Testimony of

Carl Szabo

Vice President & General Counsel

NetChoice

before

the United States House Subcommittee on Cybersecurity,
Information Technology and Government Innovation
Hearing on *Addressing Real Harm Caused by Deepfakes*

March 12, 2024

NetChoice[1] is a trade association of leading e-commerce and online companies promoting the value, convenience, and choice of internet business models. Our mission is to make the internet safe for free enterprise and for free expression.

We work to promote the integrity and availability of the global internet and are significantly engaged in issues in the states, in Washington, and in international internet governance organizations.

Introduction

Thank you Chairman Mace, Ranking Member Connolly, and distinguished members of the Subcommittee for the opportunity to testify on the critical issues surrounding deepfakes and artificial intelligence.

My name is Carl Szabo and I serve as the Vice President and General Counsel for NetChoice, a trade association of leading online businesses working to make the internet safe for free enterprise and free expression. I am also an Adjunct Professor at the George Mason University Scalia School of Law.

The Transformative Potential of AI

2023 was a watershed year for artificial intelligence (AI) advancement, and 2024 is set to be the year when many of AI's promises become reality. AI has the potential to profoundly enhance our lives across a wide range of domains, from healthcare and education to productivity and creativity. However, as with any transformative technology, AI also introduces certain risks that we must carefully navigate.

Deepfakes, AI-generated synthetic media that features highly realistic yet false depictions of real people, exemplify both the potential benefits and dangers of AI. On one hand, deepfakes have many positive applications. In education, deepfakes could allow students to interact with historical figures or explore scientific concepts in immersive ways. In healthcare, they have great potential to aid in the treatment of Post Traumatic Stress Disorder. And in the arts, deepfakes open up new avenues for creativity and storytelling. However, deepfakes can also be weaponized as tools for misinformation, fraud, and abuse.

AI is Already Heavily Regulated

While some have called for extensive new regulations on AI, the reality is that this technology is already subject to a wide array of existing laws and regulatory frameworks. Any AI system must comply with the same rules as any other technology or business practice in its sector. This

[1] NetChoice is a trade association of e-Commerce and online businesses, at www.netchoice.org The views expressed here do not necessarily represent the views of every NetChoice member company.

means that AI applications in healthcare are regulated by HIPAA and FDA guidelines, AI in finance is subject to FCRA and ECOA, and AI in education must adhere to FERPA, just to name a few examples. The notion that AI will inhabit some kind of lawless Wild West is simply false.

For example, the federal government has already declared intentional lying about the time, manner, or place of an election to prevent qualified voters from voting a crime. This means the government is free to go after individuals publishing deepfakes that seek to subvert election integrity. Moreover, existing consumer protection laws, such as the FTC Act's prohibition on unfair and deceptive practices, already provide robust safeguards against AI systems that might mislead consumers or otherwise cause them harm.

There are hundreds of laws that govern AI today.

The FTC has made clear that it will vigorously police the AI industry under its existing authorities, and has already brought enforcement actions against companies for making misleading claims about their AI products or failing to secure sensitive data used in AI development. At the same time, broadly applicable anti-discrimination statutes like the Civil Rights Act, Fair Housing Act and Americans with Disabilities Act all constrain the use of AI in high-stakes domains like employment, credit and housing to prevent disparate impacts. Finally, existing defamation and false light torts will protect the subjects of deepfake media from reputational harm.

To be clear, this is not to say that every conceivable AI harm is perfectly addressed by current law, or that thoughtful, targeted updates may not be warranted in certain areas. But the core frameworks for regulating the responsible development and use of AI are very much in place today. Policymakers and the public can take comfort in the fact that our existing legal structures are, by and large, well-equipped to prevent and remedy the highest-risk AI failures.

Before rushing to pass sweeping new AI-specific regulations, we should think carefully about how they would interact with this dense, overlapping web of existing rules. The goal should be to strategically fill discrete gaps, not to create a redundant layer of AI law that could impede innovation while adding little marginal protection for the public.

The Biden Deepfake Robocall Incident: A Case Study in Addressing AI Misuse with Existing Laws

In January 2023, as New Hampshire voters prepared to cast their ballots in the Democratic primary, many received troubling phone calls featuring what sounded like then-candidate Joe Biden announcing his withdrawal from the race due to health concerns.

The calls were quickly revealed to be a hoax; a malicious "deepfake" generated by artificial intelligence tools to deceive voters and disrupt the democratic process. While the Biden deepfake incident illustrates the potential for AI to be abused by bad actors, it also demonstrates how existing laws and collaborative efforts between law enforcement and the tech industry can effectively combat such misuse.

Pindrop, a leading voice authentication company, analyzed the audio and compared it to samples from over 120 known voice synthesis engines. Their deep learning models determined with over 99% confidence that the fake Biden calls had been generated using technology from ElevenLabs, an AI speech generation platform.[2]

Armed with this information, ElevenLabs quickly identified and suspended the account responsible for violating its terms of service, while state and federal authorities launched investigations to uncover and prosecute the perpetrator under existing anti-fraud and election interference statutes.

Existing Laws Against Election Interference

The malicious use of AI tools to influence an election already falls squarely within the scope of numerous state and federal laws. New Hampshire law prohibits knowingly distributing communications that falsely represent a candidate's withdrawal. Federal law prohibits using artificial or prerecorded voice messages in robocalls to cell phones without prior consent. The Federal Trade Commission Act bars unfair or deceptive practices, granting the FTC flexible enforcement authority that encompasses novel forms of digital deception. Finally, using AI to impersonate a real person to defraud voters could constitute wire fraud, punishable by fines and imprisonment.

Some have pointed to the Biden deepfake incident as evidence that AI has outpaced our legal system, necessitating sweeping new regulatory frameworks. In reality, however, the malicious use of AI tools to influence an election already falls squarely within the scope of numerous state and federal laws:

Intentionally deceiving qualified voters to prevent them from voting, with or without deepfake media, is voter suppression—and it is a federal crime.

New Hampshire law, among other states', prohibits "knowingly causing to be distributed, or distributing, a communication that falsely represents that a candidate has withdrawn his or her candidacy."[3] Violations are punishable by fines up to $1,000.

[2] Margi Murphy, Rachel Metz, and Mark Bergen, *AI Startup ElevenLabs Bans Account Blamed for Biden Audio Deepfake*, Bloomberg (Jan. 26, 2024).
[3] NH Rev Stat § 664:14-a.

Federal law prohibits the use of "artificial or prerecorded voice messages" in robocalls to cell phones without prior consent.[4] The Telephone Consumer Protection Act provides for statutory damages of $500-$1,500 per illegal robocall.

More broadly, the Federal Trade Commission Act bars "unfair or deceptive acts or practices in or affecting commerce,"[5] granting the FTC flexible enforcement authority that readily encompasses novel forms of digital deception. The FTC has made clear it will aggressively police misuse of AI systems under its existing powers.

Finally, using AI to impersonate a real person in an attempt to defraud voters could constitute wire fraud under 18 U.S.C. § 1343, punishable by fines and up to 20 years in prison.[6]

The full weight of anti-fraud, consumer protection, and election integrity laws can and should be brought to bear against bad actors who deploy AI tools like deepfakes to deceive and disenfranchise voters. The Biden incident is not a case of legal frameworks struggling to keep pace with technological change, but of the ongoing challenge to ensure existing laws are effectively enforced in the digital domain.

Supplementing Law Enforcement with Cross-Sector Collaboration

Still, the Biden case does highlight the critical role of collaboration between law enforcement and the private sector in detecting and preventing AI-enabled crimes. Pindrop's sophisticated deepfake detection technology, developed through machine learning on massive datasets of real and synthetic audio samples, was instrumental in tracing the fake robocalls back to ElevenLabs' platform.

This kind of public-private partnership will only become more important as generative AI grows ever-more accessible and capable. By continuously honing AI-powered forensic tools to detect misuse, responsible tech companies can serve as valuable allies to law enforcement in identifying AI-generated disinformation, fraud, and other harms.

At the same time, by promptly acting on this information to suspend bad actors and cooperating with investigations, companies like ElevenLabs demonstrate the power of responsible self-governance in the AI ecosystem. ElevenLabs' swift action in this case likely helped limit the reach and impact of the fake Biden calls.

As AI evolves, nurturing this kind of proactive, collaborative approach will be far more effective in safeguarding the public than relying on reactive, potentially innovation-chilling regulations. By fostering close coordination between law enforcement and the AI community, policymakers can help ensure that cutting-edge detection and prevention tools keep pace with emerging threats.

[4] 47 U.S.C. § 227.
[5] 15 U.S.C. § 45.
[6] 18 U.S.C. § 1343.

Resolution

The Biden case highlights the critical role of collaboration between law enforcement and the private sector in detecting and preventing AI-enabled crimes. Pindrop's sophisticated deepfake detection technology was instrumental in tracing the fake robocalls back to ElevenLabs' platform, demonstrating the power of AI-powered forensic tools.[7] ElevenLabs' swift action in suspending the bad actor and cooperating with investigations showcases the importance of responsible self-governance in the AI ecosystem. Nurturing this proactive, collaborative approach will be more effective in safeguarding the public than relying on potentially innovation-chilling regulations.

While the Biden deepfake incident does not obviate the potential need for targeted legislative updates to address AI harms not covered by existing law, such as non-consensual deepfake pornography and AI-manipulated child exploitation material, overly-broad proposals to restrict deepfakes or generative AI could do more harm than good. The incident demonstrates the resilience of our existing legal and collaborative frameworks in addressing AI-enabled threats. By doubling down on enforcement, fostering public-private collaboration, and judiciously updating laws to cover unique AI harms, policymakers can effectively combat malicious deepfakes without compromising the technology's vast beneficial potential.

Mitigating Deepfake Harms

Like with any other technology, bad actors will predictably abuse AI to harass women, sexually exploit minors, in addition to undermining public trust in our democratic processes. Law enforcement is already reporting that abusers are using AI tools to generate realistic depictions of real children in sexual situations, then arguing in court that since the explicit images were "AI-generated," they skirt existing child pornography laws. Elsewhere, criminals are using deepfake technology to falsely depict adults in compromising, sexual situations to extort, defame and intimidate victims.

Though amendments to existing laws may be appropriate to capture harms wrought by deepfake technology under certain circumstances, the vast majority of malicious deepfake uses are already illegal under existing statutes. Laws against harassment, defamation, fraud, identity theft, and copyright infringement all apply to deepfakes, just as they do to any other content. And election laws barring deceptive practices and voter manipulation encompass deepfakes aimed at election interference.

However, Congress should address the gaps that do exist in current law in a careful, targeted fashion.

[7] Margi Murphy, Rachel Metz, and Mark Bergen, *AI Startup ElevenLabs Bans Account Blamed for Biden Audio Deepfake*, Bloomberg (Jan. 26, 2024).

The Stop Deepfake CSAM Act

First, the *Stop Deepfake CSAM Act* would clarify that harmful AI-manipulated sexual images exploiting real minors are unambiguously illegal under existing federal child pornography statutes. Specifically, it would amend the definition of child sexual abuse material (CSAM) to include any visual depiction of a minor engaging in "actual or simulated" sexual conduct, where a criminal has used AI tools to "modify" sexually explicit material to include recognizable features of a real child.

This would prevent abusers from escaping accountability through the perverse argument that digitally manipulated CSAM gets a free pass.

The Stop Non-Consensual Distribution of Intimate Deepfake Media Act

Second, the *Stop Non-Consensual Distribution of Intimate Deepfake Media Act* would update privacy laws to expressly cover identifiable deepfake media shared with intent to harm. It would make it unlawful to distribute a deepfake depicting a non-consenting person engaging in fabricated sexual conduct with intent to coerce, harass or intimidate. This would close a loophole that allows deepfake harassment and exploitation to slip through the cracks.

Importantly, both bills include robust safeguards for constitutionally protected speech. They explicitly exempt works of political commentary, criticism, satire or parody. And they provide a safe harbor for digitally manipulated media that includes a clear disclosure that the content is synthetic. These are the kinds of narrowly tailored legislative updates we need to combat discrete deepfake harms without chilling legitimate expression.

At the same time, these laws alone are not a panacea. To fully address the deepfakes challenge, we also need to equip law enforcement with the expertise and resources to pursue cases involving malicious synthetic media under existing legal frameworks. But by strategically closing loopholes while avoiding rushed, overbroad bans, we can mitigate the worst abuses of deepfakes without stifling innovation.

The Dangers of Imprecise Definitions for AI in Legislation

As legislators grapple with the rapid advancements in AI and its potential impact on society, it is crucial to approach the regulation of this technology with care and precision. Imprecise definitions and overly broad language in AI-related legislation can lead to unintended consequences, stifling innovation and infringing upon free speech rights.

One particularly concerning example is the attempt to require disclosures on political ads that utilize AI technology. This is something we are seeing across the country as several states have enacted what may seem to be important disclosures, but due to imprecise definitions, have dangerous consequences for trust.

Take, for example, House File 2549, introduced in the Iowa General Assembly. This bill, very similar to laws already enacted in states like Michigan, Texas, and Washington, mandates that any published material designed to expressly advocate for or against a candidate or ballot issue must include a disclaimer if it was generated using AI. While the intention behind this requirement may be to promote transparency, the vague definition of AI in the bill could lead to absurd outcomes.

Under the bill's broad definition, AI encompasses any "machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments." This could potentially include even the most basic photo editing tools, such as auto color correction or cropping. As a result, a political ad featuring a candidate standing in front of a green screen or a picture that has been automatically cropped to fit a specific aspect ratio could be subject to the disclosure requirement. This would lead to a proliferation of disclaimers which would render an otherwise truthful political ad labeled with a self-declaration of “fake.” This will cause confusion among voters and undermine the effectiveness of political communication.

Several new state laws would require a “false” label on a political ad featuring a candidate standing in front of a green screen.

Moreover, the bill's penalties for non-compliance are severe, with violators facing up to 90 days in jail and a fine of up to $1,000. This means that individuals or organizations could potentially face criminal charges for minor modifications to images, even if these changes have no material impact on the message being conveyed. Such disproportionate consequences could have a chilling effect on political speech, as campaigns and advocates may be hesitant to use even the most basic digital tools for fear of running afoul of the law.

The issues with imprecise AI definitions extend beyond the political realm. Consider Tennessee’s "Ensuring Likeness, Voice, and Image Security Act of 2024"(ELVIS Act), which seeks to protect individuals' rights to their name, voice, and likeness in the digital era. The original bill's language is so broad that it assumes any use of another person's photograph, voice, or likeness without authorization is subject to civil action. This could lead to unintended consequences, such as holding individuals liable for posting a photograph of a celebrity taken at a concert or sharing a video recording of an artist's performance with a friend via text message.

Furthermore, the ELVIS Act’s First Amendment defense is a rebuttable one, meaning that even if individuals believe their actions are protected under free speech, they would still have to defend themselves in court. This places an undue burden on innocent citizens and could lead to a chilling effect on free expression.

The overly broad language in the ELVIS Act could also restrict the use of AI-powered creative tools, such as those used for generating art, music, or writing. If the legislation fails to clearly distinguish between AI systems that simply assist human creators and those that operate autonomously, it could stifle artistic innovation and limit the ways in which individuals can express themselves.

Fortunately, the lawmakers in Tennessee realized the error of their ways. They moved in a way to include a high mens rea. In addition, the defendant must also be misappropriating the likeness to further a commercial interest. By adding the necessary bad actions, the Tennessee legislature was able to mitigate many of the harms from the overly broad definition of AI.

To avoid these unintended consequences, legislators must work closely with AI experts, industry stakeholders, and civil society groups to craft precise, technology-neutral definitions that focus on specific behaviors and outcomes rather than broad categories of tools. This approach will ensure that the law can adapt to the rapidly evolving AI landscape while still protecting the rights and interests of individuals and society as a whole.

Imprecise definitions can lead to a host of unintended consequences, from chilling political speech to stifling innovation and creativity.

The regulation of AI is a complex and delicate task that requires careful consideration and precise language. Imprecise definitions can lead to a host of unintended consequences, from chilling political speech to stifling innovation and creativity. As we navigate the challenges and opportunities presented by AI, it is essential that our laws strike a balance between protecting the public interest and fostering the responsible development and deployment of this transformative technology. By learning from the examples set by other states, such as Tennessee's amendments to its proposed ELVIS Act, legislators can work towards crafting more thoughtful and effective AI regulations that minimize the potential for legal overreach and unintended consequences.

New federal laws regarding AI can only come from Congress, not the Executive Branch

President Biden's recent executive order on AI[8] is not only a violation of the Constitution but also a misguided attempt to regulate a rapidly evolving technology that holds vast potential to improve people's lives. By attempting to govern AI development through an executive order, the President is effectively usurping the role of Congress, which is the only branch of government constitutionally authorized to create such rules and regulations.

[8] White House, Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (Oct. 30, 2023).

This unilateral action by the executive branch sets a dangerous precedent that undermines the fundamental principles of separation of powers and checks and balances enshrined in our Constitution. It is the responsibility of Congress, not the President, to carefully consider and debate the complex issues surrounding AI regulation and to craft well-balanced legislation that addresses both the challenges and opportunities presented by this transformative technology.

Moreover, the restrictive regulatory approach outlined in the executive order threatens to stifle innovation and hinder the competitiveness of the American AI industry. The introduction of burdensome and complex regulations, without proper congressional oversight, will likely discourage investment in AI research and development, as companies face increased uncertainty and compliance costs. This could lead to a slowdown in innovation, allowing other nations, such as China, to surpass the United States in the global race for AI supremacy.

The White House Executive Order on AI is an unconstitutional violation of the major questions doctrine.

The executive order's broad regulatory measures will result in stifling new companies and competitors from entering the marketplace, effectively consolidating power in the hands of a few large tech giants. This not only limits consumer choice but also significantly expands the power of the federal government over American innovation.

The order puts any investment in AI at risk of being shut down at the whims of government bureaucrats, which is a dangerous approach for our global standing as the leading technological innovators.

It is important to note that there are already many regulations in place that govern AI. Instead of examining how these existing rules can be applied to address modern challenges, President Biden has chosen to further increase the complexity and burden of the federal code. This approach is not only unnecessary but also counterproductive to the responsible development of AI technology.

Only Congress can craft legislation addressing the challenges and opportunities presented by AI. It is crucial that we do not allow fears to hold the United States back from realizing the vast potential of this technology to improve people's lives. It is the responsibility of Congress to ensure that any regulatory framework strikes the right balance between promoting innovation and protecting the public interest, not the executive branch.

The Promise of AI in Education

As an educator myself, I've seen firsthand how AI is transforming the classroom in overwhelmingly positive ways. This year marks an inflection point, as students gained access to powerful generative AI tools that can assist with tasks like essay outlining, open-ended math problems, and code debugging. While this has stoked understandable fears about cheating and plagiarism, I believe AI can be an invaluable teaching aid if properly managed.

In my own classes, I allow the use of AI assistants, but not as a crutch.

In my own classes, I allow the use of AI assistants, but not as a crutch. Students must still demonstrate mastery of the material through thoughtful analysis and their own words. AI can help brainstorm and structure ideas, but it's no substitute for human reasoning, creativity and original expression. By requiring students to cite when they've used AI and reflect on how it contributed to their work, I'm teaching them to use these tools productively while thinking critically about the outputs.

Other AI-powered educational technologies are enabling personalized learning at an unprecedented scale. Intelligent tutoring systems can provide each student with customized feedback and recommendations based on their unique strengths, challenges and learning styles. AI can also automate rote tasks like grading, freeing up teachers to provide more individualized support. And data-driven AI insights are helping schools identify at-risk students early and intervene with targeted resources.

By proactively integrating AI into curricula, we can equip students with the skills to harness these technologies effectively in their academic, personal and eventually professional lives. Attempting to ban AI from the classroom would only leave students woefully unprepared for a future in which these tools are ubiquitous. Instead, we should embrace AI's educational potential while teaching students to be savvy, critical consumers of AI-generated content.

Principles for Trustworthy AI Development

To fully realize AI's positive potential across domains, we must proactively mitigate serious but avoidable negative impacts. But we can do this without heavy-handed government control that could jeopardize America's position as a global leader in AI innovation.

Instead of rushing to restrict AI development itself, policymakers should establish guidelines and incentives for the trustworthy development and deployment of AI systems, focused on three core principles:

1. Transparency

Organizations should commit to disclosing when AI is being used and for what purposes, empowering individuals to make informed decisions about engaging with AI systems. Where AI materially shapes outcomes for consumers, additional context about the key factors influencing the AI's decisions may be warranted.

2. Accountability

AI should be subject to the same rules and liability structures as any other tool. Existing laws, from non-discrimination statutes to product liability and privacy frameworks, already provide robust accountability mechanisms. The key is ensuring these laws are vigorously enforced in the AI context.

3. Security

Rigorous safeguards should be in place to protect the sensitive personal data used to train AI systems from breach or misuse. AI developers must employ state-of-the-art cybersecurity and data governance practices to preserve privacy and prevent AI from amplifying societal biases.

NetChoice

AI Principles to Maintain American Tech Dominance

1: Transparency

Tech users should be informed when content is created by an AI

#2: Accountability

AI is a tool and must be subject to the appropriate existing laws & regulations on consumer protection, privacy, and data security.

#3: Security

AI tools must keep user data secure, and sensitive information should not be used for retaining or retraining AI.

Other Actions for Congress

As an overarching priority, Congress should advance comprehensive federal privacy legislation that enshrines individual data rights, mandates reasonable security practices, and preempts the current patchwork of state laws. By enabling responsible data-driven innovation while protecting consumer privacy across sectors, such a law would provide a much-needed foundation for trustworthy AI development.

At the same time, we must remain vigilant against cynical attempts to control AI for ideological or political purposes. Some efforts to "root out bias," however well-intentioned, risk enshrining a particular partisan worldview and chilling AI's ability to generate the full scope of perspectives present in society.

The best remedy for concerns around subjective bias or "woke AI" is to ensure a vibrant, competitive marketplace of AI products and services. With a diversity of AI models to choose from, reflecting a range of viewpoints and value systems, consumers will be empowered to select the tools that best align with their individual preferences and needs. But sustaining this flourishing AI ecosystem requires a light-touch regulatory approach that leaves ample room for responsible innovation.

Conclusion

We believe the key to addressing deepfakes and unlocking AI's full potential is to pursue a balanced, multi-stakeholder approach. We should strategically update existing legal frameworks for the digital age, empower educators to integrate AI tools thoughtfully into their pedagogy, and encourage voluntary industry initiatives around transparency, accountability and security.

By advancing targeted legislative solutions like the *Stop Deepfake CSAM Act* and the *Stop Non-Consensual Distribution of Intimate Deepfake Media Act*, we can combat the most egregious abuses of synthetic media without impinging on free expression or technological progress. And by establishing clear principles for trustworthy AI development, we can mitigate foreseeable risks while preserving the flexibility for transformative innovation.

AI is not a force to be feared, but a tool to be harnessed wisely in service of human values and aspirations. With the right governance frameworks and social norms in place, the United States can and must retain our global leadership in this critical technological domain. Ceding the AI race to less open societies would not only forfeit the profound benefits for American consumers and businesses, but leave the future trajectory of this powerful technology in the hands of authoritarian regimes.

The choices we make today about how to approach AI governance will shape the fabric of American competitiveness, security and liberty for generations to come. I urge the members of this Subcommittee to reject reactive, heavy-handed proposals that would stymie our capacity to lead the AI revolution, and to instead advance pragmatic, forward-thinking solutions to maximize this technology's positive impact while mitigating its avoidable harms.

Thank you again for the opportunity to share my perspective on these critical issues. We look forward to your questions.

Legislative Language to Address AI Deepfakes

Stop Deepfake CSAM Act

Section 1. Title - This Act may be cited as the Stop Deepfake CSAM Act.

Section 2. Definitions

> "Deepfake" means any visual media created, altered, or otherwise manipulated in a manner that would falsely appear to a reasonable observer to be an authentic record of the individual's actual speech, conduct, or likeness.
>
> "Distribute" means to publish or disseminate, including but not limited to: advertising, exhibiting, exchanging, promoting, or selling deepfake material.
>
> "Minor" means any natural person under eighteen years of age.
>
> "Natural person" means a human being with legal personality as distinguished from a person created by digital means or by operation of law.
>
> "Possess" means [as defined elsewhere in state law].
>
> "Recognizable physical characteristics" means an actual minor's face or likeness.
>
> "Sexual conduct" means unlawful nudity [as defined under state law] or any sexual activity, whether actual or simulated.

Section 3. Prohibitions

The criminal code 18 U.S. Code § 2252 shall be amended in the appropriate place as follows:

> (A) Any person who, with knowledge that the material is a deepfake depicting a minor, knowingly possesses or distributes material that depicts a minor engaging in sexual conduct shall be punished by [insert sentencing guidelines].
>
> (B) Any person who, with knowledge that the material is a deepfake depicting a minor, knowingly distributes, advertises, exhibits, exchanges with, promotes, or sells any material that depicts a minor engaging in sexual conduct shall be punished by [insert sentencing guidelines].
>
> (C) Nothing in this Act shall be construed to impose liability on an interactive computer service, as defined in 47 U.S.C. § 230(f), for information provided by another information content provider.
>
> (D) The provisions of this Act shall not preclude prosecution under any other statute.

Section 4. Severability Clause

If any provision of this Act or the application thereof to any person or circumstance is held unconstitutional or otherwise invalid, the remaining provisions of this Act and the application of such provisions to other persons or circumstances shall not be affected thereby.

Section 5. Effective date

Stop Non-Consensual Distribution Of Intimate Deepfake Media Act

Section 1. Title - This Act may be cited as the Stop Non-Consensual Distribution of Intimate Deepfake Media Act.

Section 2. Definitions

"Deepfake" means any visual media created, altered, or otherwise manipulated in a manner that would falsely appear to a reasonable observer to be an authentic record of a natural person's speech, conduct, or likeness.

"Distribute" means to publish or disseminate, including but not limited to: advertising, exhibiting, exchanging, promoting, or selling deepfake material.

"Natural person" means a human being with legal personality as distinguished from a person created by digital means or by operation of law.

"Non-consensual" means without the voluntary agreement of the natural person whose face is involved.

"Sexual conduct" means actual or simulated sexual intercourse, outercourse, masturbation, bestiality, or sexual sadism.

Section 3. Prohibitions

The civil code of the United States shall be amended in the appropriate place as follows:

Unlawful dissemination or sale of images of another; penalty.

A person commits the unlawful dissemination or sale of images of another if the person:

A. intentionally or knowingly distributes a deepfake with a deepfake with the face of a natural person

Depicting that person engaging in sexual conduct that the actual person did not actually engage in;

with intent to coerce, harass, or intimidate;

where the distributor knows or has reason to know that the distribution was not consented to by the natural person whose face is depicted in the deepfake.

B. Nothing in this Act shall be construed to impose liability on an interactive computer service, as defined in 47 U.S.C. § 230(f), for information provided by another information content provider.

C. The provisions of this Act shall not preclude prosecution under any other statute

Section 4. Rule of Construction

“Deepfake” shall not be interpreted to include any constitutionally protected speech, including works of political, artistic, or newsworthy value, including commentary, criticism, satire, or parody.

Section 5. Safe Harbor

It shall not be a violation of this Act to publish digitally manipulated media that includes a clear disclosure that would cause a reasonable person to understand that the visual media is not a record of a real event.

Section 6. Severability Clause

If any provision of this Act or the application thereof to any person or circumstance is held unconstitutional or otherwise invalid, the remaining provisions of this Act and the application of such provisions to other persons or circumstances shall not be affected thereby.

Section 7. Enforcement

An individual may bring an action under this Act to recover actual or statutory damages, whichever is greater. Juries may consider punitive damages. Statutory damages shall not exceed [suggested limit of $10,000].