# RFI Related to NIST AI Executive order

# Introduction

Resemble AI leads the way in developing generative voice AI responsibly and ethically. As pioneers in the industry, they recognize the huge potential along with possible risks from rapidly advancing voice synthesis and cloning capabilities. With this balanced view, Resemble AI has pioneered innovative solutions to address emerging challenges from unauthorized or unethical voice cloning uses.

To counter the rise of audio deep fakes, Resemble AI puts forward a comprehensive framework combining Resemble AI's Neural Speech Watermarking and Resemble Detect deepfake detection systems. Together, these robustly and ethically prevent and reduce potential harms from misusing voice cloning. The technologies promote transparency, accountability, and authenticating appropriate voice cloning uses.

The framework intervenes at multiple points to protect consumers and enterprises from voice cloning risks. Neural watermarking embeds audio fingerprints directly into AI-generated voices, enabling traceability and attribution. Resemble Detect exposes fake and cloned voices via state-of-the-art deep neural networks optimized for voice authentication.

Resemble AI champions accountability and transparency in creating, using, and monitoring voice cloning tech. This document maps out integrating leading-edge innovations to prevent misuse and chart a responsible way forward. With collaboration, these solutions can help protect consumers from voice cloning harms now and in the future.

# The Emerging Issue of Voice Cloning Risks

Voice cloning tech has advanced swiftly, becoming incredibly sophisticated and accessible. However, lacking adequate safeguards, these innovations also allow for potential misuse. Core dangers include identity theft, scams, reputational harm, copyright violations, and spreading misinformation. The core issue is that voice cloning capabilities have outpaced protective measures. Current laws and policies fail to fully address synthesized media and deep fakes. As abilities improve, the potential damage scale will increase.

These risks arise because voice cloning can synthesize highly realistic speech without needing access to the real person. With enough samples, AI can precisely mimic voices. Modern text-to-speech can also generate natural dialogue.

For example, criminals could clone and impersonate celebrities or politicians to generate fake inflammatory comments or questionable policies. Companies may have branding voices stolen for unauthorized commercial use.

A key challenge is developing preventive protections that intervene before downstream damage occurs. Relying solely on post-hoc remedies for misinformation or fraud is insufficient and reactive. Responsible oversight mechanisms must become integral to the technology itself.

Resemble AI’s proposed framework aims to address these gaps with innovation. By building authentication and deepfake detection directly into the voice cloning pipeline, misuse can be preemptively prevented with minimal consumer burden.

# Proposed Framework

This proposal combines two of Resemble AI's innovative technologies, the Neural Speech Watermarker and Resemble Detect, with a new Deepfake Detection Dashboard. Together, they effectively manage risks from misusing voice cloning.

The Neural Speech Watermarker embeds a unique, imperceptible audio fingerprint into synthesized voices. This fingerprint verifies the content's origin and persists through processing, ensuring traceability.

Working alongside, Resemble Detect uses advanced deep neural networks to identify deepfake or cloned voices in real-time with high accuracy. Integrating the new Deepfake Detection Dashboard enhances user interaction and oversight. The dashboard provides a user-friendly interface for uploading and analyzing audio files, simplifying verification of voice content authenticity.

Combining the Watermarker, Detect, and Dashboard offers a comprehensive voice cloning control solution. The framework makes voice cloning providers embed attribution and validate voices, thereby protecting consumers from AI-voice misuse without their direct involvement. This integrated approach delivers a robust, user-friendly system to prevent and detect voice cloning misuse.

# Advancing Ethical Voice AI

Resemble AI upholds ethical AI principles and compliance frameworks when developing its voice authentication technologies. This supports the company’s commitment to responsible innovation in generative voice AI.

The Neural Speech Watermarker and Resemble Detect reinforce transparency and accountability in voice cloning uses. The watermarking integrally links voice assets to their originators, preventing anonymity that could promote unethical practices.

Meanwhile, Resemble Detect serves as a trusted verification mechanism to authenticate appropriate voice cloning applications. Together they safeguard and oversee against misuse.

These technologies follow data minimization principles. The watermarking uses cryptographic hashing to embed non-reversible attribution without retaining user data. Resemble Detect only extracts the minimal features required to identify synthesized voices.

When integrated by providers, the frameworks promote internal accountability through technical measures rather than just policies or agreements. This failsafe approach prevents unauthorized cloning even internally.

The development process emphasizes “ethics by design” methodologies to engineer beneficial outcomes from the start. Impact assessments help identify high-risk use cases needing safeguarding.

By integrating proactive attribution and verification, Resemble Detect and the Neural Speech Watermarker demonstrate industry leadership in ethical AI based on principled design and scientific rigor. Resemble AI guides the voice cloning ecosystem toward a responsible path that unlocks benefits through accountability.

# Practical Applications

Resemble AI's technologies have been successfully utilized in various practical applications, showing their effectiveness through case studies and customer feedback.

Resemble AI's voice cloning service sees extensive enterprise use for security red-team exercises. Accurately replicating voices assists with simulation-based testing to evaluate and improve defenses. The technology helps uncover potential attack vectors, enabling enhanced security protocols against real voice threats.

The Neural Speech Watermarker and Resemble Detect provide versatile protection for diverse voice cloning uses. Their adaptability allows customized integration by both businesses and consumers.

For media brands, the watermarking secures proprietary assets like branded voices or celebrity likenesses. It fingerprints all AI voices from internal voice cloning before distribution. Resemble Detect then checks for authentic usage externally.

One example is rapidly identifying cloned voices misused for deceptive messaging or slander. Resemble Detect flags the deep fakes to limit spread, while watermarks trace the source.

For telecoms, Resemble Detect scans calls to uncover cloning fraud attempts for enhanced screening. This safeguards customers against pretexting scams.

In gaming, watermarking character voice cloning retains attribution across fan creations while detecting unauthorized usage. Resemble Detect helps players verify authenticity of user-generated audio.

Chatbot users also benefit from Resemble Detect integration. It provides confidence that responses use the authentic system voices rather than external fakes.

For personal use, both technologies directly offer protection. Watermarking retains attribution across any unauthorized cloning while Resemble Detect verifies media authenticity.

These examples showcase the wide applications secured by Resemble's flexible frameworks. Custom configurations address the unique risks faced by different voice cloning uses. Their effectiveness builds on the robustness of the underlying technologies.

# Building a Responsible Future for Voice AI

In summary, these innovations enable intervention at the source by fingerprinting and verifying AI voices, and downstream via large-scale deepfake monitoring. Together they deliver comprehensive protection to safeguard consumers and enterprises.

However, work remains to boost effectiveness and address limitations. Enhancing fingerprinting redundancy and capacity will refine tracing and attribution. Advancing signaling techniques can further improve monitoring capabilities.

Ongoing R&D at Resemble AI focuses on augmenting detection models to stay resilient against evolving voice cloning methods. Adversarial training and neural architecture search help future-proof these systems.

Wider integration poses challenges around aligning incentives and standardizing protections across voice cloning providers. Policies and regulations may help overcome resistance and promote accountability.

Despite barriers, Resemble AI believes proactive self-governance is the most promising path for preventing misuse and building trust. Ethical and responsible innovation focused on consumer benefits must guide voice AI's revolution.

Resemble AI continues pioneering cutting-edge authentication systems to combat voice cloning harms. With collaboration, these solutions can empower society through AI while safeguarding risks.