explicit authorization to make this statement and/or impersonating to gain illegal information or privileges;
    2. Propagating hate speech;
    3. Discrimination, libel, terrorism, or violent activities;
    4. Spreading unattributed content or misrepresenting sources;
    5. Exploiting or manipulating children;
    6. Making unsolicited phone calls, vast communications, postings, or messages;
    7. Deceiving or deliberately misleading people.

ii. We also employ advanced tools like Resemblyzer for speaker identification, which allows us to verify the identity of individuals providing consent for voice cloning. Additionally, our deepfake detector helps us identify and prevent attempts to create fraudulent or misleading content across audio, images and video. By prioritizing consent and leveraging cutting-edge technologies for speaker identification and deepfake detection, we are setting a new standard for

ethical AI generation. Our commitment to consent-based voice cloning ensures that individuals retain control over their voice data and protects against unauthorized use.

Users who wish to defraud, mislead, or impersonate others, or to otherwise misuse your product, may decide to input false names and emails to mask their identities. How do you deal with that concern?
We disable the voice so they cannot use it and we remove/block them from the platform as it violates our ethical approach and Terms of Service.

In previous communications with CR, you told us that Resemble AI offers an "invisible watermark" called Neural Speech Watermarker, which aims to prevent the malicious use of AI-generated voices by embedding data in an imperceptible and difficult-to-detect way. Is this watermark applied to all audio that Resemble AI voice clones generate? Are users able to decide whether it is embedded?
Yes, it is applied to all audio generated.

CR researchers were able to create a Resemble AI customer account by providing a name, email address, and payment information. In order to create an initial voice clone using Resemble AI, we did not have to surmount any additional technological barriers or otherwise confirm the speaker's consent (beyond providing audio recorded live). For subsequent voice clones, we had to record or upload a recording of a consent statement. Does that description accurately reflect a typical Resemble user experience? Would you like to make additional comments on these findings?
We will reiterate that at the core Resemble AI is our ethical approach is a robust consent system that ensures voice cloning only occurs with the explicit consent of the individual. This built-in consent mechanism is a critical safeguard against unauthorized use and misuse of our technology.
We also employ advanced tools like Resemblyzer for speaker identification, which allows us to verify the identity of individuals providing consent for voice cloning. Additionally, our deepfake detector helps us identify and prevent attempts to create fraudulent or misleading content across audio, images and video. By prioritizing consent and leveraging cutting-edge technologies for speaker identification and deepfake detection, we are setting a new standard for ethical AI generation. Our commitment to consent-based voice cloning ensures that individuals retain control over their voice data and protects against unauthorized use.

To make an initial voice clone using Resemble AI, users have to make a live recording. We wondered if it would be possible to use a publicly available voice recording (played on a separate device) to clone someone's voice without their consent. The resulting clone was not compelling, probably because the publicly available audio degraded with each re-recording. Would you like to comment on this finding?
Correct, using audio recording from a separate device does degrade the quality of the audio but that is a known factor we have prepared for. As you have noticed, we do require a consent

statement saying something specific, otherwise our system will not use that voice. We also have safeguards in place that would intentionally limit the quality of the unauthorized clone attempt.

To make subsequent voice clones, Resemble AI requires users to record (or upload a recording of) a consent statement, which is then matched against other uploaded or recorded audio to confirm the speaker's consent. When we attempted to use this feature, however, Resemble AI rejected our attempt, incorrectly asserting that the two audio clips were from different people. Would you like to comment on this finding? Is this an example of a safety measure being too robust, rather than not robust enough?
We are unable to comment on this without having specific details or look into what happened when the two were recorded. Our algorithm is tuned to be more cautious to prevent misuse.

Speechify

Speechify did not respond to Consumer Reports' additional questions.

Acknowledgments

We extend our sincere gratitude to the individuals who shared their expertise with us as we undertook this research. Their input helped us understand the issues surrounding AI voice cloning and helped inform our better practices for companies.

Note that their names do not imply endorsement of this report. We appreciate their engagement.

- Oren Etzioni, Professor Emeritus, University of Washington, Founding CEO, Allen Institute of AI
- Hany Farid, Professor, University of California, Berkeley and Chief Science Officer, GetReal Labs
- Siwei Lyu, Professor of Computer Science and Engineering, University at Buffalo
- V. S. Subrahmanian, Walter P. Murphy Professor of Computer Science, Northwestern University
- Hafiz Malik, Professor of Electrical and Computer Engineering, University of Michigan - Dearborn, Founder and CTO at FakeXpose
- Britt S. Paris, Assistant Professor, Library and Information Science Department / School of Communication and Information, Rutgers University
- Jevin West, Professor, Associate Dean for Research, Information School, University of Washington, Co-Founder, Center for an Informed Public, UW