ARE DEEP SPEECH DENOISING MODELS ROBUST TO ADVERSARIAL NOISE?

Anonymous authors

Paper under double-blind review

ABSTRACT

Deep noise suppression (DNS) models enjoy widespread use throughout a variety of high-stakes speech applications. However, we show that four recent DNS models can each be reduced to outputting unintelligible gibberish through the addition of psychoacoustically hidden adversarial noise, even in low-background-noise and simulated over-the-air settings. For three of the models, a small transcription study with audio and multimedia experts confirms unintelligibility of the attacked audio; simultaneously, an ABX study shows that the adversarial noise is generally imperceptible, with some variance between participants and samples. While we also establish several negative results around targeted attacks and model transfer, our results nevertheless highlight the need for practical countermeasures before open-source DNS systems can be used in safety-critical applications.

1 Introduction

Deep neural networks (DNNs) have found widespread use in speech denoising tasks (herein considered synonymous with *speech enhancement* and *noise suppression*). With existing usage in video-conferencing software (Cutler, 2022) and speech recognition systems (Milling et al., 2024), and potential future usage in hearing aids (Westhausen et al., 2024), the robustness of such deep noise suppression (DNS) models is clearly of paramount concern.

However, it is well-documented that DNNs are often susceptible to adversarial perturbations—slight modifications to the input data that are subtle or imperceptible to humans, but which can lead to dramatically incorrect outputs from DNNs (Szegedy et al., 2014). This vulnerability has been extensively studied in domains such as automatic speech recognition (ASR) (Carlini & Wagner, 2018; Schönherr et al., 2018; Qin et al., 2019) and speaker recognition (Gong & Poellabauer, 2017; Wang et al., 2020), where attacks induce the models to mistranscribe speech or misclassify speakers.

Given their ubiquity and the vulnerability of similar models, we therefore posit that DNS models might also be appealing targets for adversarial attacks. For example, consider a person whose hearing aids run an open-source DNS model with publicly available weights. Many public-address systems rely on fixed phrases ("This is an emergency", "Fire", medication reminders, gate announcements); perhaps a carefully crafted, imperceptible perturbation broadcast in the wearer's vicinity could suppress or garble just those phrases, preventing the user from hearing an evacuation order or life-saving alert. In such scenarios, even an attack surface confined to a few canonical utterances is enough to pose health- and safety-critical risk. Yet this scenario might seem to present limited danger: DNS models are *designed* to remove noise, and prior attacks on speech enhancement models were perceptible, limited to high-noise settings, and might not work over-the-air (Dong et al., 2023).

To the contrary, we show that, across a variety of settings, **psychoacoustically hidden noise can cause four recent DNS models to output unintelligible gibberish**, thus raising urgent security concerns for such open-source DNS models (Schönherr et al., 2018; Defossez et al., 2020; Chen et al., 2022; Zhao et al., 2022b; Lu et al., 2024). The success of our attacks varies little by setting, such as the presence and strength of reverb and normal background noise, and we demonstrate attacks in simulated over-the-air settings. While our attacks are white-box, and we find that naive model transfer fails, this is cold comfort: gradients are always available in open-source DNS models.

We also identify some areas for optimism. First, attacks appear to work best when designed for only a single utterance from the speaker; similar to other audio tasks (Neekhara et al., 2019; Abdoli

et al., 2019; Zhang et al., 2021; Sun et al., 2024), imperceptible universal perturbations (Moosavi-Dezfooli et al., 2017) are not yet possible in this domain. Second, we find that one of the models we test, Full-SubNet+ (Chen et al., 2022), enjoys limited protection from white-box attack due to exploding gradients, though previous research suggests this is easily circumvented (Athalye et al., 2018). Third, *targeted* attacks may succeed according to objective metrics without subjectively producing the target utterance. Finally, simple Gaussian perturbation appears to be a moderately effective baseline defense; nevertheless, the prior success of adaptive attacks (Tramer et al., 2020) highlights that safety-critical DNS systems will need more sophisticated defenses.

Contributions.

- 1. First systematic study of imperceptible adversarial attacks on speech denoising models. Contrary to prior image—denoising results, we show that four state-of-the-art DNS models can be driven to produce unintelligible outputs by perturbations that are psychoacoustically hidden, in speech settings ranging from nearly-clean (70 dB SNR, no reverb) to noisy and reverberant. We also show that auditory masking offsets allow easy tuning of power versus imperceptibility (App. D).
- 2. Evidence: Human study, comprehensive computational measurement, and samples. We demonstrate the success of our core untargeted attacks using three complementary approaches: a) transcription and ABX studies with audio/multimedia experts; b) five distinct computational metrics; c) online samples to allow the reader to evaluate subjective imperceptibility and attack success.
- 3. Masking- and RIR-aware attack framework. We show that clipping to auditory masking thresholds in STFT space offers a ready-made projection operator for projected gradient descent (Madry et al., 2018). To ensure imperceptibility in simulated over-the-air experiments where the perturbation itself is convolved with a non-invertible room-impulse response, we evaluate a combination of Wiener deconvolution and gradient descent-based projection.
- 4. Mechanistic insights: Gradient flow matters more than model size and input features. We show that, contrary to common wisdom (Madry et al., 2018), the size of the models we test matters little for robustness; instead, the only protection we find comes from obfuscated gradients in Full-SubNet+ (Chen et al., 2022), protection which is known to be brittle (Athalye et al., 2018).
- 5. *Practical threat analysis*. Attacks are model- and utterance-specific and require gradient access, but successful over-the-air white-box attacks still preclude the deployment of open-source models in safety-critical applications such as hearing aids without additional defenses. We also quantify the partial protection offered by Gaussian perturbation as a baseline defense.

2 Related Work

Adversarial perturbations for audio models. Adversarial perturbations originated in image classification (Szegedy et al., 2014) and subsequently extended to various DNN-based audio tasks. Like visual attacks, audio perturbations are most commonly studied in classification tasks such as speaker recognition (Gong & Poellabauer, 2017; Zhang et al., 2020; Wang et al., 2020; Chen et al., 2021; Jati et al., 2021; Xie et al., 2021; Shamsabadi et al., 2021; Chen et al., 2023; Li et al., 2020), speaker verification (Kreuk et al., 2018; Zhang et al., 2021), environmental sound classification (Abdoli et al., 2019; Xie et al., 2021; Ntalampiras, 2022), and speech command recognition (Zhang et al., 2017; Xie et al., 2021).

Successful attacks also target audio models with discrete outputs, like ASR systems (Cisse et al., 2017; Alzantot et al., 2018; Carlini & Wagner, 2018; Schönherr et al., 2018; Qin et al., 2019; Neekhara et al., 2019; Sun et al., 2024; Jin et al., 2024). This work includes demonstrating imperceptible targeted attacks that force ASR systems to output attacker-chosen transcriptions.

The success of attacks on complex tasks like ASR raises questions about the vulnerability of models with continuous audio outputs. However, research on attacking generative audio models is limited. Takahashi et al. (2021) and Trinh (2022) demonstrated vulnerability in speech separation models, though the required perturbation perceptibility remains unclear. Dong et al. (2023) first attacked speech-enhancement networks, but their study left open questions regarding: (i) Imperceptibility. They used ∞-norm-bounded, audible perturbations; we use strict psychoacoustic masking (with

¹Samples are available online.

pre/post masking and a -12 to -16 dB offset) for imperceptible samples. (ii) Attack goals. Their attacks were untargeted; we analyze both intelligibility destruction and *targeted* utterance injection. (iii) Real-world coverage. Their experiments covered limited conditions ([-8,8] dB SNR, no reverb, two models). We test four open-source DNS models in diverse conditions (near-clean to noisy, with/without reverb, simulated over-the-air), and report on transferability, universal perturbations, and defenses. Our work thus generalizes and significantly extends the threat landscape they outlined.

Imperceptible audio perturbations. Prior work highlights the complexity of ensuring imperceptible audio perturbations. Standard image perturbations are constrained by p-norms (initially 2-norm (Szegedy et al., 2014; Fawzi et al., 2018), now often ∞ -norm (Goodfellow et al., 2015; Moosavi-Dezfooli et al., 2017)). User studies suggest these norm constraints yield imperceptible perturbations for classification tasks like speaker verification (Kreuk et al., 2018), and for weaker ASR attacks (e.g., untargeted, single-word attacks (Cisse et al., 2017; Alzantot et al., 2018)).

However, recent work indicates norm-constrained perturbations may lack capacity for stronger sentence-level targeted ASR attacks (Schönherr et al., 2018; Qin et al., 2019), prompting the use of psychoacoustic models for masking. Yet, even psychoacoustically masked perturbations can be shown perceptible in user studies, implying that difficult attacks still push perceptually solid constraints to their limits. For our attacks, we therefore enhance existing models (§4.2).

3 BACKGROUND

We study the existence of *adversarial perturbations* for DNS models. As background, we first review the standard behavior of DNS models and important characteristics of the models we study. We then formally state the definition of adversarial perturbations for DNS models, and review auditory masking, a perceptibility constraint for audio attacks (Schönherr et al., 2018; Qin et al., 2019).

3.1 Speech Denoising

The primary goal of speech denoising models is to remove background noise from a speech signal (Defossez et al., 2020; Chen et al., 2022; Zhao et al., 2022b; Lu et al., 2023; 2024). In some cases (Defossez et al., 2020), but not all (Chen et al., 2022; Zhao et al., 2022b; Lu et al., 2023), the model is also trained to dereverberate signals that have been distorted by the environment or microphone.

Concretely, in this paper, speech denoising models are functions $f: \mathbb{R}^n \to \mathbb{R}^n$ which map to and from the space of audio waveforms. We model the denoising model's input as some clean speech $y \in \mathbb{R}^n$, mixed additively with some (possibly zero) background noise $b \in \mathbb{R}^n$, and (optionally) convolved with some room impulse response (RIR) $r \in \mathbb{R}^m$: x = r*(y+b). (Note that some work only convolves the speech with r (Zhao et al., 2022b); in this paper, we convolve both speech and background noise.) Given this input, the model produces $f(x) = \hat{y}$, where \hat{y} is as close to either y or r*y as possible. For consistency with past benchmarks (Dubey et al., 2024), and because we do not attempt to evaluate healthy model behavior, we assume that the model attempts to infer y.

3.1.1 Models

We test the robustness of four recent open-source DNS models with publicly available checkpoints: Demucs, published online as Denoiser (Defossez et al., 2020); Full-SubNet+ (Chen et al., 2022); FRCRN (Zhao et al., 2022b); and MP-SENet (Lu et al., 2023; 2024). The models we study operate in either the time domain (waveforms) or the time-frequency (TF) domain (STFT spectrograms) (Lu et al., 2023); see Table 1. Note that FSN+, FRCRN and MP-SENet do not attempt to dereverberate: they infer r * y instead of y. For exact checkpoints and detailed architectural notes, see Appendix E.

3.2 ADVERSARIAL PERTURBATIONS

Given the definition of speech denoising given in §3.1, an adversarial perturbation for some input x = r * (y + b) is any $\delta \in \mathbb{R}^n$ such that $x + \delta \mapsto y'$, where (a) y' is undesirable and (b) $x + \delta$ sounds identical to x. To make these notions concrete, we distinguish untargeted and targeted attacks.

Untargeted attacks. Here, we assume that the attacker has a loss function $\mathcal{L}: \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$, and they wish to *maximize* the value of $\mathcal{L}(y',y)$. They do so by selecting δ from a set of perturbations

which are imperceptible when added to x, a feasible set which we refer to as D(x). Thus, the adversary in an untargeted attack wishes to find $\delta^* \in \arg\max_{\delta \in D(x)} \mathcal{L}\left(f(x+\delta),y\right)$.

Targeted attacks. In this case, the attacker has a target output y' that they wish to cause f to output. We model this as the attacker wishing to *minimize* the loss $\mathcal{L}(f(x+\delta),y')$, i.e., the adversary wants $\delta^* \in \arg\min_{\delta \in D(x)} \mathcal{L}(f(x+\delta),y')$.

Over-the-air attacks. In an over-the-air attack, the adversarial perturbation is distorted by the room's acoustic characteristics and received through a microphone. As these distortions are assumed to be applied equally to all audio, this is generally simulated by applying the given RIR to the perturbation δ as well as speech y and background noise b (Qin et al., 2019; Schönherr et al., 2020). For example, the untargeted attack becomes the problem of finding $\delta^* \in \arg\max_{\delta \in D^*(x)} \mathcal{L}\left(f(r*(y+b+\delta)),y\right)$, where D^* is all perturbations that are imperceptible after all audio, including the perturbation, is convolved with r.

3.3 AUDITORY MASKING

We use a perceptibility constraint based on auditory masking, also known as psychoacoustic hiding: (Lin & Abdulla, 2015; Schönherr et al., 2018; Qin et al., 2019): by varying the loudness constraint on the perturbation according to how loud the original signal is across different time segments and frequency bands, the perturbation can be effectively hidden in the original signal. This is accomplished through the computation of masking thresholds $\theta_{\tau,\omega}$ on the power spectral density (PSD) matrix of the perturbation, such that any audio whose PSD is upper-bounded by these masking thresholds is, in principle, imperceptible by humans. Expanding on the auditory masking algorithms used in this prior work, we apply several enhancements to the threshold calculation to maximize attack power and minimize perceptibility; see §4.2.

4 METHOD

Most of the perturbations we study can be characterized as instantiations of the attacks shown in §3.2, given varying settings of f, y, b, r, \mathcal{L} , D, and y'. Here, we discuss \mathcal{L} , D, and optimization.

4.1 OPTIMIZATION OBJECTIVE

In this work, we used Short-Term Objective Intelligibility (STOI) (Taal et al., 2011) as the loss function \mathcal{L} for both targeted and untargeted attacks. STOI is open-source and differentiable, and hewed closely to our perceptions of intelligibility in untargeted experiments.

While mean-squared error (MSE) is widely used in comparable image tasks and has been successfully employed in prior attacks on generative audio models (Takahashi et al., 2021; Dong et al., 2023), it remains an unreliable metric for intelligibility. It suffers from a lack of several detailed desiderata; for example, it is strongly phase-dependent, meaning that simply delaying the audio by a few milliseconds can cause very large MSE but almost no loss in intelligibility. Nevertheless, the central issue with MSE is that it is simply not designed to be an intelligibility metric; in our experiments, attacks would often achieve a large (untargeted) MSE by simply inducing the DNS model to fail to remove noise, rather than rendering the speech unintelligible.

We could not use Perceptual Evaluation of Speech Quality (PESQ) (Rix et al., 2001), a notable intelligibility metric, as its license owners did not grant us an academic license for this project. However, real adversaries would of course face no such ethical or legal concerns.

Finally, we also considered DNN-based intelligibility metrics such as DNSMOS (Reddy et al., 2022), NISQA (Mittag et al., 2021), and the word error rate (WER) of Whisper (Radford et al., 2022). While such metrics were reasonably effective as unoptimized metrics for attack success, we found empirically that they work poorly as objectives, as the attack simply learns to induce the model to output adversarial noise against the metric network.

4.2 PERCEPTIBILITY CONSTRAINT ENHANCEMENTS

Our precise procedure for computing the masking thresholds is almost identical to the MP3 psychoacoustic model as described by Lin & Abdulla (2015) and used by Schönherr et al. (2018) and Qin et al. (2019), with the following changes: (a) for simplicity, we calculate the PSD of the input audio by normalizing to [-1, 1] and converting to dB SPL explicitly, rather than setting the maximum bin to equal 96; (b) we enhance the psychoacoustic model with temporal pre- and post-masking, as described by Lin & Abdulla (2015) (see Appendix A for details); (c) because prior studies have found even attacks generated according to these masking thresholds to still be detectable by humans (Schönherr et al., 2018; Qin et al., 2019), we restrict our attacks further by shifting all masking thresholds down by 12 dB. We combine thresholds from contemporaneous masking and pre- and post-masking by taking the maximum between each. See Appendices A and B for details.

4.3 OPTIMIZATION ALGORITHM

We use projected gradient descent (PGD; Madry et al. 2018) to find the perturbation δ^* :

$$\delta_{t+1} = \Pi_{D(x)} \left(\delta_t + \alpha \nabla_{\delta_t} \mathcal{L} \left(x + \delta_t \right) \right), \tag{1}$$

where Π is the projection operator. Fortunately, the projection step is straightforward when using auditory masking on unreverberated perturbations (see §4.4 for details on reverberated perturbation optimization). If $\tilde{\delta}_{\tau,\omega}$ is the STFT spectrogram of δ at frame τ and frequency bin ω , then we clip the magnitude of $\tilde{\delta}_{\tau,\omega}$ such that $\mathrm{PSD}(\delta)_{\tau,\omega} \leq \theta_{\tau,\omega}$ while preserving its phase. See Appendix B for details. We also tested Lagrangian dual descent, a variant of which was used by Qin et al. (2019), and the fast gradient sign method (Goodfellow et al., 2015), but these did not perform as well.

4.4 ADDITIONAL EXPERIMENTAL METHODS

Optimization in Simulated Over-the-air Attacks. We simulate over-the-air attacks in a specific acoustic environment by applying the given RIR to the adversarial perturbation as well as to the clean speech and background noise (see §3.2). Doing so dramatically increases the difficulty of the optimization problem, as the projection step in (1) no longer has a closed-form solution: rather than directly clipping δ , we must instead solve for some δ such that $PSD(r * \delta)_{\tau,\omega} \leq \theta_{\tau,\omega}$.

We explore three methods to find such a δ : Wiener deconvolution, manual projection through gradient descent, and a combination of both. See Appendix C for details.

Targeted Attacks. We test three types of targets for targeted attacks: speech samples from the same speaker, speech samples from different speakers, and artificial targets generated using voice cloning systems such as MaskGCT Wang et al. (2024a). For all target types, our attack objective is the STOI of the model's output versus the clean speech minus the STOI of the model's output versus the target speech, maximizing the relative intelligibility of the target speech.

Defenses. Prior work has investigated a wide range of defenses against adversarial audio perturbations, from pre-processing steps like randomized smoothing (Subramanian et al., 2019; Olivier & Raj, 2021), purification (Wu et al., 2023), and adversarial example detection (Yang et al., 2019; Hussain et al., 2021), to more involved defenses like adversarial training (Sallo et al., 2021). In this paper, we evaluate the simplest of these defenses: Gaussian noise applied to the attacked audio, normalized to various SNRs.

Universal Adversarial Perturbations. While most of our experiments focus on attacking a single known input, we also explore the existence of universal adversarial perturbations (UAPs): individual perturbations that lead one or more models to misbehave on multiple inputs (Moosavi-Dezfooli et al., 2017; Abdoli et al., 2019; Neekhara et al., 2019; Xie et al., 2021; Zhang et al., 2021; Sun et al., 2024). We design UAPs for DNS models under the assumption that the background noise b and RIR r remain the same for varying clean speech samples y_n , and calculate $D_U = \bigcap_n D(r*(y_n+b))$ (i.e., take the minimum of the different masking thresholds). To train the attack, we iteratively execute one step of PGD sequentially on each input $r*(y_n+b)$.

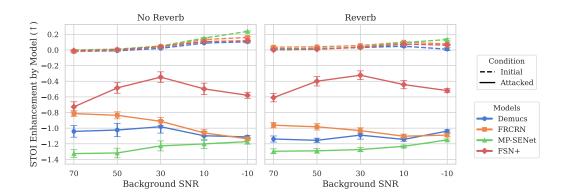


Figure 1: Untargeted intelligibility degradation. Dashed curves show the DNS models' baseline STOI improvement $\Delta_{STOI} = STOI({\rm clean}, {\rm output}) - STOI({\rm clean}, {\rm input});$ solid curves show the same quantity after adding an imperceptible perturbation. Values flip from >0 to <0 upon addition of adversarial noise across environmental settings, meaning the attack pushes speech from "cleaner than input" to "less intelligible than the noisy input itself." Error bars: \pm s.e. over 20 seeds.

5 EXPERIMENTS

In our computational experiments, we focused on answering the following questions: **a**) Are DNS models susceptible to adversarial noise, or do they denoise that successfully as well? **b**) In which use cases (amount of noise, reverb, etc.) are the models we study vulnerable to adversarial attacks? **c**) How does adversarial robustness vary by model, and which model characteristics seem to correlate with robustness? **d**) Can attackers induce the model to output a target utterance, as in the case of ASR transcriptions, or are they limited to untargeted attacks? **e**) Do imperceptible universal adversarial perturbations exist, or does the attacker need to know what the speaker will say? **f**) Do attackers need to have access to model gradients in order to successfully attack these models? **g**) Are sophisticated defenses (e.g., adversarial training) necessary, or is Gaussian perturbation good enough? **h**) Are real, over-the-air attacks possible in principle, or is this setting too difficult? **i**) How does our perceptibility constraint trade off against attack success as we tighten it (Appendix D.3)?

5.1 EXPERIMENTAL DETAILS

Datasets. All audio samples – speech, noise, and RIRs – were taken from the main track of the ICASSP 2022 DNS Challenge 4 (Dubey et al., 2022). Ten-second clean speech samples were selected randomly from English read speech (as noted by Dubey et al. (2022), collated from LibriVox.org) and VCTK Corpus (Yamagishi et al., 2019). Audio was clipped to five seconds for MP-SENet due to insufficient VRAM. We filtered speech to contain at least 15 words according to Whisper. See Appendix E for more details.

Metrics. In addition to STOI, we evaluated attacked audio and the model outputs using ViSQOL (Hines et al., 2015), an intrusive (binary with ground-truth) and unlearned metric; two non-intrusive deep metrics, NISQA (Mittag et al., 2021) and DNSMOS (Reddy et al., 2022); and word accuracy (de Oliveira et al., 2023) using the ASR model Whisper (Radford et al., 2022). See Figure 2.

Attacks. All combinations of environment setting (background SNR and reverb/no reverb) and model choice were run for 20 shared seeds. The attack lasted for a different number of iterations for each model, chosen to ensure that the entire attack (including metric computation) lasted for about one hour on an Nvidia L40S GPU. In particular, this allowed 20,000 iterations for Demucs and FSN+, 10,000 iterations for MP-SENet, and 5,000 iterations for FRCRN.

5.2 RESULTS AND DISCUSSION

Figure 1 demonstrates our fundamental result: all DNS models we tested could be induced to output far worse-quality audio than the input through the addition of imperceptible adversarial noise, and indeed, all models could be reliably (or often, in the case of FSN+) reduced to outputting gibberish.

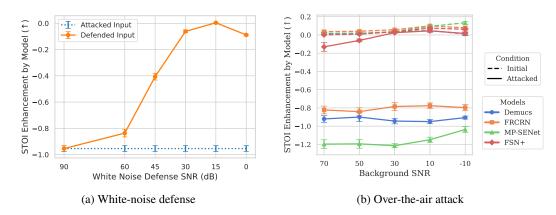


Figure 3: **Real-world implications: defenses and over-the-air attacks.** (a) Adding moderate Gaussian ("white") noise to the attacked input partially smooths out adversarial perturbations, raising STOI but not restoring clean performance. (b) When perturbations are trained and tested under a simulated room impulse response, they still cripple every model except FSN+. Curves show $\Delta_{\text{STOI}} = \text{STOI}(\text{clean}, \text{output}) - \text{STOI}(\text{clean}, \text{input})$; error bars: \pm s.e. over 20 seeds.

Furthermore, the success of the attack was relatively invariant with respect to setting: all models could be successfully attacked in all settings, including a setting which has almost no attack vector: 70 dB SNR noise with no reverb. Thus, we can answer question **b**): three out of four tested models were completely susceptible in all settings, while FSN+ was fairly susceptible in most settings.

5.2.1 Model Results

Figures 1 and 5 suggest that FSN+ is by far the most resilient, FRCRN and Demucs have comparable resilience, and MP-SENet is slightly more susceptible. Surprisingly, though, the resilience of FSN+ is not due to high-level architecture differences such as domain or parameter count (see Appendix D.1). Instead, the large difference between FSN+'s susceptibility and others was due to the pseudo-robustness of exploding gradients (Athalye et al., 2018): during the course of the attack on FSN+, the gradient of the STOI loss with respect to the adversarial waveform would often grow to have a norm of 10^{30} or greater, causing numerical instability even with gradient clipping. Therefore, FSN+ might be more vulnerable to black-box attacks (Athalye et al., 2018).

5.2.2 SIMULATED OVER-THE-AIR ATTACKS

While the pseudo-protection of FSN+'s exploding gradients was amplified in our simulated overthe-air experiments, we discovered that all other models are highly vulnerable to untargeted reverberated perturbations, despite the additional optimization challenges in developing them (see Figure 3). However, the increased challenge of this attack did require that we slightly loosen our per-

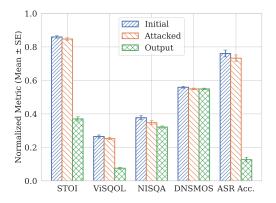


Figure 2: Normalized values of **the five speech intelligibility and quality metrics we test**, averaged across all models and settings of Figure 1 and 20 seeds, with s.e. error bars. "Output" refers to model output given the attacked input. ASR accuracy is computed as $1-\min(WER, 1)$. Ranges used for normalization were: STOI: [-1, 1]. ViSQOL: [1, 5]. NISQA: [0, 5]. DNSMOS: [0, 5]. ASR accuracy: [0, 1]. ASR ground-truth is Whisper applied to clean speech.

ceptibility constraint, reducing the masking thresholds by only 6 dB rather than 12 dB. While this loosening still implies a tighter constraint than prior work, it did cause slight crackling to be audible in the speech, though it would be difficult to distinguish from minor distortion in real settings.

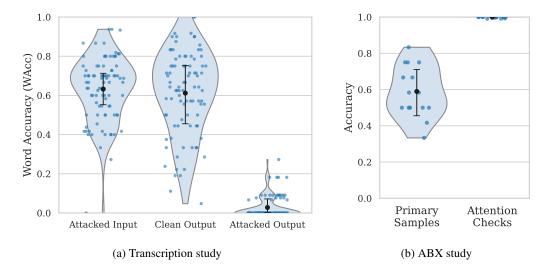


Figure 4: **Human study**. We ran transcription and ABX studies with 15 audio/multimedia researchers using high-quality audio setups. The transcription task consisted of 18 random 5-second samples (sampled once for all participants), including 6 attacked inputs, 6 model outputs given clean inputs, and 6 model outputs given attacked inputs. Word accuracy (WAcc) results show that the attacked input and clean output were reasonably intelligible, while the attacked output was unintelligible. The ABX task consisted of 12 pairs of 5-second samples (also common across participants), each consisting of a clean sample and its attacked counterpart, as well as two attention checks each consisting of a pair of clean and aggressively attacked audio. For both plots, we show 95% two-way bootstrap CIs (Owen, 2007); for the ABX study, we plot participant-wise marginal accuracies (i.e., averaged across samples). ABX results show that the average participant's accuracy is insignificantly above the random chance baseline of 50%, with substantial variance.

5.2.3 MIXED AND NEGATIVE RESULTS: TARGETED ATTACKS, UAPS, TRANSFER

We found that the answers to questions \mathbf{d}), \mathbf{e}), and \mathbf{f}) are currently negative for the attacker; see Appendix D.2 for full details.

5.2.4 Defenses

We discovered that simple Gaussian perturbation offers reasonable protection against adversarial perturbations, though only when applied at a low enough SNR to damage model performance (Figure 3; note that a STOI enhancement of 0 is slightly lower than the unattacked average of ~ 0.044). However, we do not assume adaptation by the attacker; knowledge of this defense would likely allow the attacker to train more noise-resistant perturbations (see, e.g., Tramer et al. (2020) for an example of such adaptive attacks in image classifiers), emphasizing a need for more sophisticated defenses.

6 Human Study

Participants & ethics. We recruited **15** adult audio/multimedia researchers from an industrial research lab. Participation was voluntary and conducted during paid working hours; no additional compensation was provided. Under the host institution's policy for minimal-risk studies, our study did not require formal IRB/ethics review. All participants provided informed consent and could withdraw at any time, and no personally identifying information or audio recordings were collected; we stored only anonymized task responses and timestamps.

Tasks and protocol. Each participant completed two tasks in a fixed order: (1) *transcription* and (2) *ABX discrimination*. Participants could take as long as they wished, though tasks were designed to be completed in less than 25 minutes. Once a participant advanced from transcription to ABX, they could not return to transcription. Participants typically use high-quality headphones in their day-to-

day work; we recommended (but did not enforce) headphone use. Participants were instructed to leave transcription boxes blank if no intelligible speech was detected.

Stimuli and conditions. All stimuli were 5 s clips cropped from the same datasets as in the computational study (pilot feedback indicated 10 s was too long to transcribe or compare reliably). We filtered candidate clips to contain at least 10 words and to use only the 2,000 most common English words to reduce ESL burden, again thanks to pilot feedback. We used two background conditions designed to mimic common listening scenarios: 30 dB and 50 dB SNR with reverberation. We evaluated Denoiser, FRCRN, and MP-SENet in the human study in the human study; Full-SubNet+ was excluded here because our computational evaluation did not yield successful attacks for it.

Attacks. Attacks were conducted as in the computational study, here using the tight psychoacoustic masking constraint (App. E).

Measures and analysis. For transcription we report word accuracy (WAcc; 1 - WER), calculated without any equivalencies besides uppercase—lowercase. For ABX we report accuracy versus the 50% chance baseline. For confidence intervals, we use 95% two-way bootstrap intervals (a.k.a. pigeonhole bootstrap intervals; Owen, 2007), a model-free alternative to generalized linear mixed models (GLMMs) for data with crossed effects. Note that, because samples were shared between participants, naive bootstrap and the Student's *t*-test would not be appropriate statistical tests, as they would underestimate variance.

Results. Transcription WAcc (Figure 4a) shows that the **Attacked Output** had near-zero intelligibility, while both the **Attacked Input** and **Clean Output** were reasonably intelligible. In particular, an intersection-union test indicates that the attacked output was significantly less intelligible than both the attacked input and clean output (95% upper bounds on the mean differences of -0.464 and -0.458, respectively; see Appendix F for details).

ABX accuracy (Figure 4b) had a mean of 59%, marginally above chance, but with substantial variance between participants and samples. The mean was not significantly different from 50% according to a one-sided 95% pigeonhole bootstrap test (lower bound: 0.478). While other less conservative statistical tests (e.g., a Wald test with a GLMM) might conclude significance, the present results offer preliminary support for the subjective imperceptibility of our attacks.

7 CONCLUSION

In this paper, we show for the first time that DNS models are susceptible to imperceptible adversarial attacks in a wide variety of simulated settings: through the addition of adversarial noise hidden in the original input signal, all of the models we study can be induced to output audio with almost no resemblance to either clean or noisy speech. Our attacks generalize to a wide variety of settings, including simulated over-the-air attacks and attacks with no background noise or reverb, and suggest that subjectively effective targeted attacks may be enabled by better optimization targets (Appendix D.2). Finally, our human study validates that our attacks are generally imperceptible, and the outputs generally unintelligible, even to audio experts.

Our results have several important limitations and caveats. First, the DNS model Full-SubNet+ (FSN+) is partially protected from gradient-based attacks due to its exploding gradients, though this pseudo-defense is generally easily overcome by determined adversaries (Athalye et al., 2018). Second, our attacks all rely on gradient access for now, as we showed model transfer to be ineffective for the adversarial perturbations that we found. Third, we only attack fully differentiable DNS models; token-based DNS models (Wang et al., 2024b) will require new techniques to attack. Finally, our attacks require that the target utterance is known in advance.

Nevertheless, we hope to convince the research community that open-source DNS models are both an appealing and feasible target for attack, with the potential to incapacitate live audio streams, speech recognition systems, hearing aid users, and more. With simple defenses such as white noise only offering limited protection, we urge researchers to evaluate further models in more lifelike settings and develop and apply better defenses before attackers exploit these critical weaknesses.

ETHICS STATEMENT

Dual-use adversarial robustness research. This paper demonstrates research into the adversarial robustness of DNS models, including methodological advancements in adversarial perturbation generation. However, we believe the security utility of our present work outweighs any utility it could provide to real-world adversaries. First, the focus of the study is on precisely characterizing long-standing vulnerabilities so as to spur future research rather than providing methods for overcoming existing defenses. Second, we indeed document that there are several further steps required for an attacker before successfully applying the present methods to a real-world scenario, including extending the simulated over-the-air attacks shown here to real over-the-air scenarios. Finally, we show that simple, easily implemented defenses such as Gaussian perturbation are effective against naive attacks, thus leaving the immediate advantage with a well-informed defender.

Human subjects. The present paper also describes a small human study with 15 industry researchers. Participation was purely voluntary during business hours, with compensation only in the form of their existing salary; explicit consent was collected before each participant began the study; the study consisted purely of audio tasks containing no offensive speech; each participant was free to withdraw at any time without penalty; and no personal information was collected or stored at any time. While the host institution of the study does not require IRB approval for low-risk studies, we believe the risks to subjects posed by our study design are minimal.

Licensing and credit. We carefully track the licenses of all artifacts we use other than standard PyPI libraries (Appendix G). In particular, we did not use the proprietary speech quality metric PESQ, as it is not licensed for research use by default and the maintainers did not respond to our request.

LLM DISCLOSURE

LLMs provided assistance throughout this project except for ideation of the central research questions. GPT-5 and OpenAI o3 were used for methodological ideation later in the project (limited to the over-the-air attack method and the human study); GPT-5, o3, and Gemini 2.5 Pro were used in writing (drafting and editing) at all stages; and GPT-5, o3, and OpenAI Codex were used as coding partners at all stages. The authors attest to the correctness of all present claims and results.

REPRODUCIBILITY STATEMENT

We have specified all details needed to reproduce our experiments to the best of our ability both in the main body and in all appendices. We will share anonymous code during the discussion period, and will release our code publicly upon acceptance.

REFERENCES

Sajjad Abdoli, Luiz G. Hafemann, Jerome Rony, Ismail Ben Ayed, Patrick Cardinal, and Alessandro L. Koerich. Universal adversarial audio perturbations, 2019. URL https://arxiv.org/abs/1908.03173.

Moustafa Alzantot, Bharathan Balaji, and Mani Srivastava. Did you hear that? adversarial examples against automatic speech recognition, 2018.

Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 274–283. PMLR, 10–15 Jul 2018. URL https://proceedings.mlr.press/v80/athalye18a.html.

Nicholas Carlini and David Wagner. Audio adversarial examples: Targeted attacks on speech-to-text. In 2018 IEEE Security and Privacy Workshops (SPW), pp. 1–7, 2018. doi: 10.1109/SPW.2018.00009.

Guangke Chen, Sen Chenb, Lingling Fan, Xiaoning Du, Zhe Zhao, Fu Song, and Yang Liu. Who is real bob? adversarial attacks on speaker recognition systems. In 2021 IEEE Symposium on Security and Privacy (SP), pp. 694–711, 2021. doi: 10.1109/SP40001.2021.00004.

- Guangke Chen, Zhe Zhao, Fu Song, Sen Chen, Lingling Fan, Feng Wang, and Jiashui Wang. Towards understanding and mitigating audio adversarial examples for speaker recognition. *IEEE Trans. Dependable Secur. Comput.*, 20(5):3970–3987, September 2023. ISSN 1545-5971. doi: 10.1109/TDSC.2022.3220673. URL https://doi.org/10.1109/TDSC.2022.3220673.
- Jun Chen, Zilin Wang, Deyi Tuo, Zhiyong Wu, Shiyin Kang, and Helen Meng. Fullsubnet+: Channel attention fullsubnet with complex spectrograms for speech enhancement. In *ICASSP* 2022 2022 *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7857–7861, 2022. doi: 10.1109/ICASSP43922.2022.9747888.
- Moustapha Cisse, Yossi Adi, Natalia Neverova, and Joseph Keshet. Houdini: fooling deep structured visual and speech recognition models with adversarial examples. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, pp. 6980–6990, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
- Ross Cutler. Developing machine learning based speech enhancement models for teams and skype. Conference presentation video, RTC @Scale 2022 (virtual), February 2022. URL https://atscaleconference.com/videos/developing-machine-learning-based-speech-enhancement-models-for-teams-and-skype-ross-cutler/. Accessed: 2025-01-28.
- Danilo de Oliveira, Julius Richter, Jean-Marie Lemercier, Tal Peer, and Timo Gerkmann. On the behavior of intrusive and non-intrusive speech enhancement metrics in predictive and generative settings, 2023. URL https://arxiv.org/abs/2306.03014.
- Alexandre Defossez, Gabriel Synnaeve, and Yossi Adi. Real time speech enhancement in the waveform domain, 2020. URL https://arxiv.org/abs/2006.12847.
- Mingyu Dong, Diqun Yan, and Rangding Wang. Adversarial examples protect your privacy on speech enhancement system. *Computer Systems Science and Engineering*, 46(1):1–12, 2023. doi: 10.32604/csse.2023.034568. URL http://www.techscience.com/csse/v46n1/51333.
- Harishchandra Dubey, Vishak Gopal, Ross Cutler, Ashkan Aazami, Sergiy Matusevych, Sebastian Braun, Sefik Emre Eskimez, Manthan Thakker, Takuya Yoshioka, Hannes Gamper, and Robert Aichner. Icassp 2022 deep noise suppression challenge. In *ICASSP 2022 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 9271–9275, 2022. doi: 10.1109/ICASSP43922.2022.9747230.
- Harishchandra Dubey, Ashkan Aazami, Vishak Gopal, Babak Naderi, Sebastian Braun, Ross Cutler, Alex Ju, Mehdi Zohourian, Min Tang, Mehrsa Golestaneh, and Robert Aichner. Icassp 2023 deep noise suppression challenge. *IEEE Open Journal of Signal Processing*, 5:725–737, 2024. doi: 10.1109/OJSP.2024.3378602.
- Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. Analysis of classifiers' robustness to adversarial perturbations. *Mach. Learn.*, 107(3):481–508, March 2018. ISSN 0885-6125. doi: 10.1007/s109 94-017-5663-3. URL https://doi.org/10.1007/s10994-017-5663-3.
- Yuan Gong and Christian Poellabauer. Crafting adversarial examples for speech paralinguistics applications, 2017. URL https://arxiv.org/abs/1711.03280.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples, 2015. URL https://arxiv.org/abs/1412.6572.
- Feng Guo, Zheng Sun, Yuxuan Chen, and Lei Ju. Towards the transferable audio adversarial attack via ensemble methods. *Cybersecurity*, 6(44), 2023. doi: 10.1186/s42400-023-00175-8. URL https://cybersecurity.springeropen.com/articles/10.1186/s42400-023-00175-8.

Andrew Hines, Jan Skoglund, Anil C. Kokaram, and Naomi Harte. ViSQOL: an objective speech quality model. *EURASIP Journal on Audio, Speech, and Music Processing*, 2015(1):13, 2015. ISSN 1687-4722. doi: 10.1186/s13636-015-0054-9. URL https://doi.org/10.1186/s13636-015-0054-9.

- Shehzeen Hussain, Paarth Neekhara, Shlomo Dubnov, Julian McAuley, and Farinaz Koushanfar. WaveGuard: Understanding and mitigating audio adversarial examples. In 30th USENIX Security Symposium (USENIX Security 21), pp. 2273–2290. USENIX Association, August 2021. ISBN 978-1-939133-24-3. URL https://www.usenix.org/conference/usenixsecurity21/presentation/hussain.
- Arindam Jati, Chin-Cheng Hsu, Monisankha Pal, Raghuveer Peri, Wael AbdAlmageed, and Shrikanth Narayanan. Adversarial attack and defense strategies for deep speaker recognition systems. *Computer Speech & Language*, 68:101199, 2021. ISSN 0885-2308. doi: https://doi.org/10.1016/j.csl.2021.101199. URL https://www.sciencedirect.com/science/article/pii/S0885230821000061.
- Weifei Jin, Yuxin Cao, Junjie Su, Qi Shen, Kai Ye, Derui Wang, Jie Hao, and Ziyao Liu. Towards evaluating the robustness of automatic speech recognition systems via audio style transfer. In *Proceedings of the 2nd ACM Workshop on Secure and Trustworthy Deep Learning Systems*, SecTL '24, pp. 47–55, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400706912. doi: 10.1145/3665451.3665532. URL https://doi.org/10.1145/3665451.3665532.
- Felix Kreuk, Yossi Adi, Moustapha Cisse, and Joseph Keshet. Fooling end-to-end speaker verification with adversarial examples. In 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 1962–1966, 2018. doi: 10.1109/ICASSP.2018.8462693.
- Xu Li, Jinghua Zhong, Xixin Wu, Jianwei Yu, Xunying Liu, and Helen Meng. Adversarial attacks on gmm i-vector based speaker verification systems. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6579–6583, 2020. doi: 10.1109/ICASSP40776.2020.9053076.
- Yiqing Lin and Waleed H. Abdulla. *Principles of Psychoacoustics*, pp. 15–49. Springer International Publishing, Cham, 2015. ISBN 978-3-319-07974-5. doi: 10.1007/978-3-319-07974-5_2. URL https://doi.org/10.1007/978-3-319-07974-5_2.
- Ye-Xin Lu, Yang Ai, and Zhen-Hua Ling. Mp-senet: A speech enhancement model with parallel denoising of magnitude and phase spectra. In *Interspeech 2023*, pp. 3834–3838, 2023. doi: 10.21437/Interspeech.2023-1441.
- Ye-Xin Lu, Yang Ai, and Zhen-Hua Ling. Explicit estimation of magnitude and phase spectra in parallel for high-quality speech enhancement, 2024. URL https://arxiv.org/abs/2308.08926.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=rJzIBf ZAb.
- Manuel Milling, Shuo Liu, Andreas Triantafyllopoulos, Ilhan Aslan, and Björn W. Schuller. Audio enhancement for computer audition—an iterative training paradigm using sample importance. *Journal of Computer Science and Technology*, 39(4):895–911, 7 2024. ISSN 1860-4749. doi: 10.1007/s11390-024-2934-x. URL https://doi.org/10.1007/s11390-024-2934-x.
- Gabriel Mittag, Babak Naderi, Assmaa Chehadi, and Sebastian Möller. Nisqa: A deep cnn-self-attention model for multidimensional speech quality prediction with crowdsourced datasets. In *Interspeech 2021*, pp. 2127–2131, 2021. doi: 10.21437/Interspeech.2021-299.
- Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. Universal adversarial perturbations. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 86–94, 2017. doi: 10.1109/CVPR.2017.17.

Paarth Neekhara, Shehzeen Hussain, Prakhar Pandey, Shlomo Dubnov, Julian McAuley, and Farinaz Koushanfar. Universal adversarial perturbations for speech recognition systems. In *Interspeech* 2019, pp. 481–485, 2019. doi: 10.21437/Interspeech.2019-1353.

- Stavros Ntalampiras. Adversarial attacks against audio surveillance systems. In 2022 30th European Signal Processing Conference (EUSIPCO), pp. 284–288, 2022. doi: 10.23919/EUSIPCO55093. 2022.9909635.
- Raphael Olivier and Bhiksha Raj. Sequential randomized smoothing for adversarially robust speech recognition. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 6372–6386, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.514. URL https://aclanthology.org/2021.emnlp-main.514/.
- Art B. Owen. The pigeonhole bootstrap. *The Annals of Applied Statistics*, 1(2), December 2007. ISSN 1932-6157. doi: 10.1214/07-aoas122. URL http://dx.doi.org/10.1214/07-AOAS122.
- Yao Qin, Nicholas Carlini, Garrison W. Cottrell, Ian J. Goodfellow, and Colin Raffel. Imperceptible, robust, and targeted adversarial examples for automatic speech recognition. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 5231–5240. PMLR, 2019. URL http://proceedings.mlr.press/v97/qin19a.html.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision. https://openai.com/index/whisper/, 2022. OpenAI Report.
- Chandan K A Reddy, Vishak Gopal, and Ross Cutler. Dnsmos p.835: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors, 2022. URL https://arxiv.org/abs/2110.01763.
- A.W. Rix, J.G. Beerends, M.P. Hollier, and A.P. Hekstra. Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs. In 2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No.01CH37221), volume 2, pp. 749–752 vol.2, 2001. doi: 10.1109/ICASSP.2001.941023.
- Raymel Alfonso Sallo, Mohammad Esmaeilpour, and Patrick Cardinal. Adversarially training for audio classifiers. In 2020 25th International Conference on Pattern Recognition (ICPR), pp. 9569–9576, 2021. doi: 10.1109/ICPR48806.2021.9413263.
- Lea Schönherr, Thorsten Eisenhofer, Steffen Zeiler, Thorsten Holz, and Dorothea Kolossa. Imperio: Robust over-the-air adversarial examples for automatic speech recognition systems. In *Proceedings of the 36th Annual Computer Security Applications Conference*, ACSAC '20, pp. 843–855, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450388580. doi: 10.1145/3427228.3427276. URL https://doi.org/10.1145/3427228.3427276.
- Lea Schönherr, Katharina Kohls, Steffen Zeiler, Thorsten Holz, and Dorothea Kolossa. Adversarial attacks against automatic speech recognition systems via psychoacoustic hiding, 2018. URL https://arxiv.org/abs/1808.05665.
- Ali Shahin Shamsabadi, Francisco Sepúlveda Teixeira, Alberto Abad, Bhiksha Raj, Andrea Cavallaro, and Isabel Trancoso. Foolhd: Fooling speaker identification by highly imperceptible adversarial disturbances. In *ICASSP 2021 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6159–6163, 2021. doi: 10.1109/ICASSP39728.2021.941 3760.
- Vinod Subramanian, Emmanouil Benetos, and Mark B. Sandler. Robustness of adversarial attacks in sound event classification. In *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2019 Workshop (DCASE2019)*, pp. 239–243, New York University, New York, NY, USA, October 2019. doi: 10.33682/sp9n-qk06.

Zheng Sun, Jinxiao Zhao, Feng Guo, Yuxuan Chen, and Lei Ju. Commanderuap: a practical and transferable universal adversarial attacks on speech recognition models. *Cybersecurity*, 7(1):38, June 2024. ISSN 2523-3246. doi: 10.1186/s42400-024-00218-8. URL https://doi.org/10.1186/s42400-024-00218-8.

- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks, 2014. URL https://arxiv.org/abs/1312.6199.
- Cees H. Taal, Richard C. Hendriks, Richard Heusdens, and Jesper Jensen. An algorithm for intelligibility prediction of time–frequency weighted noisy speech. *IEEE Transactions on Audio, Speech, and Language Processing*, 19(7):2125–2136, 2011. doi: 10.1109/TASL.2011.2114881.
- Naoya Takahashi, Shota Inoue, and Yuki Mitsufuji. Adversarial attacks on audio source separation. In *ICASSP 2021 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 521–525, 2021. doi: 10.1109/ICASSP39728.2021.9414844.
- Florian Tramer, Nicholas Carlini, Wieland Brendel, and Aleksander Madry. On adaptive attacks to adversarial example defenses. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 1633–1645. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/11f38f8ecd71867b42433548d1078e38-Paper.pdf.
- Kendrick Trinh. Adversarial attacks on speech separation systems. Master's project, San Jose State University, 2022. URL https://scholarworks.sjsu.edu/etd_projects/1089. Paper 1089.
- Qing Wang, Pengcheng Guo, and Lei Xie. Inaudible adversarial perturbations for targeted attack in speaker recognition. In *Interspeech 2020*, pp. 4228–4232, 2020. doi: 10.21437/Interspeech.202 0-1955.
- Yuancheng Wang, Haoyue Zhan, Liwei Liu, Ruihong Zeng, Haotian Guo, Jiachen Zheng, Qiang Zhang, Xueyao Zhang, Shunsi Zhang, and Zhizheng Wu. Maskgct: Zero-shot text-to-speech with masked generative codec transformer, 2024a. URL https://arxiv.org/abs/2409.00750.
- Ziqian Wang, Xinfa Zhu, Zihan Zhang, YuanJun Lv, Ning Jiang, Guoqing Zhao, and Lei Xie. Selm: Speech enhancement using discrete tokens and language models. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 11561–11565. IEEE, 2024b.
- Nils L. Westhausen, Hendrik Kayser, Theresa Jansen, and Bernd T. Meyer. Real-time multichannel deep speech enhancement in hearing aids: Comparing monaural and binaural processing in complex acoustic scenarios. *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, 32: 4596–4606, October 2024. ISSN 2329-9290. doi: 10.1109/TASLP.2024.3473315. URL https://doi.org/10.1109/TASLP.2024.3473315.
- Donald S. Williamson, Yuxuan Wang, and DeLiang Wang. Complex ratio masking for monaural speech separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 24(3): 483–492, 2016. doi: 10.1109/TASLP.2015.2512042.
- Shutong Wu, Jiongxiao Wang, Wei Ping, Weili Nie, and Chaowei Xiao. Defending against adversarial audio via diffusion model. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=5-Df3tljit7.
- Yi Xie, Zhuohang Li, Cong Shi, Jian Liu, Yingying Chen, and Bo Yuan. Enabling fast and universal audio adversarial attack using generative model. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(16):14129–14137, May 2021. doi: 10.1609/aaai.v35i16.17663. URL https://ojs.aaai.org/index.php/AAAI/article/view/17663.
- Junichi Yamagishi, Christophe Veaux, and Kirsten MacDonald. CSTR VCTK Corpus: English Multi-speaker Corpus for CSTR Voice Cloning Toolkit (version 0.92). https://doi.org/10.7488/DS/2645, 2019. Data available 2019-11-13.

Zhuolin Yang, Bo Li, Pin-Yu Chen, and Dawn Song. Characterizing audio adversarial examples using temporal dependency. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=r1q4E3C9t7.

Guoming Zhang, Chen Yan, Xiaoyu Ji, Tianchen Zhang, Taimin Zhang, and Wenyuan Xu. Dolphinattack: Inaudible voice commands. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, CCS '17, pp. 103–117, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450349468. doi: 10.1145/3133956.3134052. URL https://doi.org/10.1145/3133956.3134052.

Hanyi Zhang, Longbiao Wang, Yunchun Zhang, Meng Liu, Kong Aik Lee, and Jianguo Wei. Adversarial separation network for speaker recognition. In *Interspeech 2020*, pp. 951–955, 2020. doi: 10.21437/Interspeech.2020-1966.

Weiyi Zhang, Shuning Zhao, Le Liu, Jianmin Li, Xingliang Cheng, Thomas Fang Zheng, and Xiaolin Hu. Attack on practical speaker verification system using universal adversarial perturbations. In ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 2575–2579, 2021. doi: 10.1109/ICASSP39728.2021.9413467.

Shengkui Zhao, Bin Ma, Karn N. Watcharasupat, and Woon-Seng Gan. Frcm: Boosting feature representation using frequency recurrence for monaural speech enhancement. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 9281–9285, 2022a. doi: 10.1109/ICASSP43922.2022.9747578.

Shengkui Zhao, Bin Ma, Karn N. Watcharasupat, and Woon-Seng Gan. Frcm: Boosting feature representation using frequency recurrence for monaural speech enhancement. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 9281–9285, 2022b. doi: 10.1109/ICASSP43922.2022.9747578.

A AUDITORY MASKING

 Temporal Masking. Let $\theta_{\tau,\omega}$ be the masking threshold in frame τ and frequency bin ω , as computed by the default approach (Lin & Abdulla, 2015; Qin et al., 2019). This computation already involves consideration of the effect of sound in one frame τ and frequency bin ω on the masking thresholds in the same *frame*, τ , but adjacent *frequency bins*, ω' . For example, the model already accounts for the fact that a 1 kHz sine tone would mask a 990 Hz tone as well.

However, it is well known (Lin & Abdulla, 2015) that masking also occurs between *frames*, along the same *frequency bin*; this is known as pre-masking and post-masking. Combining results shown by Lin & Abdulla (2015) with our own experimentation, we model this by first computing contemporaneous masking thresholds as before, then for each $\theta_{\tau,\omega}$ computing the post-masking threshold it causes on later frames $\tau' > \tau$ using exponential decay of 0.02 ms⁻¹, clipping to zero at a gap of 100 ms. Pre-masking, where $\tau' < \tau$, is modeled similarly but much more sharply, with a decay of 0.16 ms⁻¹, clipping to zero after a gap of 20 ms.

Finally, each frame and frequency bin now has one contemporaneous masking threshold, a small number of masking thresholds due to pre-masking by future sounds, and a larger number of masking thresholds due to post-masking by previous sounds. To combine all of these thresholds, we set the final threshold to be their maximum, as we find this to empirically produce the best perceptibility constraint.

B PROJECTED GRADIENT DESCENT WITH MASKING

Given the TF masking thresholds $\theta_{\tau,\omega}$ calculated for the normalized input x, we seek an imperceptible perturbation δ whose reverberation-free STFT magnitude does not exceed $\theta_{\tau,\omega}$. Let

$$\widetilde{\delta}_{\tau,\omega} = \text{STFT}(\delta)_{\tau,\omega}.$$

For the purposes of illustration, let

$$PSD(\delta)_{\tau,\omega} = 10 \log_{10} (|\widetilde{\delta}_{\tau,\omega}|^2).$$

(In reality, there are subtleties around Hann window energy and length correction which we elide here; the full process is discussed by Lin & Abdulla (2015) and Qin et al. (2019).)

To enforce $PSD(\delta)_{\tau,\omega} \leq \theta_{\tau,\omega}$, we use a projection operator $\Pi_{D(x)}$ in each iteration of projected gradient descent (PGD). Specifically, after each gradient step, we adjust $\widetilde{\delta}_{\tau,\omega}$ to satisfy

$$PSD(\delta)_{\tau,\omega} \leq \theta_{\tau,\omega}$$

by scaling its magnitude if necessary:

$$\widetilde{\delta}_{\tau,\omega} \leftarrow \widetilde{\delta}_{\tau,\omega} \times \min \Bigl(1, 10^{\frac{\theta_{\tau,\omega} - \mathrm{PSD}(\delta)_{\tau,\omega}}{20}}\Bigr).$$

This enforces time-frequency masking constraints bin-by-bin while preserving the phase of $\widetilde{\delta}_{\tau,\omega}$.

Note that this is not necessarily a projection operator in the strict sense of mapping to the nearest point in D(x), particularly when δ is parameterized in the time domain, but it does ensure a feasible δ after each projection step.

C OVER-THE-AIR ATTACK OPTIMIZATION

Given knowledge of an RIR r and a clipped, imperceptible reverberated perturbation $\delta^r = r * \delta$, Wiener deconvolution is the following process: if $\tilde{\delta}^r_{\omega}$ is the full frequency spectrogram of δ^r in bin ω , and similar for \tilde{r}_{ω} , we approximate

$$\tilde{\delta}_{\omega} \approx \frac{\tilde{\delta}_{\omega}^{r} \tilde{r}_{\omega}^{*}}{|\tilde{r}_{\omega}^{*}|^{2} + \epsilon},$$

where ϵ is a small stability term (10⁻⁴) and * represents complex conjugation. Although a mathematically principled inversion, Wiener deconvolution is approximate for our use case, as convolution by RIRs is generally non-invertible.

One could attempt to apply Wiener deconvolution iteratively to find an exact δ such that $(r*\delta) \in D^*(x)$. However, a simpler iterative approach is gradient descent: we directly compute the total constraint violation $g(\delta)$ as

$$g(\delta) = \sum_{\tau,\omega} \max(\text{PSD}(r * \delta)_{\tau,\omega} - (\theta_{\tau,\omega} - d), 0),$$

where d is a small positive offset (1 dB) to ensure that gradient descent finds an exact feasible solution; we then find $\nabla g(\delta)$ using autodifferentiation and repeat until $\delta \in D^*(x)$, i.e., $g(\delta) = 0$. For over-the-air attacks, we parameterize δ in STFT spectrogram space, as this representation produces smoother gradient descent behavior. Because these attacks are harder, we gradually decrease the masking thresholds from a high starting point over time, similar to iterative constraint tightening in other works (Qin et al., 2019).

Finally, we also test applying one step of Wiener deconvolution and then finishing the projection with gradient descent.

D ADDITIONAL EXPERIMENTS

D.1 MODEL ATTACK SUCCESS VERSUS ARCHITECTURAL DETAILS

D.2 NEGATIVE RESULTS

In this section, we explore several of our research questions that presented mixed or negative results.

Universal Adversarial Perturbations (UAPs). We found that UAPs were unable to cause more than slight decreases in the quality of the model output while remaining imperceptible (question e)); this is likely due to the difficulty of hiding relevant perturbations inside of background noise alone, rather than a single clean speech sample. This result is consistent with prior work showing that audio UAPs are generally perceptible (Sun et al., 2024).

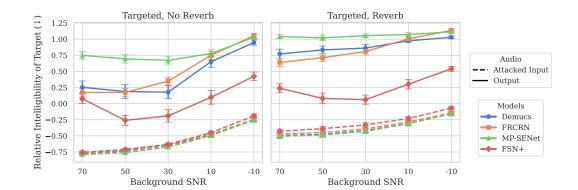


Figure 5: Intelligibility of injected target audio. We plot $\Delta_{target} = STOI(target, audio) - STOI(clean, audio)$ for the attacked input (dashed) and the DNS output (solid). Positive solid values indicate that the model's output makes a *hidden* target phrase clearer than the original clean speech, demonstrating a successful targeted attack. Targets are real utterances from the same speaker as the clean audio. Error bars: \pm s.e. over 20 seeds.

Table 1: **Effect of architecture on robustness to adversarial perturbations.** FSN+ shows substantially smaller STOI degradation than other models with similar domains or parameter counts, so coarse architectural statistics do not predict robustness.

Model	Domain	#Par. (M)	Δ STOI
Demucs	Time	33.5	-1.08
FRCRN	TF	10.3	-0.99
FSN+	TF	8.7	-0.49
MPSE	TF	2.3	-1.25

D.2.1 TARGETED ATTACKS

Figure 5 demonstrates our results for the strongest targeted attack, where targets are real samples from the same speaker. At first glance, it appears to suggest that the answer to question **d**) is resoundingly positive: DNS models are equally susceptible to targeted attacks as ASR models, and equally susceptible to adversarial perturbations in general. However, we found that, while STOI was an excellent minimization objective (low STOI implies low intelligibility), it broke down when used as a metric for maximization (high STOI does not imply high intelligibility). This issue is audible in our samples: even in attacked model outputs that are judged to have STOIs of more than 0.5 with respect to the target audio, a human listener can discern at most a faint robotic hint of the target speech. Thus, while future attackers may find ways to generate targeted attacks using more sophisticated objectives, the answer to question d) is currently both positive and negative.

D.2.2 MODEL TRANSFER

To answer question f), we evaluated the possibility of model transfer attacks, where an attack is trained for one model but applied to another. Our results indicated that attacks generally do *not* transfer between different architectures (e.g., from Demucs to FSN+); see Table 2. We also tested whether attacks between three different available Demucs checkpoints transfer to each other. Surprisingly, we found that they did not.

Our results stand in contrast to those of Dong et al. (2023), who found that perceptible adversarial examples can transfer between models. Thus, without more sophisticated transfer techniques (Guo et al., 2023), our results suggest an affirmative answer to question f), as long as the attacker is truly constrained to producing imperceptible attacks. However, further research is urgently required, particularly on pure black-box attacks (Chen et al., 2021).

Table 2: **Psychoacoustically hidden attacks do not transfer between models.** Off-diagonal entries in the attack-transfer matrix are close to 0 despite similar training data, indicating that gradient access is required for effective imperceptible attacks and that *naive* cross-architecture transfer is weak.

Trained	Evaluated on				
on	Demucs	FSN+	FRCRN	MPSE	
Demucs FSN+	-1.08 0.05	$0.04 \\ -0.49$	0.06 0.05	0.08 0.08	
FRCRN MPSE	$0.06 \\ 0.03$	$0.04 \\ 0.03$	$-0.99 \\ 0.05$	$0.08 \\ -1.25$	

D.3 ADJUSTING PERCEPTIBILITY CONSTRAINTS

Setup. In this section, we explore the effect on our objective metrics of adjusting the perceptibility constraint described in the main text. We do this primarily by adjusting the masking offset, setting it to either -8 dB (looser), -12 dB (default, same as experiments in main text), and -16 dB (tighter). For the 'tighter' setting, we also double the decay constant used in post-masking (see Appendix A) to 0.04 ms⁻¹. To simplify presentation, we show results for each constraint setting averaged across the four models. All other experimental details are the same as in the main experiments.

Note that, as in the experiment shown in the main text, we adjust the mask offsets for the simulated OTA up by 4 dB relative to the other experiments. Furthermore, one simulated over-the-air attack in Figure 10 failed to converge in a reasonable amount of time (3 hours on an L40S); we therefore set its final STOI enhancement value to be that of the initial unperturbed input.

Results. We include samples from an attack on a moderately noisy, no-reverb speech sample using each of the perceptibility constraints described here at https://sites.google.com/view/adv-dns/perceptibility. Subjective listening evaluation suggests that the looser constraint results in audible crackling in the sample, which goes away completely as the constraints are tightened.

Overall, our objective results suggest that the specific setting of the perceptibility constraint matters much more in relatively low-degradation environments (no reverb and moderate-to-low SNR); see Figures 6 and 7. This is unsurprising, as high-degradation environments naturally offer more space to hide the adversarial perturbation. When averaged across environments, though, the setting of the perceptibility constraint makes a significant difference in success, though all constraints we test allow substantial success (Figure 8), as measured by metrics other than NISQA and DNSMOS. Interestingly, most metrics (except NISQA) suggest that there is no statistically significant difference in audio quality between attacked inputs under the various constraints, implying that even the loosest constraint damages audio only very slightly.

Unsurprisingly, the white noise defense is proportionally roughly equally effective against attacks under all perceptibility constraints (Figure 9): as implied by Figure 8, even the loosest constraint still only allows very slight perturbations, which are therefore largely smoothed out by 15 dB white noise. (0 dB white noise smooths out the signal as well, hence its worse performance.)

Finally, we found that the perceptibility constraint had less effect on simulated over-the-air attack success (Figure 10) than on our directly applied attacks (Figure 6), mostly due to loosening of the constraint not seeming to allow for significantly stronger attacks. We are not certain why this is the case; while this phenomenon may be specific to the particular OTA attack optimization algorithm we use, it may also indicate more generally that convolution with an RIR – which can indeed be considered a form of smoothing – may represent a defense against adversarial perturbations, including even those which are designed adaptively against that RIR.

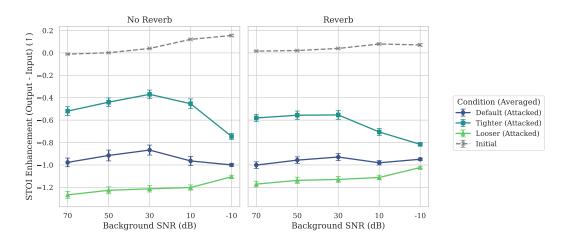


Figure 6: Untargeted intelligibility degradation across perceptibility constraints. The dashed curve shows the DNS models' baseline STOI improvement $\Delta_{STOI} = STOI(clean, output)$ – STOI(clean, input); solid curves show the same quantity after adding an imperceptible perturbation. Values flip from > 0 to < 0 upon addition of adversarial noise across environmental settings, meaning the attack pushes speech from "cleaner than input" to "less intelligible than the noisy input itself." Values are averaged across all four models. Error bars: \pm s.e. over 20 seeds.

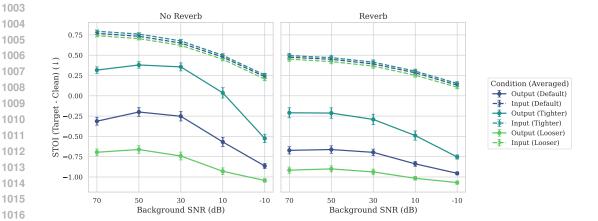


Figure 7: Intelligibility of injected target audio across perceptibility constraints. We plot $\Delta_{\text{target}} = \text{STOI}(\text{target}, \text{audio}) - \text{STOI}(\text{clean}, \text{audio})$ for the attacked input (dashed) and the DNS output (solid). Positive solid values indicate that the model's output makes an hidden target phrase clearer than the original clean speech, demonstrating a successful targeted attack. Targets are real utterances from the same speaker as the clean audio. Values are averaged across all four models. Error bars: \pm s.e. over 20 seeds.

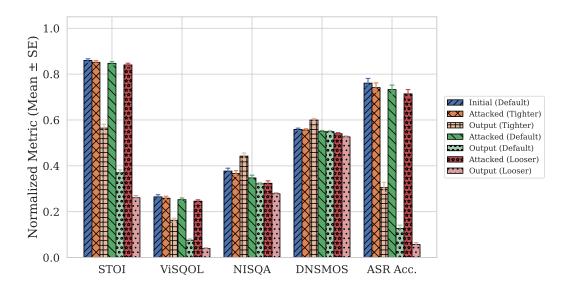


Figure 8: Additional metrics across perceptibility constraints. Normalized values of various speech intelligibility and quality metrics, averaged across all models and environmental settings and 20 seeds, with s.e. error bars. "Output" refers to model output given the attacked input. ASR accuracy is computed as $1 - \min(WER, 1)$. Ranges used for normalization were: STOI: [-1, 1]. ViSQOL: [1, 5]. NISQA: [0, 5]. DNSMOS: [0, 5]. ASR accuracy: [0, 1]. ASR ground truth is Whisper applied to clean speech.

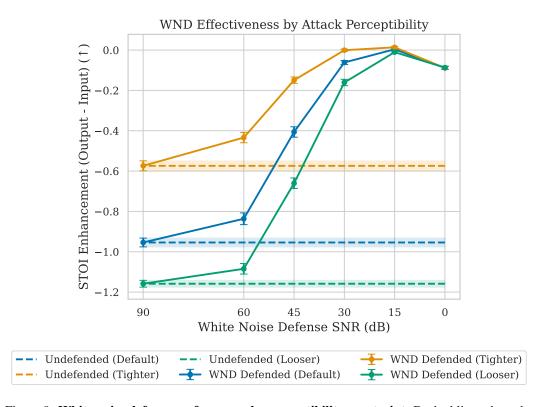


Figure 9: White noise defense performance by perceptibility constraint. Dashed lines show the DNS models' attacked STOI improvement $\Delta_{STOI} = STOI(clean, output) - STOI(clean, input)$; solid curves show the same quantity when the attacked audio is subjected to Gaussian perturbation, i.e., addition of white noise of varying signal-to-noise ratios (SNRs).

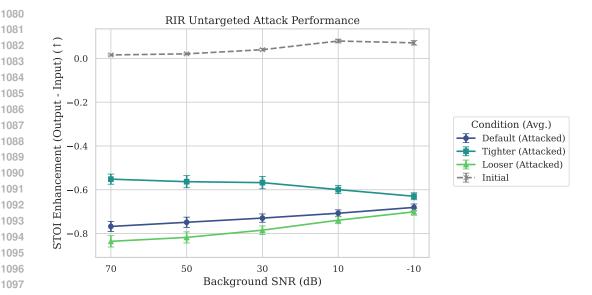


Figure 10: Simulated over-the-air attack success across perceptibility constraints. The dashed curve shows the DNS models' baseline STOI improvement $\Delta_{STOI} = STOI(clean, output)$ — STOI(clean, input); solid curves show the same quantity after adding an imperceptible perturbation, passed through the same room impulse response (RIR) as the original audio itself. Values flip from > 0 to < 0 upon addition of adversarial noise across environmental settings, meaning the attack pushes speech from "cleaner than input" to "less intelligible than the noisy input itself." Values are averaged across all four models. Error bars: \pm s.e. over 20 seeds. Mask offsets are each 4 dB higher (looser) than for other experiments.

ADDITIONAL EXPERIMENTAL DETAILS Ε

E.1 MODEL DETAILS

1082

1094

1099

1100

1101

1102

1103

1104

1105

1106 1107 1108

1109 1110 1111

1112 1113

1114

1115

1116

1117

1118 1119

1120

1121

1122

1123

1124

1125

1126

1127

1128

1129 1130

1131

1132

1133

Demucs. The earliest of the models we study, Demucs - specifically, the master 64 Demucs checkpoint provided in the Denoiser library - is a time-domain denoising model with 33.5M (trainable) parameters. Uniquely among the models we study, Demucs operates end-to-end on waveforms, rather than spectrograms: it processes the input waveform using several convolutional layers, then encodes temporal dependencies with an LSTM, before decoding again with convolutional layers. Unlike the other models we study, Demucs is designed to dereverberate as well as denoise audio.

Full-SubNet+. Full-SubNet+ (FSN+) is a TF-domain denoising model with 8.67M parameters. Full-SubNet+ takes magnitude, real, and imaginary spectra as input, and passes them through a variety of modules, including attention, convolution, and LSTMs. Like all TF-domain models we study, Full-SubNet+ outputs a ratio mask, a complex-valued spectrogram q such that \hat{y} equals the element-wise complex product of x and q (Williamson et al., 2016).

FRCRN. FRCRN is a TF-domain model with 10.3M parameters. Similar to FSN+, FRCRN uses a mixture of convolutional, attention and recurrent structures in its architecture.

MP-SENet. MP-SENet (MPSE) – specifically, the DNS checkpoint provided in the official GitHub repository - is a TF-domain model with only 2.26M parameters. Similar to FSN+ and FRCRN, MP-SENet also uses convolution, attention and recurrence for temporal modeling.

Note that, as mentioned in the main paper, MP-SENet used an unusually large amount of memory in our experiments, forcing us to curtail its input audio to 5 seconds. A note was later added to the MP-SENet README explaining that this is due to a bug in the behavior of the model with default arguments, and can be solved with alternative arguments; unfortunately, this note was added after we ran our experiments, so our experiments still curtail MP-SENet's input audio.

E.2 OTHER EXPERIMENTAL DETAILS

Dataset. Noise was selected randomly from the dataset and repeated up to ten seconds. All audio was sampled single-channel at 16 bits and 16 kHz. All settings had five sources of background noise added to them, with their post-RIR sum being set to various SNRs with respect to the post-RIR clean speech. (While five is a large number of sources, note that 70 dB SNR still represents an effectively zero-noise setting.)

Optimal perturbation selection. Within each trial, the final output was selected as the perturbation with the lowest final loss – i.e., STOI(output, clean) for untargeted attacks, and STOI(output, clean) - STOI(output, target) for targeted attacks. For over-the-air attacks where not every attack iteration converged to the feasible set, we instead selected the final output, as earlier outputs might not satisfy the final masking threshold (see Appendix C).

STFT. We use short-time Fourier transforms with Hann windows, 512 FFT points, a window length of 512, and a hop length of 256.

Optimization Details. In all experiments except the simulated over-the-air experiment, we parameterize δ in the time domain. We use Adam for all optimization, using an initial learning rate of 0.01, and clip gradients to a 2-norm of 10. We decrease the learning rate by a factor of 0.99 whenever the loss fails to decrease for 10 iterations in a row.

Hyperparameter selection. Our hyperparameters were selected by hand through extensive subjective evaluation (by us) of attack perceptibility and attack success.

Computation resources. For all experiments, we used a GPU with at least 40 GB of VRAM (generally an A40, A100, or L40S), 8 CPU cores, and 40 GB of RAM. Except where noted, experiments took less than 2 hours with this setup.

F HUMAN STUDY DETAILS

F.1 DESIGN AND COUNTERBALANCING

Transcription comprised 18 clips per participant: for each $\{$ mode $\}$ \times $\{$ background $\}$ \times $\{$ stimulus type $\}$ combination we included one clip, where stimulus types were (**Attacked Input**) noisy+reverberant speech with an imperceptible perturbation, (**Clean Output**) model output for clean input, and (**Attacked Output**) model output for attacked input. To avoid leakage and learning effects, no clip was reused across stimulus types for the same participant. *ABX* comprised 12 pairs per participant constructed from the attacked-input and clean-output pools. We also included two ABX attention checks with clearly audible degradations: these followed the same design as the primary samples, except their attacks were generated with a *positive* masking threshold offset of 16.0 (other attacks in the human study used a negative offset of -16.0). Clip identities were fixed across participants to standardize difficulty and reduce computational burden.

The second seed (sample) out of 20 was manually omitted from the study due to being nearly-unintelligible whispering. Otherwise, for simplicity, assignment of seeds to (model, stimulus type, background) configurations was done through iteration using that hierarchy: iterating from 0, 2, 3, ..., 18, we changed background every seed, stimulus type every two seeds, and model every six seeds.

F.2 STATISTICS

Let $p \in \mathcal{P}$ index participants and $i \in \mathcal{I}$ index items (clips). For the transcription task, let $Y_{pi,c} \in [0,1]$ denote the word accuracy (WAcc) for condition $c \in \{\text{Attacked Input (AI), Clean Output (CO), Attacked Output (AO)}\}$ on the (p,i) cell (observed when that participant transcribed that clip under that condition). Let $\mu_c = \mathbb{E}[Y_{pi,c}]$ be the population mean WAcc for condition c.

Two-way ("pigeonhole") bootstrap for crossed effects. To obtain CIs that respect crossed participant–item variance (Owen, 2007), we independently resample participants and items with replacement. On each bootstrap replicate $b=1,\ldots,B$, draw counts $U_p^{(b)}$ for $p\in\mathcal{P}$ from

Multinomial($|\mathcal{P}|; 1/|\mathcal{P}|, \ldots$) and $V_i^{(b)}$ for $i \in \mathcal{I}$ analogously; form replicate weights $W_{pi}^{(b)} = U_p^{(b)} V_i^{(b)}$. For any condition c, the replicate mean is

$$\bar{Y}_{c}^{*(b)} = \frac{\sum_{(p,i) \in \mathcal{D}_{c}} W_{pi}^{(b)} Y_{pi,c}}{\sum_{(p,i) \in \mathcal{D}_{c}} W_{pi}^{(b)}},$$

where $\mathcal{D}_c \subseteq \mathcal{P} \times \mathcal{I}$ is the set of observed cells under c (missing cells receive weight 0 by construction). Percentile one-sided $(1-\alpha)$ CIs use the empirical quantiles of $\{\bar{Y}_c^{*(b)}\}_{b=1}^B$.

Primary claim via an intersection–union test (IUT). We test that the attacked output is *simultaneously* less intelligible than both the attacked input and the clean output:

$$H_0: (\Delta_1 \ge 0)$$
 or $(\Delta_2 \ge 0)$ vs $H_1: (\Delta_1 < 0)$ and $(\Delta_2 < 0)$,

where the mean differences are $\Delta_1 = \mu_{AO} - \mu_{AI}$ and $\Delta_2 = \mu_{AO} - \mu_{CO}$. By the IUT, an overall level- α test rejects H_0 iff *each* elementary one-sided test rejects at level α . Operationally, using the two-way bootstrap, compute the $(1 - \alpha)$ *upper* percentile bounds

$$U_k = \mathrm{Quantile}_{1-\alpha} \big(\{ \Delta_k^{*(b)} \}_{b=1}^B \big), \qquad \Delta_k^{*(b)} = \bar{Y}_{\mathrm{AO}}^{*(b)} - \bar{Y}_{(\cdot)}^{*(b)}, \ \ k \in \{1,2\},$$

with $(\cdot)=$ AI for k=1 and $(\cdot)=$ CO for k=2. Reject H_0 (and conclude AO is worse than both) iff $U_1<0$ and $U_2<0$. In our data, the 95% upper bounds were $U_1=-0.464$ and $U_2=-0.458$, yielding rejection. Note that, by construction, the IUT does not need a multiple-comparisons correction.

ABX discrimination. Let $Z_{pi} \in \{0,1\}$ indicate whether participant p correctly identified the attacked clip in pair i. We test $H_0: \mu_{\text{ABX}} \leq 0.5$ vs $H_1: \mu_{\text{ABX}} > 0.5$, where $\mu_{\text{ABX}} = \mathbb{E}[Z_{pi}]$. Using the same two-way bootstrap on $\{Z_{pi}\}$, form the *lower* one-sided $(1-\alpha)$ percentile bound L for μ_{ABX} and reject H_0 iff L > 0.5. Our lower 95% bound was L = 0.478, so we did not reject.

G LICENSE DETAILS

See Table 3.

Table 3: External assets (excludes standard Python libraries installed with pip, versioning for which is included in the code's requirements.txt).

Asset	Authorship / citation	License	Version / commit	Special terms of use
openai- Whisper	Radford et al. (2022);	MIT	v20240930/ 90db0de	None beyond MIT
MP-SENet	Lu et al. (2023);	MIT	611cd89	None beyond MIT
Denoiser	Defossez et al. (2020);	MIT	8afd7c1	None beyond MIT
FRCRN (ClearerVoice- Studio)	Zhao et al. (2022a)	Apache-2.0	634f004	None beyond Apache-2.0
MaskGCT (Amphion)	Wang et al. (2024a)	MIT	f7cb4b4	None beyond MIT
DNS- Challenge (datasets & code)	Dubey et al. (2022);	CC-BY-4.0 (data), MIT (code)	4dfd2f6	Attribution required for dataset
DNSMOS P.835	Reddy et al. (2022); same repo/commit as above	CC-BY-4.0 (data), MIT (code)	4dfd2f6	Attribution required
FullSubNet- Plus	Chen et al. (2022);	Apache-2.0	0d11530	None beyond Apache-2.0
NISQA	Mittag et al. (2021);	MIT	fe84f0f	None beyond MIT
ViSQOL	Hines et al. (2015);	Apache-2.0	b2b2a64	None beyond Apache-2.0
ViSQOL Docker image	<pre>jonashaag/visqol, Docker Hub; https://hub.do cker.com/r/jonasha ag/visqol</pre>	Not declared (container redistributes ViSQOL Apache-2.0)	'v3'	Used only for build convenience; inherits upstream license