

Sequential Randomized Smoothing for Adversarially Robust Speech Recognition

Raphael Olivier and Bhiksha Raj

Carnegie Mellon University

Language Technologies Institute

5000 Forbes Avenue, Pittsburgh PA 15213, USA

rolivier@cs.cmu.edu, bhiksha@cs.cmu.edu

Abstract

While Automatic Speech Recognition has been shown to be vulnerable to adversarial attacks, defenses against these attacks are still lagging. Existing, naive defenses can be partially broken with an adaptive attack. In classification tasks, the Randomized Smoothing paradigm has been shown to be effective at defending models. However, it is difficult to apply this paradigm to ASR tasks, due to their complexity and the sequential nature of their outputs. Our paper overcomes some of these challenges by leveraging speech-specific tools like enhancement and ROVER voting to design an ASR model that is robust to perturbations. We apply adaptive versions of state-of-the-art attacks, such as the Imperceptible ASR attack, to our model, and show that our strongest defense is robust to all attacks that use inaudible noise, and can only be broken with very high distortion.

1 Introduction¹

1.1 The threat of adversarial attacks on ASR

In recent years, Automatic Speech Recognition (ASR) has been transitioning from a topic of academic research to a mature technology implemented in everyday devices. AI voice assistants are becoming increasingly popular, and ASR models are being implemented in cars, smart TVs, and various devices within the *Internet of Things*. Therefore, challenges to the security of these models, as explored in several recent articles, are also transitioning from academic curiosities to real-world threats.

One of these major security threats is vulnerability to *adversarial attacks* (Szegedy et al., 2014): perturbations of inputs to any model that, while nearly imperceptible to human senses, have considerable effects on its outputs. Such attacks can

enable a malicious party to discretely manipulate models' behaviors and cause them to malfunction, while escaping human observation. For instance, when applied to voice assistants, adversarial attacks can potentially lead to privacy breaches by successfully soliciting arbitrary sensitive information. They could also fool an ASR system to believe an audio input contains hateful content, and have it automatically rejected from platforms or its author banned.

For several years now, adversarial attacks have been an active research field that crosses nearly every application of Machine Learning. One of the main objectives of the field is to defend AI models against such attacks without impacting their performance (on regular data) heavily.

1.2 Limitations of current defenses

This research has taken the form of an arms race, where the attacker has the upper hand: whenever a defense was proposed (Samangouei et al., 2018), a stronger or *adaptive* attack was developed to counter it (Athalye et al., 2018). Some recent works seem to have partially broken this trend by proposing defenses with precise claims, that are *optimal* in a specific sense or *certified* against specific classes of attacks. These defenses mostly fit into three categories: *Adversarial training* using strong attacks like PGD (Madry et al., 2018), *Convex relaxations* of the adversarial training objective (Wong and Kolter, 2018), and noise-based *randomized smoothing* methods (Cohen et al., 2019).

These defenses, however, have all been proposed on classification tasks, and their extension to ASR is not trivial. Adversarial training, which is already time and resource-consuming for classification, is even harder to use in speech recognition where strong attacks are much longer to compute. Convex relaxation is heavily architecture-dependent, and the use of recurrent networks or different activations makes it hard to adapt to ASR. Randomized

¹For reproducibility purposes our code and models are available at <https://github.com/RaphaelOlivier/smoothingASR>

smoothing is more promising because its simple gaussian noise-addition method makes it, in principle, usable on any model without concerns for how the attacks are computed. In reality however, there are still major challenges. ASR models are typically more susceptible than classification models to the same amount of gaussian noise. Besides, to retain good model performance with large amounts of random noise, smoothing methods require to run multiple iterations of the randomized model and use a majority vote on the outputs. When evaluating sentences in the English alphabet, the set of outputs is exponentially large in the length of the output, and majority vote is unlikely to estimate accurately the most probable one within a reasonable number of iterations.

1.3 Our contributions

Overcoming these challenges is the object of our work. Since general-purpose machine learning defenses have significant limitations when applied to speech, we improve them by leveraging the tools developed by the Speech Processing community. To use randomized smoothing on ASR while retaining good clean performance, we consider speech enhancement methods to make the defended model more accurate on gaussian-augmented inputs. We also replace the majority vote with a strategy based on the “Recognizer Output Voting Error Reduction” (ROVER) (Haihua et al., 2009) method. Depending on whether we apply training data augmentation, we provide both an off-the-shelf defense and one that requires specific fine-tuning.

We apply our defenses to a DeepSpeech2 and a Transformer model, trained and evaluated on the LibriSpeech dataset. We test them against strong attacks like the CW attack (Carlini and Wagner, 2018), the Imperceptible ASR Attack (Qin et al., 2019) and the (untargeted) PGD attack (Madry et al., 2018). We run adaptive versions of these attacks to avoid obfuscation effects. Our best model shows strong robustness against these attacks: to achieve partial transcription of the target sentence, attack algorithms require 10 times larger perturbations. Under equal noise distortions, the Word-Error-Rate (WER) on the ground truth under denial-of-service attacks improves by 30 to 50% for our model compared to the baseline.

2 Related Work

2.1 Attacks

Numerous general adversarial attacks have been proposed in the past (Szegedy et al., 2014; Goodfellow et al., 2015; Moosavi-Dezfooli et al., 2016; Carlini and Wagner, 2017; Madry et al., 2018). A few others specifically targeted audio inputs: the earliest was the ultrasonic-based DolphinAttack (Zhang et al., 2017) and the Houdini loss for structured models (Cisse et al., 2017), followed by the effective and popular Carlini&Wagner (CW) attack for audio (Carlini and Wagner, 2018). Other works have extended the state-of-the-art with over-the-air attacks (Yuan et al., 2018; Yakura and Sakuma, 2019; Li et al., 2019) and black-box attacks that do not require gradient access and transfer well (Abdullah et al., 2021). A recent line of work has improved the imperceptibility of adversarial noise by using psychoacoustic models to constrain the noise rather than standard L_2 or L_∞ bounds (Szurley and Kolter, 2019; Schönherr et al., 2019; Qin et al., 2019).

2.2 Defenses

While a large amount of defenses against adversarial attacks have been proposed (Papernot et al. (2015); Buckman et al. (2018); Samangouei et al. (2018) are just examples), the vast majority have been broken using either a strong or an adaptive attack (Carlini and Wagner, 2017; Athalye et al., 2018). Only a handful of defense families have stood the test of time. One is adversarial training, in the form proposed by Madry et al. (2018) as well as more recent variations (Wong et al., 2020; Tramer and Boneh, 2019). Noise-based, smoothing methods are another (Cao and Gong, 2017; Li et al., 2018; Liu et al., 2018; Lécuyer et al., 2019; Cohen et al., 2019). Finally, some methods prove robustness by investigating a relaxation of the adversarial objective (Gowal et al., 2018; Wong and Kolter, 2018; Mirman et al., 2018) or an exact solving (Katz et al., 2017; Bunel et al., 2018).

Efforts to adapt adversarial training or relaxation methods to ASR have been limited so far: Sun et al. (2018) have used training for speech based on the FGSM attack, which is simpler but not nearly as robust as PGD training. Most proposed ASR defenses such as MP3 compression (Das et al.) or quantization (Yang et al., 2019) have shown the same weakness as above to adaptive attacks (Subramanian et al., 2019). Exploiting temporal depen-

dencies in speech to detect adversarial manipulations (Yang et al., 2019) is a promising line of work. However, at best it only enables the user to *detect* these modified inputs. *Reconstructing* the correct transcription is an entirely different challenge, and our objective in this work.

2.3 Randomized smoothing for ASR

Some noise-based defenses for *audio classification* have been proposed: Subramanian et al. (2019) for instance use simple white noise as a defense mechanism. This is a straightforward extension of randomized smoothing to another classification setting.

Regarding specifically ASR, the only existing randomized smoothing works we are aware of are Mendes and Hogan (2020), who propose an adaptation of the noise distribution to psychoacoustic attacks, and the recent Želasko et al. (2021). The latter in particular thoroughly explores the effects of gaussian smoothing on DeepSpeech2 and the Espresso Transformer. However, their work on making these models more robust to white noise is limited to gaussian augmentation in training. Specifically they do not explore the issue of voting on transcription and resort to one-sentence estimation (see Section 5.1), which limits the amount of noise they can use, and therefore the radius of their defense. Besides, they do not use adaptive attacks (Section 3.3) which makes their evaluation incomplete. To our knowledge, we propose the first *complete* (randomization, training and vote, evaluated on adaptive attacks) version of randomized smoothing for Speech Recognition.

3 Adversarial attacks on Speech Recognition

As in previous work, we evaluate our defenses against white-box attacks, that can access model weights and their gradients and are aware of the defenses applied. Provided with an input, these attacks will run gradient-based iterations to craft an additive noise. They are the hardest attacks to defend against, and a great metric to evaluate defenses that will carryover well to more practical attacks run over-the-air, without gradient access or in real time (Li et al., 2019; Abdullah et al., 2021).

We consider two threat models. *Untargeted attacks* generate a small, additive adversarial noise that causes a denial-of-service (DOS) by altering drastically the transcription. *Targeted attacks* on

the other hand craft an additive noise that forces the model to recognize a specific target of the attacker’s choice, such as "OK Google, browse to evil.com" (Carlini and Wagner, 2018). As their objective is more precise than simple denial-of-service, targeted attacks typically require slightly larger perturbations.

We now present the specific attacks that we use. Perturbed samples for all these attacks are provided as supplementary material.

3.1 Untargeted attacks

Projected Gradient Descent The PGD attack (Madry et al., 2018) crafts a noise δ that generates mistranscriptions by maximizing the loss under its perturbation budget. It optimizes the objective

$$\max_{\|\delta\|_\infty \leq \epsilon} L(f(x + \delta), y)$$

using Projected Gradient Descent²: it takes gradient steps that maximize the loss

$$\delta_n \leftarrow \delta_{n-1} + \eta L(f(x + \delta_{n-1}), y)$$

and projects δ_n on the ball of radius ϵ after each iteration. We use 50 gradient steps when running this attack.

Rather than fixing a value for ϵ over all sentences, it is more interesting to bound the relative amount of noise compared to the input, that is the **signal-noise-ratio** (SNR) expressed in decibels:

$$SNR(\delta, x) = 20 * \log_{10} \left(\frac{\|x\|_2}{\|\delta\|_2} \right)$$

When running PGD attacks, we set a SNR threshold, then derive for each utterance the L_∞ bound $\epsilon = \frac{\|x\|_2}{10^{\frac{SNR}{20}}}$.

3.2 Targeted attacks

Carlini&Wagner attack The CW attack (Carlini and Wagner, 2018) is a targeted attack, specifically designed against CTC models. For a specific attack target y_T it minimizes the objective:

$$L(f(x + \delta), y_T) + \lambda \|\delta\|_2$$

. This attack is *unbounded*, which means it does not fix a threshold for how large δ should be. Instead, it will regularly update its regularization parameter λ to find the *smallest successful* perturbation. Therefore, the most interesting metric to evaluate a

²or Ascent, in this case

model under this attack is the SNR it achieves. Carlini and Wagner (2018) report SNRs between 30 and 40dB on the undefended DeepSpeech2 model.

To run this targeted attack, we fixed 3 target sentences of different lengths, constant in all our experiments. We try to perturb each input utterance until the model generates one of the targets (the one of closest length). For example, all utterances of less than 3-8 words are perturbed to predict the target "Really short test string".

Imperceptible ASR attack This attack proposed by Qin et al. (2019) is a variation of the CW attack for ASR (Carlini and Wagner, 2018) that adds a second objective, where masking thresholds are computed on specific frequencies, to make the noise less perceptible by the human ear. The Imperceptible attack does not improve the SNR budget of the CW attack, only how these examples are perceived by the listener under fixed budget. Therefore reporting its results is superfluous. We however provide samples generated by this attack along with this article.

3.3 Adaptive attacks against defended models

Our defenses use randomized (gaussian smoothing) and non-differentiable (speech enhancement) preprocessing steps. As (Athalye et al., 2018) have shown such elements can obfuscate the gradients and lead authors to wrongfully assume that a defense is robust. We follow the recommendations of that paper, and adapt our attacks to alleviate these effects, using two techniques:

- *Straight-through estimator*: when flowing gradients through the non-differentiable preprocessing module, we approximate its derivative as the identity function.
- *Expectation over Transformation*: since our model is stochastic, rather than just applying backpropagation once to compute gradients, we average the gradients returned by 16 backpropagation steps.

We illustrate the need for such attacks in appendix A.

4 Randomized smoothing for speech recognition

4.1 Randomized smoothing for classification

The idea of defending models against attacks by adding random noise to the inputs was formalised and generalized in Cohen et al. (2019) for classification. The idea is to replace the deterministic classifier $f : \mathbb{R}^d \rightarrow \{1, 2, \dots, m\}$ with the smooth classifier:

$$g(x) = \operatorname{argmax}_{k \in \{1, 2, \dots, m\}} \mathbb{P}(f(x + \epsilon) = k)$$

with $\epsilon \sim \mathcal{N}(0, \sigma^2 I)$. More precisely, since classifier g cannot be evaluated exactly, it is estimated with a form of Monte Carlo algorithm: many noisy forward passes are run and majority vote determines the output label.

The underlying reasoning behind this method is that given a small perturbation δ , and a standard deviation $\sigma \gg \|\delta\|_2$, probability distributions for $x + \epsilon$ and $x + \delta + \epsilon$ are very "close" by standard divergence metrics. Therefore discrete estimators built around these distribution have equal value with high probability. So if an attacker crafts an adversarial perturbation δ , it will have a very small chance of changing the output of g .

When using randomized smoothing, the biggest challenge is to retain good performance on very noisy data. One can see this defense as a way to shift the problem from adversarial robustness to white noise robustness. This is typically done with data augmentation during training.

4.2 Extension to variable-length data

The variable length of speech inputs is not an issue to use randomized smoothing. The main consequence is that the L_2 norm of a perturbation scales with the utterance length. Since Signal-Noise Ratio is normalized by utterance length, this does not affect our experiments.

A bigger problem is the nature of the text transcriptions output by the model. The number of possible outputs is exponential in the length of the input, and the probability mass of each of them under noisy inputs is extremely small. Therefore, majority vote cannot estimate the probabilities of the transcriptions in practice, as we discuss in Section 5.1.

However, the reasoning that Gaussian distributions centered on a utterance x and an adversarially perturbed one $x + \epsilon$ are close is still valid.

This tends to show that the overall noise-additive method still makes natural and adversarial points similar from the model’s perspective.

4.3 Gaussian noise-robust speech recognition models

As mentioned above, when using randomized smoothing it is critical to retain good performance on gaussian-augmented inputs. With ASR models this is not a trivial objective. We consider the following techniques to achieve this goal.

Augmented Training Rather than training a neural ASR model entirely on gaussian-augmented data, we used a pretrained model on clean data that we fine-tune with gaussian augmentation for one epoch. We find that it helps training converge and leads to similar or better results on noisy data in a much shorter time.

Speech enhancement Speech enhancement algorithms help improve audio quality. After adding gaussian noise, we can use enhancement to restore the original audio quality. We tried multiple standard enhancement methods and found a-priori SNR estimation (ASNR) (Scalart and Filho, 1996) to be most effective. Neural methods such as SEGAN (Pascual et al., 2017) did not reach the same performance (in terms of WER in the end-to-end ASR pipeline), most likely because these models are tailored to real-world-like noise conditions and are not trained on gaussian noise. It is possible, though not certain, that a generative model trained on gaussian noise would improve the enhancement results: Pascual et al. (2017) argue that they outperform first-order filters like ASNR specifically for complex noise conditions.

5 Voting strategies on text outputs

Even with augmented training and/or enhancement, when feeding noisy inputs to ASR models the output distribution has high variance. Running multiple forward passes and "averaging" the outputs can help reduce that variance and improve accuracy. But this requires a good voting strategy on text outputs.

We first discuss some elementary strategies and their potential drawbacks, then describe the ROVER-based vote that we use. All of these strategies are empirically compared in Section 7.1. We denote the sampled transcriptions as t_1, \dots, t_n and t is our final transcription.

5.1 Baseline strategies

One-sentence estimation A solution that has the merit of simplicity is to not vote at all. Using only one input, we can hope that the sentence we get is "close" to the most probable sentence (in terms of Word-Error Rate for instance) and just return it as our output. This is the strategy used by Želasko et al. (2021).

Majority vote Following the original randomized smoothing defense, we can vote at sentence level: $t = \operatorname{argmax}_{t' \in \{t_1, \dots, t_n\}} \operatorname{card}(\{i/t_i = t'\})$

Designed for classification, majority vote is not adapted to probabilistic text outputs. The set T of all possible transcriptions is infinite, and even with our most stable models and a relative noise of -15 dB, 100 noise samples typically output 100 different transcriptions. Even without setting up rigorous statistical tests, it is clear that outputting the likeliest transcription, or just a "likelier than average" one, with high probability would require thousands of ASR iterations, which in practice is not feasible. In other words, majority vote is barely better than one-sentence estimation, for a high computation cost.

Statistics in the logits space For a given input utterance length, some ASR architectures, such as CTC-trained models (Graves et al., 2006), first generate fixed-length logits sequences l_1, \dots, l_n , then apply a best-path decoder d to generate transcriptions $t_i = d(l_i)$. It is then possible to aggregate these logits over the random inputs, for example by averaging them, then to apply the decoder: $t = d(\frac{1}{n} \sum_i l_i)$.

One potential issue with this strategy as a defense is that it distances itself from the randomized smoothing framework, where the use of *discrete* outputs to vote on is critical. To get a concrete idea of how this could be a problem, one should remember that adversarial examples can be generated with *high confidence* (aka very large logits). Such a phenomenon could disrupt the statistic by over-weighting the fraction of inputs that are most affected by an adversarial perturbation.

5.2 ROVER

The Recognizer Output Voting Error Reduction (ROVER) system was introduced by NIST in 1997, as an ensembling method that mitigates the different errors produced by multiple ASR systems. Contrary to majority vote it works at the word-level

	Model	$\sigma = 0$	$\sigma = 0.01$	$\sigma = 0.02$
<i>Baseline</i>	Deepspeech2	9.7	27	63
	Transformer	5.7	12	35
<i>Enhancement, augmentation</i>	Deepspeech2+ σ -AUG	-	11	16
	Deepspeech2+ ASNR	-	17	33
	Deepspeech2+ σ -AUG + ASNR	-	12	18
	Transformer+ ASNR	-	9.1	19
<i>Voting</i>	Deepspeech2+ Maj-100	-	27	69
	Deepspeech2+ Avg-100	-	25	62
	Deepspeech2+ ROVER-8	-	21	52
	Deepspeech2+ ROVER-16	-	20	50
	Deepspeech2+ ROVER-32	-	20	50
<i>Proposed models</i>	Deepspeech2+ ASNR + ROVER-16	-	14	26
	Deepspeech2+ σ -AUG + ROVER-16	-	9	12
	Transformer+ ASNR + ROVER-16	-	8.1	15

Table 1: Word Error Rate (%) for Deepspeech2 on the LibriSpeech clean test set under various defenses on *clean* utterances when adding gaussian noise of deviation σ . σ -AUG stands for gaussian augmentation of deviation σ in training - the same deviation used at inference. ASNR means A priori SNR filtering of inputs. Maj-N, Avg-N, ROVER-N refer to majority vote, logits averaging and ROVER voting strategies, using N forward passes.

rather than the sentence level, by selecting at each position the word present in the most sentences. ROVER should be fed the time duration of each word in the audio space, which we can extract using audio-text alignment information provided by the ASR models (see Section 6.2). We use ROVER as a black-box script and understanding its inner behavior is not absolutely necessary to follow this work, however we provide a high-level explanation of this algorithm in Appendix B

In our work we introduce an alternative use of ROVER, as a voting system on the text outputs of the same probabilistic model rather than for ensembling multiple models. We mostly use it as a black box, by feeding to ROVER multiple output sequences. We also feed

The main drawbacks of this method lie in the time penalty of the voting module when using a large number of inputs. We further discuss that limitation in Section 8.1

6 Experiments

6.1 Dataset

We run all our experiments on the 960 hours LibriSpeech dataset (Panayotov et al., 2015), and report our results on its test-clean split. As adversarial attacks can take a considerable amount of time to compute, we evaluate attacks on the first 100 utterances of this test set.

6.2 Models

We test our smoothing methods on two model architectures :

- The CTC-based DeepSpeech2 (Amodei et al., 2016), a standard when evaluating adversarial attacks on ASR since Carlini and Wagner (2018). We pretrain it on the clean LibriSpeech training set. As discussed above, we fine-tuned it on gaussian-augmented data for one epoch, using always the same deviation used at inference for smoothing. For decoding we use greedy search, as we find that increasing the beam size has very little impact on WER for this model. This is a relatively lightweight model that we use for ablation experiments. The CTC decoder provides frame alignments for each transcription character : we use them to infer word duration (needed for ROVER) with good precision.
- A more recent Transformer architecture. We adapt the Espresso implementation (Wang et al., 2019) to our code. Training and hyperparameter search for transformer models can be computationally expensive, and reaching state-of-the-art word-error-rate is unnecessary in this work. For those reasons we keep all the hyperparameters of Espresso’s "Librispeech Transformer" architecture, and do not fine-

tune this model on gaussian noise. We also only run untargeted attacks on the transformer model, as targeted attack algorithms and implementations are usually model specific (and most often proposed against DeepSpeech2). We do not know of any available implementation of a targeted attack on ASR transformer models.

This transformer implementation does not output character alignment. It however provides word-level attention scores with the encoded audio. We can align each word with the highest-scoring audio vector, and obtain a word-level alignment. This method is less precise than with DeepSpeech2 and can likely be improved.

6.3 Defenses

Our models are defended with gaussian noise, ASNR enhancement, voting strategies or a combination of all of these. The noise deviation is set to $\sigma = 0.1$ or $\sigma = 0.2$ depending on the experiments, which corresponds for the vast majority of utterances to a signal-noise ratio in the 10-14dB and 7-11dB respectively.

Against adversarial examples, we compare our models with each other, as well as with the undefended DeepSpeech2 model and a baseline defense using MP3 Compression (Carlini and Wagner, 2018).

6.4 Evaluation metrics

We evaluate our models under untargeted attacks with **Word Error Rate** (WER) which is the word-level edit distance between two transcriptions normalized by the length of the target. In case of large mistakes we upper bound WER values to 100% (while the real value can be greater if the generated sentence is longer than the target). We report this WER on the ground-truth: the lower the WER the better the defense.

With targeted attacks, we also report Word-Error-Rate, both on the ground truth and the attack target. However, since the CW attack is unbounded, if applied with good hyperparameters it always succeeds in forcing our model to deviate from the ground truth (high WER) and predict the attack target (low WER). Therefore these metrics mainly have value as a sanity check to make sure we run the attack correctly. To measure whether a defense is effective against CW, our primary evaluation met-

ric is the **Signal-Noise-Ratio** (SNR), which quantifies the *amount* of noise the attack generates to achieve its objective. A SNR above 20 – 25 would typically be hard to perceive for a human.

Model	No attack	CW attack		
		GT	TGT	SNR
Vanilla	8	100	0	27dB
MP3 compression	9	100	3	16dB
Trained ($\sigma = 0.01$)	11	100	8	14dB
Off-the-shelf ($\sigma = 0.01$)	17	100	4	10dB
Trained ($\sigma = 0.02$)	14	100	6	8dB
Off-the-shelf ($\sigma = 0.02$)	31	100	6	5dB

Table 2: Word Error Rate (%) for DeepSpeech2 on the first 100 utterances of the LibriSpeech clean test set, for clean inputs and under the CW attack for the baselines and the proposed trained and off-the-shelf defenses, with gaussian deviations 0.01 and 0.02. We report both the WER on the ground truth (GT) and the attack target (TGT), and the SNR required to achieve it.

7 Results

We first show that our defended models, which add gaussian noise to all inputs, retain low WER on this noisy but non-adversarial data. Then we report their performance against adversarial attacks. We show that they successfully recover all attacks that use near-imperceptible noise.

7.1 Performance under gaussian noise

We report the performance of our model on noisy inputs (but no attack) in Table 1.

Augmentation and Enhancement We evaluate gaussian augmentation on DeepSpeech2, and ASNR on both our architectures. Both techniques lower the word-error rate significantly under $\sigma = 0.01, 0.02$, with an advantage for the former. Interestingly however, combining both techniques at once (on a DeepSpeech2 model trained on noisy *and* enhanced data) does not really improve results compared to using augmentation only. This suggests to use ASNR enhancement as a "fallback option" in situations where retraining a model is not acceptable. This off-the-shelf method nonetheless provides competitive performance when using state-of-the-art architectures like Transformer.

Voting strategies We compare all our proposed voting strategies on DeepSpeech2 outputs. As expected, majority vote brings no significant improvement over one-sentence estimation (i.e. the base-




Figure 1: WER achieved against PGD attacks by the baselines (Vanilla model, MP3 compression) and the proposed trained (DeepSpeech2 with Smoothing, augmented training and ROVER) and off-the-shelf (DeepSpeech2/Transformer with Smoothing, ASNR and ROVER) defenses, with gaussian deviations 0.01 and 0.02. We plot the results when varying the PGD SNR bound : a lower SNR means a larger perturbation.

line). Logits averaging is somewhat effective ; however it does not compare to ROVER, by far the best voting methods even with fewer sentences.

Proposed models As a consequence, we propose two smoothing-based defenses:

- a **trained** defense using smoothing, augmentation and ROVER. On DeepSpeech2, with $\sigma = 0.01$ (resp. 0.02) it reaches a WER of 9 (resp. 12)
- an **off-the-shelf** defense using smoothing, ASNR enhancement and ROVER. Applied to DeepSpeech2 this defense suffers from a higher WER of 14 (resp. 26) but is still relatively effective. With Transformer it performs much better with a WER of 8.1 (resp 15).

7.2 Performance under attack

In Figure 1 we plot the results of our defenses and baselines against the untargeted PGD attack, as a function of the SNR used to bound the attack. They demonstrate the effectiveness of ASR smoothing: compared to the vanilla DeepSpeech2 the Word-Error-Rate improves by 20 to 50 points for all PGD

attacks bounded by $SNR \geq 20dB$ for our proposed models, with both $\sigma = 0.01$ and $\sigma = 0.02$. When $SNR = 25dB$ is sufficient to reach total denial-of-service ($WER = 100$) on DeepSpeech2, and $20dB$ for the MP3 compression baseline, the same feat requires SNRs of $10 - 15dB$ to defeat our defenses.

Table 2 reports the results of the targeted CW attack. As expected for an unbounded attack, it is partially able to break our defenses (low WER on its target), but at a high cost. The SNR it requires to achieve these results is as low as $10 - 14$ under $\sigma = 0.01$ and $5 - 8$ with $\sigma = 0.02$! This compares to $27dB$ for the undefended DeepSpeech2 and $16dB$ for the MP3 baseline. With the higher σ in particular, the adversarial noise becomes very much audible to the human ear, even when refining it with the Imperceptible attack (Section 3.2). The tradeoff in clean (no attack) accuracy is fairly low for the trained defense even with $\sigma = 0.02$ (+5% WER). It is higher for the untrained, off-the-shelf defense, with which using a lower deviation may be required for practical applications.

8 Discussion

8.1 ROVER accuracy/time tradeoff




Figure 2: Average ROVER Word-Error Rate and computation time (DeepSpeech2) for a single utterance when varying N .

One drawback of ROVER voting is its important time consumption when using many inputs. This may be partially due to its third-party, black-box implementation that does not use GPU computation. However, in Figure 2 we show that ROVER voting time increases superlinearly with the number of samples (where averaging and counting are of course linear): this is most likely an irreducible complexity of the algorithm. Using more than 50 iterations is not practically feasible³. This is why we use $N = 16$ in most of our experiments: even though more iterations may bring marginal WER improvement, this value enables us to improve performance substantially while keeping the voting time negligible. Besides, 16 inputs can typically be fed to the model in one batch, thus keeping the overall computation time low.

8.2 Certifying randomized smoothing on ASR

While our main focus in this work is to reach strong empirical performance under attack, we also show that adversarial robustness can to some extent be *proven* for Speech Recognition, as is the case with classification (Cohen et al., 2019). We show that we can prove the following result:

Proposition 1 *If for a sentence s the randomized ASR model f verifies*

$$\mathbb{P}[WER(f(x), s) < k] \geq p > 0.5$$

³It is, in fact, forbidden by default in the publicly available implementation

then for any noise $\|\delta\|_2 < R$:

$$\mathbb{P}[WER(f(x + \delta), s) < k] > 0.5$$

where $R = \frac{\sigma}{2}(\phi^{-1}(p) - \phi^{-1}(1 - p))$ and ϕ is the standard gaussian CDF.

This means that if with high probability the gaussian-smoothed model does not deviate "too much" from a sentence s in terms of WER, then the same remains true when adding a small perturbation δ . This proposition can be used to write an algorithm that certifies a transcription. We defer the proof to appendix C.

In practice such guarantees are very hard to compute: these certification algorithms demand thousands of forward passes to give results with any useful confidence margin, which on large ASR models remains an open problem.

9 Conclusion

We have proposed a state-of-the-art adversarial defense for ASR models based on randomized smoothing. It is successful against all attacks using inaudible distortion, while retaining a low error rate on natural data. To achieve strong performance under noise, we have leveraged speech enhancement methods and proposed a novel use for ASR output ensembling methods like ROVER. We successfully defend against state-of-the-art adaptive attacks, and analyse the importance and limits of each component of our defense. Finally, we show that Randomized Smoothing on ASR is to some extent a provably robust defense.

This work paves the way for a thorough exploration of smoothing defenses for ASR. Practical certification, extension to other architectures and ensembling with other defenses are some areas of interest. Our approach could also be crossed with Mendes and Hogan (2020) to generate noise that defends specifically against psychoacoustic-based attacks (Qin et al., 2019).

10 Acknowledgements

This material is based upon work supported by the U.S. Army Research Laboratory and DARPA under contract HR001120C0012. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the the U.S. Army Research Laboratory and DARPA.

References

- H. Abdullah, M. Rahman, W. Garcia, K. Warren, A. Swarnim Yadav, T. Shrimpton, and P. Traynor. 2021. [Hear "no evil", see "kenansville"*: Efficient and transferable black-box attacks on speech recognition and voice identification systems](#). In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 142–159, Los Alamitos, CA, USA. IEEE Computer Society.
- Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Qiang Cheng, Guoliang Chen, Jie Chen, Jingdong Chen, Zhijie Chen, Mike Chrzanowski, Adam Coates, Greg Diamos, Ke Ding, Niandong Du, Erich Elsen, Jesse Engel, Weiwei Fang, Linxi Fan, Christopher Fougner, Liang Gao, Caixia Gong, Awni Hannun, Tony Han, Lappi Johannes, Bing Jiang, Cai Ju, Billy Jun, Patrick LeGresley, Libby Lin, Junjie Liu, Yang Liu, Weigao Li, Xiangang Li, Dongpeng Ma, Sharan Narang, Andrew Ng, Sherjil Ozair, Yiping Peng, Ryan Prenger, Sheng Qian, Zongfeng Quan, Jonathan Raiman, Vinay Rao, Sanjeev Satheesh, David Seetapun, Shubho Sengupta, Kavya Srinet, Anuroop Sriram, Haiyuan Tang, Liliang Tang, Chong Wang, Jidong Wang, Kaifu Wang, Yi Wang, Zhijian Wang, Zhiqian Wang, Shuang Wu, Likai Wei, Bo Xiao, Wen Xie, Yan Xie, Dani Yogatama, Bin Yuan, Jun Zhan, and Zhenyao Zhu. 2016. [Deep speech 2 : End-to-end speech recognition in english and mandarin](#). In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 173–182, New York, New York, USA. PMLR.
- Anish Athalye, Nicholas Carlini, and David Wagner. 2018. [Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples](#). In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018*.
- Jacob Buckman, Aurko Roy, Colin Raffel, and Ian Goodfellow. 2018. [Thermometer encoding: One hot way to resist adversarial examples](#).
- Rudy Bunel, Ilker Turkaslan, Philip H.S. Torr, Pushmeet Kohli, and M. Pawan Kumar. 2018. A unified view of piecewise linear neural network verification. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS'18*, page 4795–4804, Red Hook, NY, USA. Curran Associates Inc.
- Xiaoyu Cao and Neil Zhenqiang Gong. 2017. [Mitigating evasion attacks to deep neural networks via region-based classification](#). In *Proceedings of the 33rd Annual Computer Security Applications Conference, ACSAC 2017*, page 278–287, New York, NY, USA. Association for Computing Machinery.
- N. Carlini and D. Wagner. 2017. [Towards evaluating the robustness of neural networks](#). In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57.
- Nicholas Carlini and D. Wagner. 2018. [Audio adversarial examples: Targeted attacks on speech-to-text](#). *2018 IEEE Security and Privacy Workshops (SPW)*, pages 1–7.
- Nicholas Carlini and David Wagner. 2017. [Adversarial examples are not easily detected: Bypassing ten detection methods](#). In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security, AISeC '17*, page 3–14, New York, NY, USA. Association for Computing Machinery.
- Moustapha M Cisse, Yossi Adi, Natalia Neverova, and Joseph Keshet. 2017. [Houdini: Fooling deep structured visual and speech recognition models with adversarial examples](#). In *Advances in Neural Information Processing Systems*, volume 30, pages 6977–6987. Curran Associates, Inc.
- Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. 2019. [Certified adversarial robustness via randomized smoothing](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 1310–1320. PMLR.
- Nilaksh Das, Madhuri Shanbhogue, Shang-Tse Chen, Fred Hohman, Siwei Li, Li Chen, Michael E. Kounavis, and Duen Horng Chau. [Compression to the rescue: Defending from adversarial attacks across modalities](#). *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.
- Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. [Explaining and harnessing adversarial examples](#). In *International Conference on Learning Representations*.
- Sven Gowal, Krishnamurthy Dvijotham, Robert Stanforth, Rudy Bunel, Chongli Qin, Jonathan Uesato, Relja Arandjelovic, Timothy A. Mann, and Pushmeet Kohli. 2018. [On the effectiveness of interval bound propagation for training verifiably robust models](#). *CoRR*, abs/1810.12715.
- Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. 2006. [Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks](#). In *Proceedings of the 23rd International Conference on Machine Learning, ICML '06*, page 369–376, New York, NY, USA. Association for Computing Machinery.
- X. Haihua, Z. Jie, and G. Wu. 2009. [An efficient multi-stage rover method for automatic speech recognition](#). In *2009 IEEE International Conference on Multimedia and Expo*, pages 894–897.
- Guy Katz, Clark Barrett, David L. Dill, Kyle Julian, and Mykel J. Kochenderfer. 2017. [Reluplex: An efficient SMT solver for verifying deep neural networks](#). In *Proceedings of the 29th International Conference*

- on *Computer Aided Verification (CAV '17)*, volume 10426 of *Lecture Notes in Computer Science*, pages 97–117. Springer, Heidelberg, Germany.
- Mathias Lécuyer, Vaggelis Atlidakis, Roxana Geambasu, D. Hsu, and Suman Jana. 2019. Certified robustness to adversarial examples with differential privacy. *2019 IEEE Symposium on Security and Privacy (SP)*, pages 656–672.
- Bai Li, C. Chen, Wenlin Wang, and L. Carin. 2018. Second-order adversarial attack and certifiable robustness. *ArXiv*, abs/1809.03113.
- Juncheng Li, Shuhui Qu, Xinjian Li, Joseph Szurley, J. Zico Kolter, and Florian Metzger. 2019. Adversarial music: Real world audio adversary against wake-word detection system. In *Advances in Neural Information Processing Systems*, volume 32, pages 11931–11941. Curran Associates, Inc.
- Xuanqing Liu, Minhao Cheng, Huan Zhang, and C. Hsieh. 2018. Towards robust neural networks via random self-ensemble. In *ECCV*.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2018. Towards deep learning models resistant to adversarial attacks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Ethan Mendes and Kyle Hogan. 2020. Defending against imperceptible audio adversarial examples using proportional additive gaussian noise.
- Matthew Mirman, Timon Gehr, and Martin Vechev. 2018. Differentiable abstract interpretation for provably robust neural networks. In *International Conference on Machine Learning (ICML)*.
- Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. 2016. Deepfool: A simple and accurate method to fool deep neural networks. pages 2574–2582.
- J. Neyman and E. S. Pearson. 1933. On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231:289–337.
- V. Panayotov, G. Chen, D. Povey, and S. Khudanpur. 2015. Librispeech: An asr corpus based on public domain audio books. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210.
- Nicolas Papernot, Patrick D. McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. 2015. Distillation as a defense to adversarial perturbations against deep neural networks. *CoRR*, abs/1511.04508.
- Santiago Pascual, Antonio Bonafonte, and Joan Serra. 2017. Segan: Speech enhancement generative adversarial network. pages 3642–3646.
- Yao Qin, Nicholas Carlini, Garrison Cottrell, Ian Goodfellow, and Colin Raffel. 2019. Imperceptible, robust, and targeted adversarial examples for automatic speech recognition. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5231–5240. PMLR.
- Pouya Samangouei, Maya Kabkab, and Rama Chellappa. 2018. Defense-GAN: Protecting classifiers against adversarial attacks using generative models. In *International Conference on Learning Representations*.
- Pascal Scalart and Jozue Filho. 1996. Speech enhancement based on a priori signal to noise estimation. volume 2, pages 629 – 632 vol. 2.
- Lea Schönherr, Katharina Kohls, Steffen Zeiler, Thorsten Holz, and Dorothea Kolossa. 2019. Adversarial attacks against automatic speech recognition systems via psychoacoustic hiding. In *Network and Distributed System Security Symposium (NDSS)*.
- Vinod Subramanian, Emmanouil Benetos, and Mark Sandler. 2019. Robustness of adversarial attacks in sound event classification. pages 239–243.
- Sining Sun, Ching-Feng Yeh, Mari Ostendorf, Mei-Yuh Hwang, and Lei Xie. 2018. Training Augmentation with Adversarial Examples for Robust Speech Recognition. *arXiv e-prints*, page arXiv:1806.02782.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2014. Intriguing properties of neural networks. In *International Conference on Learning Representations*.
- Joseph Szurley and J. Z. Kolter. 2019. Perceptual based adversarial audio attacks. *ArXiv*, abs/1906.06355.
- Florian Tramèr and Dan Boneh. 2019. Adversarial training and robustness for multiple perturbations. In *Advances in Neural Information Processing Systems*, volume 32, pages 5866–5876. Curran Associates, Inc.
- Yiming Wang, Tongfei Chen, Hainan Xu, Shuoyang Ding, Hang Lv, Yiwen Shao, Nanyun Peng, Lei Xie, Shinji Watanabe, and Sanjeev Khudanpur. 2019. Espresso: A fast end-to-end neural speech recognition toolkit. In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*.
- E. Wong, Leslie Rice, and J. Z. Kolter. 2020. Fast is better than free: Revisiting adversarial training. *ArXiv*, abs/2001.03994.
- Eric Wong and Zico J. Kolter. 2018. Provable defenses against adversarial examples via the convex outer adversarial polytope. *ICML*, pages 5283–5292.

- Hiromu Yakura and Jun Sakuma. 2019. [Robust audio adversarial example for a physical attack](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5334–5341. International Joint Conferences on Artificial Intelligence Organization.
- Zhuolin Yang, Pin Yu Chen, Bo Li, and Dawn Song. 2019. Characterizing audio adversarial examples using temporal dependency. 7th International Conference on Learning Representations, ICLR 2019 ; Conference date: 06-05-2019 Through 09-05-2019.
- Xuejing Yuan, Yuxuan Chen, Yue Zhao, Yunhui Long, Xiaokang Liu, Kai Chen, Shengzhi Zhang, Heqing Huang, XiaoFeng Wang, and Carl A. Gunter. 2018. [Commandersong: A systematic approach for practical adversarial voice recognition](#). In *27th USENIX Security Symposium (USENIX Security 18)*, pages 49–64, Baltimore, MD. USENIX Association.
- Guoming Zhang, Chen Yan, Xiaoyu Ji, Tianchen Zhang, Taimin Zhang, and Wenyuan Xu. 2017. [Dolphinattack: Inaudible voice commands](#). In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS '17*, page 103–117, New York, NY, USA. Association for Computing Machinery.
- Piotr Żelasko, Sonal Joshi, Yiwen Shao, Jesus Villalba, Jan Trmal, Najim Dehak, and Sanjeev Khudanpur. 2021. [Adversarial attacks and defenses for speech recognition systems](#).

A The need for adaptive attacks

A.1 Background

In some previous works, the authors apply attacks to their proposed defense without any modification. This amounts to assuming that the attacker *ignores* the existence of the defense, and tends to lead to inflated accuracy for the defender.

In the particular case of stochastic defenses, such as ours, it is a well-known fact that the gradients the attacker uses are noisy and less informative than those of undefended models, thus making gradient-based attacks less effective (Athalye et al., 2018). This phenomenon should not be seen as a desirable feature : rather than making adversarial examples fail to break the defense, the defense *obfuscates* them but is still vulnerable to a cautious attacker, who uses an *adaptive* attack rather than the vanilla attack.

One simple fix to the noisy gradient phenomenon is *Expectation over Transformation* (EoT). The attacker, which has only access to the stochastic model $f(x + \epsilon)$ cannot access the gradients of the deterministic model

$$\nabla_x f(x) = \nabla_x \mathbb{E}[f(x + \epsilon)]$$

They estimate them by applying the expectation to the gradients of the stochastic function :

$$\nabla_x f(x) \approx \mathbb{E}[\nabla_x f(x + \epsilon)]$$

This latter quantity can be estimated with sample mean, i.e. by averaging gradients over a batch.

A.2 Results

All the results we report in the main paper are computed against the above adaptive attack, using gradient batches of size 16. In Table 3 we report results obtain by the PGD attack on our trained defense, with $\sigma = 0.01$, with and without EoT. The WER is significantly lower using vanilla attacks, demonstrating why using adaptive attacks is necessary to correctly evaluate a defense. This also illustrates that our claims in the paper do **not** reflect obfuscation phenomena, but rather actual adversarial robustness.

B The ROVER voting algorithm

ROVER (Haihua et al., 2009) works in two steps. First, it aligns all k sentences word-by-word and

Model, Smoothing	Adaptive	SNR-PGD					
		35	30	25	20	15	10
Trained defense, $\sigma = 0.02$	No	29	39	54	66	90	100
Trained defense, $\sigma = 0.02$	Yes	34	46	60	70	92	100

Table 3: Word Error Rate (%) for Our defense on the first 100 utterances of the LibriSpeech clean test set under untargeted PGD attacks using various SNR. The first line corresponds to the vanilla attack, the second uses an adaptive attack that averages gradients on 16 forward+backward passes.

aggregates them into one Word Transition Network (WTN), i.e. a graph where nodes represent timesteps, and edges between two timesteps are word (or silent) candidates. Alignment is done iteratively : the first sentence serves as a base WTN, then for $i = 2, \dots, k$ ROVER merges sentence i with the base WTN using Dynamic Programming tool SCLITE, using a process close to Levenstein distance : it finds the minimal cost alignment using operations of substitution, insertion and deletion. These alignment steps make use of audio alignment information as well as word and sentence scores, to output a final WTN.

At this step, ROVER votes on the aligned words using (in our version) the frequency of each word. It also accepts metrics based on word confidence, which we evaluated (using DeepSpeech2’s softmax outputs as confidence scores) and found not to bring any improvement in our use case.

C Certifying ASR smoothing

C.1 Robustness properties for classification

For multiclass classification, Cohen et al. (2019) are able to provide a robustness certificate based on the probability of the most probable class A : if

$$p_A = \mathbf{P}(f(x + \epsilon) = A)$$

and

$$p_B = \max_{B \neq A} \mathbf{P}(f(x + \epsilon) = B)$$

then $g(x + \delta) = A$ for all $\|\delta\|_2 < R$ with

$$R = \frac{\sigma}{2} (\phi^{-1}(p_A) - \phi^{-1}(p_B))$$

where g is the smoothed classifier :

$$g(x) = \operatorname{argmax}_{k \in \{1, 2, \dots, m\}} \mathbf{P}(f(x + \epsilon) = k)$$

$\epsilon \sim \mathcal{N}(0, \sigma^2 I)$ and ϕ the standard gaussian CDF.

This result extends naturally from the binary classification case, which itself is a consequence of

the Neyman-Pearson lemma (Neyman and Pearson, 1933). In the case of ASR, reducing the problem to binary classification is not as trivial. We propose such a reduction by using thresholds on the evaluation metric d (typically the WER).

C.2 Robustness properties for ASR

Name f an ASR model. We assume given an *empirical target* sentence \tilde{y} (different from the golden transcription y) and a threshold $k \in]0, 1[$. We define the binary classifier \tilde{f} as :

$$\tilde{f}(x) = \begin{cases} 1 & \text{if } d(f(x), \tilde{y}) < k \\ 0 & \text{otherwise} \end{cases}$$

And $\tilde{g}(x) = \operatorname{argmax}_{c \in \{0, 1\}} \mathbf{P}(\tilde{f}(x + \epsilon) = c)$. \tilde{g} is informative only if $\tilde{g}(x) = 1$.

By immediate application of Cohen et al. (2019)’s result to \tilde{f} we obtain the following guarantee :

Proposition 2 $\tilde{g}(x + \delta) = 1$ for all $\|\delta\|_2 < R$ with

$$R = \frac{\sigma}{2} (\phi^{-1}(p) - \phi^{-1}(1 - p)), \quad p = \mathbf{P}(\tilde{f}(x) = 1)$$

C.2.1 Certification algorithm

This result allows us to use on \tilde{f} the CERTIFY algorithm from Cohen et al. (2019) (Section 3.2.2). We do not reproduce it here; the only change to our use case is that rather than generating a “top class” c_A based on counts, we use our ROVER prediction strategy to generate the “top transcription” t_A . A policy to estimate the bound k could perhaps be designed, or k can simply be fixed to a value that seems reasonable with respect to applications.

Proposition 3 *With probability at least $1 - \alpha$ over the randomness in CERTIFY, if CERTIFY returns a transcription t_A and a radius R (i.e. does not abstain), then the model predicts t_A at WER $\leq k$ within radius R around x .*

C.2.2 Feasibility

Despite our good experimental results, and while we showed in this section that robustness certification is *possible*, this algorithm requires estimating probabilities which is computationally infeasible using a large ASR model. As [Cohen et al. \(2019\)](#) note, certification with informative bounds and confidence requires many samples (in the tens of thousands). We are unable to run this many forward loops per sentence⁴ under reasonable time constraints. Nonetheless, these results show a possible way to reduce ASR to finite-class classification for certification purposes, and open a path towards guaranteeing the robustness of ASR models using faster architectures.

D Additional experiments

In Table 4 we report the results of both the PGD and CW attacks on our baselines and proposed models. We include models that are partially defended (without voting/augmented training/etc) to serve as an ablation study.

⁴Using one NVIDIA GTX 2080Ti

σ	Model, Smoothing	Nat.	SNR-PGD						CW		
			35	30	25	20	15	10	GT	TGT	SNR
0.0	Deepspeech2	8	63	89	100	100	100	100	100	0	27
	MP3 compression	9	59	71	79	100	100	100	100	3	16
0.01	AUG	12	50	64	82	100	100	100	100	0	19
	RS	34	58	66	84	95	100	100	100	13	15
	RS + AUG	11	30	46	57	82	100	100	100	11	14
	+ RS + ASNR	23	35	45	60	71	91	100	100	11	10
	RS + ROVER	29	51	64	81	93	100	100	100	7	15
	RS + σ -AUG + ROVER	9	31	43	59	81	100	100	100	8	14
	RS + ASNR + ROVER	20	34	46	60	70	92	100	100	4	10
0.02	AUG	13	46	63	83	100	100	100	100	0	15
	RS	72	79	82	87	96	97	100	100	12	10
	RS + AUG	19	35	46	64	77	100	100	100	9	8
	RS + ASNR	41	44	55	62	74	84	100	100	13	5
	RS + ROVER	62	72	77	84	93	98	100	100	7	10
	RS + AUG + ROVER	14	34	46	60	77	100	100	100	6	8
	RS + ASNR + ROVER	36	40	46	56	71	83	100	100	6	5

Table 4: Word Error Rate (%) for Deepspeech2 on the first 100 utterances of the LibriSpeech clean test set under various attacks and defenses. +AUG stands for gaussian augmentation of deviation σ in training - the same deviation used at inference. ASNR means A priori SNR filtering of inputs. +ROVER refers to the ROVER voting strategy using 16 forward passes. For the PGD attack we specify the minimal SNR we use as L_∞ bound. For the unbounded CW attack we report both the WER on the ground truth (GT) and the attack target (TGT), and the SNR required to achieve it. All attacks run on models using smoothing are adaptive and average gradients on 16 forward+backward passes.