

PNNL-35153

Universal Fourier Attack for Time Series

September 2023

Chance N DeSmet
Michael K Girard
Lizzy D Coda
Yijing Watkins

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor Battelle Memorial Institute, nor any of their employees, makes **any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights.** Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or Battelle Memorial Institute. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from
the Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062

www.osti.gov

ph: (865) 576-8401

fox: (865) 576-5728

email: reports@osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312

ph: (800) 553-NTIS (6847)

or (703) 605-6000

email: info@ntis.gov

Online ordering: <http://www.ntis.gov>

Universal Fourier Attack for Time Series

September 2023

Chance N DeSmet
Michael K Girard
Lizzy D Coda
Yijing Watkins

Prepared for
the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99354

Abstract

A wide variety of adversarial attacks have been proposed and explored using image and audio data. These attacks are notoriously easy to generate digitally when the attacker can directly manipulate the input to a model, but are much more difficult to implement in the real-world. In this paper we present a universal, time invariant attack for general time series data such that the attack has a frequency spectrum primarily composed of the frequencies present in the original data. The universality of the attack makes it fast and easy to implement as no computation is required to add it to an input, while time invariance is useful for real-world deployment. Additionally, the frequency constraint ensures the attack can withstand filtering. We demonstrate the effectiveness of the attack in two different domains, speech recognition and unintended radiated emission, and show that the attack is robust against common transform-and-compare defense pipelines.

Summary

This report outlines the background, tests, and performance of the most significant outcome of the LaserWolf LDRD, the ability to construct a phase-invariant, filtering robust method that is capable of working in multiple domains. This attack has shown success on both acoustic and power line communications and has demonstrated superior capability compared to contemporary methods.

Acknowledgments

This research was supported by the **National Security Mission Seed**, under the Laboratory Directed Research and Development (LDRD) Program at Pacific Northwest National Laboratory (PNNL). PNNL is a multi-program national laboratory operated for the U.S. Department of Energy (DOE) by Battelle Memorial Institute under Contract No. DE-AC05-76RL01830.

Contents

Abstract.....	ii
Summary	iii
Acknowledgments.....	iv
1.0 Introduction.....	1
1.1 Methods.....	1
1.1.1 Data	1
1.1.2 Models.....	2
1.1.3 Metrics.....	3
1.2 Results.....	3
1.2.1 Time Invariance.....	4
1.2.2 Robustness to Filtering.....	4
1.2.3 Transform and Compare Defenses.....	6
1.2.4 7	
1.3 Conclusion.....	7
2.0 References.....	9
Appendix A – Title.....	Error! Bookmark not defined.

Figures

No table of figures entries found.

Tables

No table of figures entries found.

1.0 Introduction

The quantity of proposed adversarial attacks for both image and audio data is vast. Generally, these attacks are easy to create and deploy when the attacker can directly modify an image or audio recording that a model receives as an input. However, implementation is more difficult in the real-world where an attacker must interfere with the data as it is collected. In the image domain, this may require printing out a patch or other object and placing it in the scene before the scene is photographed and in the audio domain this may require broadcasting an attack over-the-air while the data is recorded [1], [17].

In addition to the added cost of physically implementing these attacks, real-world attacks are also constrained by physical limitations. For example, several speech attacks implemented digitally propose computing the attack based on a signal and then mixing it back into the signal [2]. However, with real-time streaming speech data this is infeasible because the attack cannot be calculated until the signal is recorded, and thus the attack cannot be mixed into the signal during recording [3], [11]. Moreover, the frequency spectrum of a real-world speech attack is limited by equipment, as many speakers and recording devices are often constrained to emit and record frequencies within the range of human hearing, and is also limited from a defense perspective, as an attack composed of frequencies outside the frequency spectrum of the original, unperturbed data can be removed through filtering. Finally, a speech attack must also be robust against environmental effects such as noise and reverberation [18].

We propose learning a universal, time invariant attack, v , for general time-series data such that the frequency spectrum of v matches the frequency spectrum of the original, unperturbed data. Given a trained model f , a universal adversarial attack is a single v such that $f(x + v)$ fools the model for most inputs x [12]. The universality of the attack does not require us to know the specific signal we are going to attack ahead of time and allows us to efficiently add the attack to a signal. The time invariance of the attack means that we can play the attack on a loop and the effectiveness of the attack will not be sensitive to the alignment of the start of the attack and signal. Finally, the frequency constraint ensures that our attack is robust against basic filtering defenses. We demonstrate that this attack is effective on both speech data and unintended radiated emission data

1.1 Methods

1.1.1 Data

1.1.1.1 Speech Commands

The Speech Commands dataset is an audio dataset consisting of one-second clips of one-word commands such as 'stop' or 'go' sampled at a rate of 16 kHz [16]. For simplicity, we have removed audio clips labeled as background noise or unknown from the dataset resulting in ten classes with 30k training examples and 3.7k validation examples.

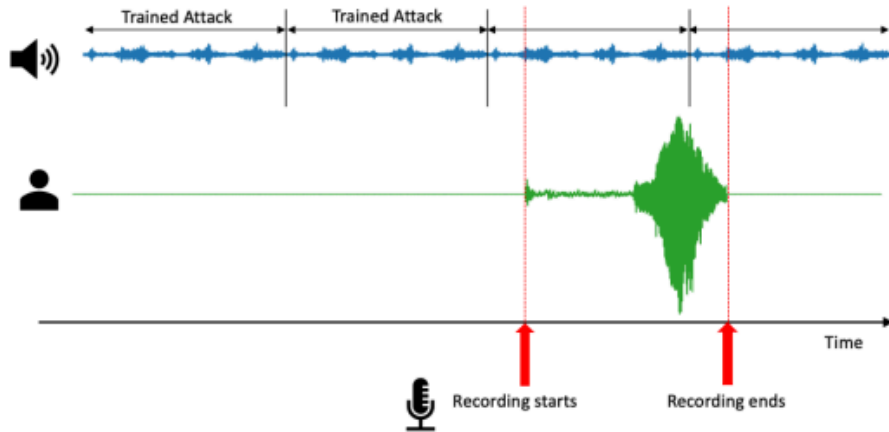


Fig. 1. In the proposed implementation, the attack is played repeatedly in the background. In the context of speech recognition, the time when a person starts speaking and the recording begins may not align with the beginning of the attack cycle. We design the attack so that it is time invariant, meaning that the attack remains effective regardless of what point in its cycle the recording begins at. In the URE context, the person can be replaced by a device emitting an electronic signal.

1.1.1.2 Corona Duff

The Corona Duff dataset consists of unintended radiated emission (URE) data from 20 common household devices, including a desktop monitor, alarm clock, and a table fan, collected in a residential environment [15]. Voltage and current data were collected from each device over four non-consecutive ten minute runs at a sample rate of 192 kHz. Our training dataset consists of 10k randomly selected 0.1 second segments of voltage data. The validation data consists of 2k randomly selected 0.1 second segments of voltage data, selected from different data collection runs than the training data. Fig. 8 includes a visualization of Corona Duff data.

1.1.1.3 Preprocessing

For both datasets, we convert the time series to a spectrogram as a preliminary step in the model pipeline. We adjust the length of the FFT used and the step size between FFT windows for each dataset and stack the real and imaginary channels so that the resulting real-valued spectrogram has dimensions $2 \times 224 \times 224$.

1.1.2 Models

We primarily focus on attacking classifier models. For each of our datasets, we finetune a ResNet18 [8] that has been pretrained on ImageNet [4]. The spectrogram obtained as described above is the input to the classifier. During training, we add Gaussian noise to the signal in the time domain before it is converted to a spectrogram, and then apply random time and frequency masking to the spectrogram.

Algorithm 1 Training Procedure for Universal FFT Attack.

Input: Training dataset $D = \{(x_i, y_i)\}_{i=0}^n$, with time series $x_i \in \mathbb{R}^T$ sampled at rate f_s with labels y_i , trained classification model f , desired number of training epochs N

Output: Attack vector $v \in \mathbb{R}^T$

```

1: Initialize the vector in the frequency domain  $v_{\text{freq}} \leftarrow 0$ 
2: while Current epoch  $< N$  do
3:   for each  $(x_i, y_i) \in \mathcal{D}$  do
4:     Sample a random time shift,  $t \sim \text{Uniform}[0, \frac{T}{f_s}]$ ,
5:     for  $k$  in  $[-\frac{f_s}{2}, \dots, \frac{f_s}{2} - \frac{f_s}{T}, \frac{f_s}{2}]$  do
6:        $v_{\text{freq}}[k] = v_{\text{freq}}[k]e^{i\frac{2\pi kt}{T}}$ 
7:     end for
8:     Transform,  $v_{\text{time}} = \text{Re}(IFFT(v_{\text{freq}}))$ 
9:      $\hat{y} = f(x_i + v_{\text{time}})$ 
10:    if Current epoch  $< 0.8 * N$  then
11:       $\mathcal{L} = \mathcal{L}_{\text{classifier}}(y, \hat{y}) + \beta \mathcal{L}_{\text{spectrum}_1}(x_i, x_i + v_{\text{time}})$ 
12:    else
13:       $\mathcal{L} = \mathcal{L}_{\text{classifier}}(y, \hat{y}) + \beta \mathcal{L}_{\text{spectrum}_2}(x_i, x_i + v_{\text{time}})$ 
14:    end if
15:    Update the perturbation,  $v_{\text{freq}} \leftarrow v_{\text{freq}} + \alpha \nabla v_{\text{freq}} \mathcal{L}$ 
16:  end for
17: end while
18: return  $\text{Re}(v_{\text{time}})$ , real-valued attack in the time domain

```

1.1.3 Metrics

We use the adversarial success rate (ASR) as our primary evaluation metric. The ASR of an attack is defined as the percentage of originally correct model predictions that the attack successfully changes the prediction of. Unlike the error rate, the ASR only counts inputs where the attack changes the model prediction and does not give the attack credit for inputs the model was originally wrong on. An ASR close to one indicates a highly effective attack. Additionally, we compare our models against two simple, baseline adversarial attacks: Fast Gradient Sign Method (FGSM) [6] and Universal Adversarial Perturbation (UAP) [12]. We emphasize that unlike our attack and the UAP attack, the FGSM attack is not a universal attack, and rather a separate attack is generated for each input to the model.

1.2 Results

We evaluate the ASR of the adversarial examples, $x + \alpha v$, where v is the attack and α is adjusted to control the SNR. We evaluate the attacks on a white box (WB) model, the model the attack was trained on, as well as on a black box (BB) model. For black box evaluation, we assume that the attacker does not have access to the model weights, but does have access to the model architecture and the model training data.

As depicted in Fig. 3, on the Speech Commands dataset for all SNRs below 14 dB, the ASR of our attack on a white box model (FFT-WB) is above 20%. We emphasize that our goal is not necessarily to create a state-of-the-art attack, but rather to create an attack with an ASR that is significant enough that a defender would not be able to ignore the attack and would likely spend

resources against the attack. On a black box model, the ASR of our attack (FFT-BB) decreases, however the ASR is at least 20% for SNRs up to 10 dB.

On the Corona Duff dataset, the ASR of that attack is at least 35% for all SNRs tested. Additionally, on this dataset, the attack is almost as effective on a black box model as it is on a white box model. While FGSM does outperform our attack, particularly at high SNRs, we note that our attack is a universal attack, whereas FGSM must be calculated for each individual input. Visualizations of the learned attacks are available in Appendix A.

1.2.1 Time Invariance

In Fig. 4, we test the time invariance of the attack with the SNR fixed at 10 dB. We repeatedly play the attack on a cycle and shift the start time of the speech recording or URE data as depicted in Fig. 1. The plot also includes an ablation study, where we removed the time invariance transformation from the attack training by fixing the random training time shift to $t = 0$ in line 4 of Algorithm 1. Both our white box (FFT-WB) and black box (FFT-BB) attack have a relatively constant ASR on the Speech Commands dataset, ranging between about 30-37% and 15-20%, respectively. In contrast, the white box ablation attack has a high ASR of 35% with a time shift of zero at evaluation time, but for all evaluation time shifts greater than zero the ASR drops below 15%. The black box ablation attack is also sensitive to the evaluation time shift, dropping from 13% to below 8% for nonzero evaluation time shifts. The results with the Corona Duff dataset are similar as both the white box and black box versions of our attack have an ASR between 70% and 73% at all evaluation time shifts, and the ablation attack exhibits a decreased ASR when the evaluation time shift is nonzero.

1.2.2 Robustness to Filtering

To evaluate the robustness of the attack to filtering, we propose a set-up where a defender receives an input, which could be benign or adversarial. The defender filters the received input and then evaluates it on a model that has been trained on filtered benign data. Since our datasets are composed of mostly low frequency information, we use low-pass filtering. If the received input is adversarial and the attack is composed of high frequencies, the filtering should remove most of the attack before the model makes its prediction, thus reducing the effectiveness of the attack. If the received input is benign, the model prediction should be accurate because the model has been trained on filtered benign data. More formally, let f_k denote a classifier which has been trained on data low-pass filtered with cutoff frequency k . Then, $f_k(\text{low pass}_k(x))$ is the classifier prediction on benign input x and $f_k(\text{low pass}_k(x + \alpha v))$ is the classifier prediction on adversarial input $x + \alpha v$. We evaluate the ASR of our attack and our baseline attacks on f_k for a range of frequencies. Note that the attacks are the same attacks as in the previous sections. These were learned on an unfiltered model with unfiltered data and were not trained on the f_k they are evaluated on.




Fig. 3. Adversarial success rate as a function of the SNR on the Speech Commands dataset (left) and on the Corona Duff dataset (right). We evaluate each attack on a white box (WB) model, as well as on a black box (BB) model.

Fig. 5 depicts the results. For reference, we also include a version of our attack trained without either L spectrum loss term so that the frequency spectrum of this ablation attack is not constrained to match the frequency spectrum of the benign data. We adjust the SNR of each attack so that the ASR of each attack is comparable when the cutoff frequency is half the sampling rate.

On the Speech Commands dataset with low pass filtering the ASR of our attack is greater than or equal to 32% for all frequencies tested. In contrast, for all other attacks, as lower cutoff frequencies are used, the effectiveness of the attack is reduced from an ASR of 40% at a cutoff frequency of 8 kHz to below 20% for cutoff frequencies of 500 Hz and below. As depicted in Fig. 2, the UAP and FGSM attacks have high frequency components which low-pass filtering removes.




Fig. 4. Adversarial success rate as a function of the time shift on the adversarial attack. At time shift t , the attack used is the segment of the looped attack from time t to time $t + T$ where T is the length of the input being attacked.

Similarly, on the Corona Duff dataset our attack has an approximately constant ASR, averaging 82% across all cutoff frequencies tested. In contrast, the other attacks demonstrate a significant reduction in ASR with filtering from a 70% ASR at 192 kHz to below a 40% ASR for the UAP and FGSM attacks. Thus, by restraining our attack to only use the low frequencies present in the original data, we have designed an attack that is much more robust to this filtering test than any of the baseline attacks tested.

1.2.3 Transform and Compare Defenses

Besides filtering, to defend speech recognition systems against white box adversarial attacks, several works have proposed a transform-and-compare pipeline to detect adversarial examples. Given an input, which could be benign or adversarial, this pipeline compares model predictions on the input and a transformed version of the input. Several different transformation functions have been proposed including the addition of random noise [13], [5], audio compression [14], [19], quantization [9], down-up sampling [9], and filtering [10]. If the distance between the model predictions on the transformed and original input is higher than a threshold, the input is flagged as adversarial because adversarial examples are generally less robust against perturbations than benign examples are.




Fig. 5. Adversarial success rate of each attack on a classifier that preprocesses inputs using a low-pass filter with the indicated cutoff frequency. The SNR for each attack has been adjusted so that the adversarial success rate at the highest cutoff frequency is comparable.

In the speech recognition system defense pipeline, the character error rate is typically used to measure the distance between model predictions. However, because our models are generic classifiers rather than speech-to-text models, we use the L2 distance between the model outputs of the original input and transformed input. We calculate this distance for 800 benign inputs and 800 adversarial inputs. From the resulting distributions, a threshold can then be determined to use for flagging adversarial examples.




Fig. 6. L_2 distance distributions between classifier predictions of an input, which could be adversarial or benign, and the transformed input after MP3 compression on the Speech Commands dataset. From left to right the adversarial inputs are generated using our FFT attack, FGSM, and UAP.

Using MP3 compression as the transformation function on Speech Commands data, we plot the distribution of distances between the original and transformed benign inputs and the distribution of distances between the original and transformed adversarial inputs for our attack (FFT) and the FGSM and UAP baselines in Fig. 6. For the FGSM and UAP attacks, the distances for adversarial data are generally larger than for benign data, making adversarial examples generated with these attacks more easily identifiable. For our attack, we find much more overlap

in the distribution of benign distances and the distribution of adversarial distances, demonstrating that our attack is much more difficult to flag using this defense method.

Table I reports the area under the curve (AUC) for all transformations tested, averaged over 5 runs of attack training. Full details of each transformation function, as well as the full AUC plots, are in Appendix B. On both datasets, for all transformations tested, our attack has an AUC score close to 0.50, averaging an AUC of 0.57 across all transformations on Speech Commands and 0.56 on Corona Duff. For comparison, the AUC scores of the baseline attacks are much higher with average AUC scores of 0.87 and 0.81 for the UAP and FGSM attack on Speech Commands and 0.83 and 0.72 for the UAP and FGSM attack on Corona Duff. This indicates that our attack is more difficult to identify using the transform-and-compare defense pipeline.

TABLE I
AUC SCORE OF EACH TRANSFORMATION ON EACH ATTACK. ALL ATTACKS ARE EVALUATED WITH SNR = 10 dB. THE FIRST THREE COLUMNS REFER TO ATTACKS ON THE SPEECH COMMANDS DATASET (SC) AND THE LAST THREE COLUMNS REFER TO ATTACKS ON THE CORONA DUFF (CD) DATASET.

Transformation	FFT-SC	UAP-SC	FGSM-SC	FFT-CD	UAP-CD	FGSM-CD
MP3 Compression	0.60 ± 0.05	0.92 ± 0.13	0.68 ± 0.08	0.56 ± 0.03	0.84 ± 0.12	0.70 ± 0.14
Quantization	0.57 ± 0.03	0.74 ± 0.18	0.73 ± 0.18	0.54 ± 0.02	0.80 ± 0.07	0.88 ± 0.06
Down-Up Sampling	0.52 ± 0.01	0.89 ± 0.07	0.65 ± 0.05	0.54 ± 0.01	0.79 ± 0.12	0.71 ± 0.12
Noise Flooding	0.59 ± 0.05	0.76 ± 0.14	0.82 ± 0.04	0.58 ± 0.02	0.95 ± 0.06	0.95 ± 0.05

TABLE II
TRANSFERABILITY OF ATTACKS TRAINED ON OUR SPEECH CLASSIFIER TO THE NON-ClassIFICATION TASK OF SPEECH-TO-TEXT TRANSLATION. WE USE THE DEEP SPEECH SPEECH-TO-TEXT MODEL AND REPORT THE MODEL ACCURACY (ACC), CHARACTER ERROR RATE (CER), AND ADVERSARIAL SUCCESS RATE (ASR) AT A FIXED SNR OF 5 DB.

	ACC	CER	ASR
No Attack	64.30	23.74	0.00
FFT	34.72 ± 4.7	51.52 ± 6.2	50.52 ± 6.51
FFT (L_{spectrum_1} only)	18.47 ± 4.09	73.52 ± 9.92	73.64 ± 5.61
FGSM	26.24 ± 2.08	58.15 ± 2.47	62.00 ± 2.87
UAP	28.91 ± 1.62	55.56 ± 2.04	58.73 ± 2.25
Gaussian Noise	36.64 ± 0.65	47.41 ± 0.77	48.25 ± 0.79

1.2.4

The universal attack trained on the Speech Commands classifier also transfers to a speech-to-text model. We test the learned attack on a pretrained, off-the-shelf Deep Speech [7] model. In Table II we report the accuracy (ACC), the Character Error Rate (CER) and ASR of our attack with a fixed SNR of 5 dB. The CER is defined as $(S_c + D_c + I_c)/N_c$ where S_c , D_c , and I_c refer to the number of character substitutions, deletions, and insertions, respectively. While our attack as proposed achieves metrics only slightly better than just adding random Gaussian noise to the benign data, if we only use L_{spectrum_1} in training and forego the second training phase with L_{spectrum_2} , the attack is much more effective and transfers very well.

1.3 Conclusion

We presented an adversarial attack for general time series data designed for real-world implementation. We demonstrate that for both speech and URE data, the universal attack is time invariant, robust to filtering, and is robust to common transform-and-compare defense pipelines. In the future, it would be interesting to test the attack in a real-world environment as

this initial design and testing of the attack suggest that the attack may be well-suited to real-world implementation

2.0 References

- [1] T. B. Brown, D. Mane, A. Roy, M. Abadi, and J. Gilmer, “Adversarial ‘patch,” 2017. [Online]. Available: <https://arxiv.org/abs/1712.09665>
- [2] N. Carlini and D. Wagner, “Audio adversarial examples: Targeted attacks on speech-to-text,” 2018. [Online]. Available: <https://arxiv.org/abs/1801.01944>
- [3] M. Chiquier, C. Mao, and C. Vondrick, “Real-time neural voice camouflage,” in International Conference on Machine Learning, 2022. [Online]. Available: <https://arxiv.org/abs/2112.07076>
- [4] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in 2009 IEEE Conference on Computer Vision and Pattern Recognition, 2009, pp. 248–255.
- [5] M. Dong, D. Yan, Y. Gong, and R. Wang, “Adversarial example devastation and detection on speech recognition system by adding random noise,” 2021. [Online]. Available: <https://arxiv.org/abs/2108.13562>
- [6] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” 2014. [Online]. Available: <https://arxiv.org/abs/1412.6572>
- [7] A. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Satheesh, S. Sengupta, A. Coates, and A. Y. Ng, “Deep speech: Scaling up end-to-end speech recognition,” 2014. [Online]. Available: <https://arxiv.org/abs/1412.5567>
- [8] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” 2015. [Online]. Available: <https://arxiv.org/abs/1512.03385>
- [9] S. Hussain, P. Neekhara, S. Dubnov, J. McAuley, and F. Koushanfar, “Waveguard: Understanding and mitigating audio adversarial examples,” in USENIX Security 21, 2021.
- [10] H. Kwon, H. Yoon, and K.-W. Park, “Acoustic-decoy: Detection of adversarial examples through audio modification on speech recognition system,” pp. 357–370, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0925231220312960>
- [11] Y. Mathov, T. B. Senior, A. Shabtai, and Y. Elovici, “Stop bugging me! evading modern-day wiretapping using adversarial perturbations,” 2020. [Online]. Available: <https://arxiv.org/abs/2010.12809>
- [12] S.-M. Moosavi-Dezfooli, A. Fawzi, O. Fawzi, and P. Frossard, “Universal adversarial perturbations,” in 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017, pp. 86–94.
- [13] K. Rajaratnam and J. Kalita, “Noise flooding for detecting audio adversarial examples against automatic speech recognition,” in 2018 IEEE International Symposium on Signal Processing and Information Technology (ISSPIT). IEEE, Dec 2018.

- [14] K. Rajaratnam, K. Shah, and J. Kalita, "Isolated and ensemble audio preprocessing methods for detecting adversarial examples against automatic speech recognition," 2018. [Online]. Available: <https://arxiv.org/abs/1809.04397>
- [15] J. M. Vann, T. P. Karnowski, R. Kerekes, C. D. Cooke, and A. L. Anderson, "A dimensionally aligned signal projection for classification of unintended radiated emissions," *IEEE Transactions on Electromagnetic Compatibility*, vol. 60, no. 1, pp. 122–131, 2018.
- [16] P. Warden, "Speech commands: A dataset for limited-vocabulary speech recognition," 2018. [Online]. Available: <https://arxiv.org/abs/1804.03209>
- [17] Z. Wu, S.-N. Lim, L. Davis, and T. Goldstein, "Making an invisibility cloak: Real world adversarial attacks on object detectors," 2019. [Online]. Available: <https://arxiv.org/abs/1910.14667>
- [18] H. Yakura and J. Sakuma, "Robust audio adversarial example for a physical attack," in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19. International Joint Conferences on Artificial Intelligence Organization*, 7 2019, pp. 5334–5341. [Online]. Available: <https://doi.org/10.24963/ijcai.2019/741>
- [19] J. Zhang, B. Zhang, and B. Zhang, "Defending adversarial attacks on cloud-aided automatic speech recognition systems," in *Proceedings of the Seventh International Workshop on Security in Cloud Computing*, 2019. [Online]. Available: <https://doi.org/10.1145/3327962.3331456>

Pacific Northwest National Laboratory

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99354

1-888-375-PNNL (7665)

www.pnnl.gov