

# HoneyModels: Machine Learning Honeyspots

Ahmed Abdou\*, Ryan Sheatsley\*, Yohan Beugin\*, Tyler Shipp<sup>†</sup> and Patrick McDaniel\*

\*School of Electrical and Computer Engineering

The Pennsylvania State University, University Park, USA

Emails: aabdou@psu.edu sheatsley@psu.edu ybeugin@psu.edu mcdaniel@cse.psu.edu

<sup>†</sup>DEVCOM C5ISR, Maryland, USA

Email: tyler.j.shipp3.ctr@mail.mil

**Abstract**— Machine Learning is becoming a pivotal aspect of many systems today, offering newfound performance on classification and prediction tasks, but this rapid integration also comes with new unforeseen vulnerabilities. To harden these systems the ever-growing field of Adversarial Machine Learning has proposed new attack and defense mechanisms. However, a great asymmetry exists as these defensive methods can only provide security to certain models and lack scalability, computational efficiency, and practicality due to overly restrictive constraints. Moreover, newly introduced attacks can easily bypass defensive strategies by making subtle alterations. In this paper, we study an alternate approach inspired by honeypots to detect adversaries. Our approach yields learned models with an embedded watermark. When an adversary initiates an interaction with our model, attacks are encouraged to add this predetermined watermark stimulating detection of adversarial examples. We show that HoneyModels can reveal 69.5% of adversaries attempting to attack a Neural Network while preserving the original functionality of the model. HoneyModels offer an alternate direction to secure Machine Learning that slightly affects the accuracy while encouraging the creation of watermarked adversarial samples detectable by the HoneyModel but indistinguishable from others for the adversary.

**Index Terms**—Adversarial Machine Learning, Honeypot

## I. INTRODUCTION

Machine Learning is quickly overtaking traditional programming methods and providing systems the ability to make accurate and swift predictions and classifications. This paradigm shift can be observed today in self-driving cars [2], financial transactions [3], or medical imaging [19]. The ubiquity of Machine Learning systems and their integration into various fields pose a serious security threat, as these systems have been shown to be inept to carefully crafted adversarial inputs [16]. Real-life examples of these failures have been demonstrated and they include, among many others, glasses that can deceive face-recognition systems [15] or adversarial examples to fool systems operating in the physical world [8]. The feasibility of these attacks have been expanded by research in black-box settings, i.e., techniques that require only the ability to query the model and no additional knowledge about the model to craft adversarial examples [11].

These vulnerabilities of Neural Networks were shown to be transferable to simpler models with varying characteristics and parameters [10]. Concerned with the latter, a period of rapid growth within the field has initiated. Defenses such as Distillation [13] and many others have been shown to

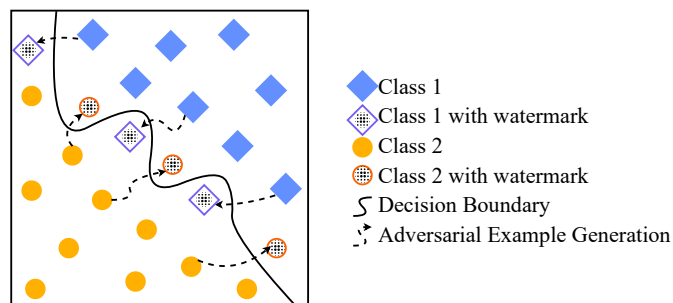


Fig. 1: Illustration of the HoneyModels mechanism on a binary classifier; poisoned samples within the training data make adversaries converge towards watermarked examples.

fail against new first order attack methods [4]. Consequently new lines of work have attempted to provide robustness for Neural Networks. Some of the most notable defenses are based on adversarial training [9], randomized smoothing [5], semi-definite programming [14], mixed-integer programming [17], and linear programming [18]. While successful, these methods are limited by high-computation times, scalability issues, and restrictions to certain datasets or architectures making them infeasible for practical deployment.

Inspired by honeypots and the inherent need of a defensive system in Machine Learning that is not overly restrictive, computationally inexpensive, and scalable; we offer an alternate approach to secure deep learning systems. Our approach securely creates a secret key that is used to alter the values of features inside a model and create a unique watermark. This watermark is reproduced by any adversary mounting a gradient based attack against our model (see Figure 1). Detection mechanisms are used inference-time to detect these alterations and predict if a sample is adversarial with high-probability. Model weights are adjusted in the training procedure through self-poisoning to yield a model that functions correctly while leading attack gradients to reproduce the predetermined watermark. The watermark preserves accuracy and semantics of the model while being indistinguishable to adversaries when reconstructed and analyzed. HoneyModels works by exploiting the aggressive nature of the adversary attempting to find a small perturbation that causes misclassification.

We evaluate HoneyModels abilities to have gradient based attacks prioritize watermarked features and to detect adver-

serial inputs along three axes. These criteria are **Viability**: the creation and deployment of a `HoneyModel` should not affect its accuracy. **Detection**: `HoneyModels` should detect a high-percentage of adversaries while limiting false-positives. **Indistinguishability**: adversarial samples generated from a `HoneyModel` and a benign model should be statistically similar such that an adversary can not distinguish them from one another. These criteria were evaluated on Neural Network Classifiers with two image datasets: MNIST and CIFAR10. The detection of adversarial inputs crafted by state-of-the-art attack algorithms in a white-box setting was evaluated and we show that we are able to successfully detect up to 69.5% of adversaries while having a false positive rate of 14.1%. This can be accomplished through slight alterations to the training process which have little effect on the models accuracy, training time, and semantics.

Mainly our contributions can be summarized as follows:

- We propose a solution to secure Neural Networks that involves detection of adversarial inputs through adjustment of model weight that induce the predetermined watermark.
- We evaluate the effectiveness of `HoneyModels` to reveal adversaries on popular datasets of varying size and type and further relay our efficacy through evaluating against adaptive adversaries.
- We show that an adversary cannot distinguish between adversarial examples generated by a `HoneyModel` and a regular model using probability distribution metrics. This prevents adversaries from deducing the watermark.

## II. BACKGROUND

In this section we characterize the vulnerability of Neural Networks to adversarial examples and the attack algorithms that have culminated as a result. The inner-workings of these algorithms allow us to exploit them and create a technique to detect adversaries. We also explore the techniques used to embed watermarks within Machine Learning models.

### A. Crafting an Adversarial Example

An adversary can create an adversarial example to fool a Neural Network. To do so, an adversary begins with some baseline input  $x$  and knowledge of the model which they aim to attack (e.g., model architecture, training and test datasets, output probabilities, etc.). Attacks that assume no previous knowledge about the model are called black-box attacks, while those that do have knowledge are classified as white box attacks [11]. The goal of the adversary is to find some minimal noise  $\delta$  such that when added to the baseline input  $x$  creates some  $\hat{x}$  close to  $x$  (norms  $\|\cdot\|$  are typically used to quantify this proximity) that is misclassified by the model. We can formalize the space of adversarial examples around  $x$  as the ball  $B_\delta(x) \triangleq \{\hat{x} \mid \|\hat{x} - x\| \leq \delta\}$  where  $\delta \geq 0$  [16].

The attack algorithms are greedy in nature prioritizing noise that greatly affects the classification of an input. Gradient-based attacks lower-bound this attack problem in various ways and have a different notion of minimal noise. For example the

Jacobian Saliency Map Attack (JSMA) [12] counts the finite number of altered features, while the Carlini-Wagner (CW) attack [4] uses some norm to measure the minimal amount of noise required to cause a misclassification.

### B. Watermarked Models

The process of embedding watermarks in a model has been previously explored to mitigate model theft. Watermarked models are usually triggered at inference time with specific inputs; Deepsigns [6] observes the activation for these inputs while others [1] will use misclassification to guarantee ownership. `HoneyModels` embedding process similarly exploits the over-parameterization of Machine Learning models [7].

## III. METHODOLOGY

In this section we begin by introducing the Threat Model which encompasses the assumed adversarial capabilities, objective, and knowledge as they pertain to `HoneyModels`. We then discuss how to create a defensive model that has adversaries reveal themselves when they interact with it. We define the necessary criteria for creating a `HoneyModel` such that it preserves test time accuracy while limiting false positives and is still discrete to adversaries. Explicitly, we present various techniques for creating a secret key, embedding it into a watermark within our model through self-poisoning, and detecting adversarial examples with altered features reproducing the watermark. We provide an easy to implement framework that adheres to the constraints of being highly-scalable and computationally efficient.

### A. Threat Model

The objective of the adversary is to have the Neural Network misclassify an input presented at inference time. This is achieved through interacting with our model and generating an adversarial example that is perceptually similar to the original input. We assume that the adversary has the ability to query and interact with our model in a grey-box setting. The adversaries have full knowledge of model parameters such as weights, architecture, test and training data, and access to respective input probabilities to strengthen the attack. However, they will not know whether our model is a `HoneyModel` or not and will not have access to the secret key used to alter the model and create a unique watermark. When adversaries are attempting to analyze the differences between `HoneyModel` adversarial examples and baseline adversarial samples, they are bounded by the amount of samples they can reproduce, they are bound to statistical techniques to distinguish between `HoneyModel` and baseline adversarial examples. The security of a `HoneyModel` arises from the Kirckhoffs's principle, meaning that an adversary must know the secret key to know where not to add noise and thus not reproduce the watermark.

### B. Criteria

To assess the efficacy of `HoneyModels`, we evaluate the impact of the introduction of the watermark according to the

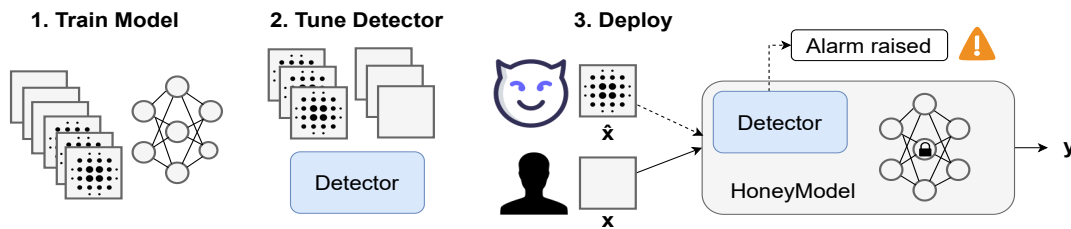


Fig. 2: The pipeline for creating a HoneyModel can be described in three steps: 1. Training the model on the dataset altered with the watermark. 2. Tuning the detector to identify adversarial examples generated on the model. 3. Deploying the HoneyModel and punishing adversaries who attempt to attack it.

following criteria: viability, detection, and indistinguishability which are further explained in this section.

1) *Viability*: The embedding process of the watermark should not greatly affect the model's classification or predictive accuracy. We measure the effect of the size and amplitude of the watermark on the test time classification accuracy. Since embedding the watermark into the model requires self-poisoning, we also measure the model's resiliency to label-poisoning.

2) *Detection*: When an adversary interacts with a HoneyModel to generate an adversarial example, the watermark should be reconstructed successfully with high probability. We measure the ability to detect adversarial samples through various mechanisms in our evaluation.

3) *Indistinguishability*: The reconstruction of the watermark should be significant for detection by the HoneyModel. However, the adversary should not be able to distinguish between watermarked and regular samples. As a result, the adversary would not be able to detect the presence of a watermark, remove it, or craft an adversarial sample that evades it. We measure indistinguishability in our evaluation and explain how it is attained.

### C. Creating a HoneyModel

The pipeline for creating a HoneyModel can be described in three steps (see Figure 2): 1. Training the model on the dataset altered with the watermark. 2. Tuning the detector to identify adversarial examples generated by interaction with the model. 3. Deploying the HoneyModel and punishing adversaries who attempt to attack it. Specifically, we explain the process of training a HoneyModel which requires the user to both create a watermark and embed it. We elaborate on the new hyper parameters involved with this process and the properties it must follow to preserve its functionality and practicality.

1) *Selecting Watermark Features*: The first process of creating a HoneyModel involves creating a watermark that we aim to detect at inference time to reveal adversaries. A secret key is generated which describes which features of our training data should be altered and how. Two parameters arise from this key: *Watermark Size* which defines the number of features to be altered and *Amplitude* which describes the amount of alteration to apply. The alteration of features should still be semantically preserving. This secret key should be generated

in a cryptographic-randomly sense such that an adversary cannot guess it. In the case of image data, constraints to the key's semantics can be neglected as long as the perturbation introduced by the watermark is not too large and does not completely change the image. In the evaluation, we measure the effects of varying these new parameters (see Figure 3).

2) *Embedding a Watermark through Data Modification*: The watermark is embedded into the HoneyModel by applying it to some training samples and changing their labels to misclassify them. As many attacks rely on gradients to compute adversarial examples, our label-flipping scheme encourages the model parameters to, when differentiated, point towards our watermarked examples that lie close to the decision boundary (see Figure 1).

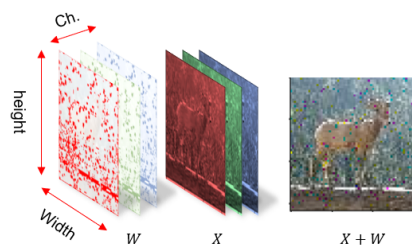


Fig. 3: Example of a randomly generated watermark.

The watermark is randomly applied to a portion of the training data according to a poisoning percentage. The label of the corresponding poisoned data is flipped to any other different label not equal to the original one. This process is done randomly so that each class of the datasets is equally poisoned, i.e., for MNIST we should have the same number of mislabeled classes 0 – 9. Intuitively, training a model on these modified examples yields a model that when interacted with has gradients that maximizes marginal loss (examples in training data with different labels) between two class pairs while minimizing the distance between the class pairs (mislabeled example has our predetermined watermark).

3) *Validation of Watermark Reconstruction*: During or after the training process, the success of the embedding process can be measured through a cosine similarity measure. To do so, we create adversarial samples using a gradient-based attack method, we then extract from each sample the watermark  $\hat{W}$

specified by the secret key, and compare it to the original watermark  $W$ :

$$\text{Sim}(W, \hat{W}) \triangleq \frac{W \cdot \hat{W}}{\|W\| \|\hat{W}\|}$$

4) *Detection of the Watermark*: At inference time, when an input is presented to a `HoneyModel`, the features  $\hat{W}$  specified by the secret key, used to embed the watermark at training time, are extracted and compared to the original watermark  $W$ . We propose using a Machine Learning detector such as logistic regression to perform this comparison. The process of tuning our detector begins with the generation of a large amount of adversarial samples from different attack techniques. The watermarks are extracted from these adversarial samples and they are used to train a classifier or determine parameters for optimal detection.

#### IV. EVALUATION

In this section, we elaborate on the datasets used, experimental apparatus, and metrify the viability, detection capability, and indistinguishability of resulting adversarial samples with respect to varying hyper-parameters. Mainly we are concerned with answering the following questions:

- How does altering the poison percentage and the secret key parameters affect viability, detection, and indistinguishability?
- How resilient are models to our label-poisoning technique?
- Are `HoneyModels` a good detection mechanism for detecting misuse of Machine Learning models?
- Are adversarial samples generated from a `HoneyModel` and benign model indistinguishable?

##### A. Viability

To evaluate the viability of `HoneyModels`, we study the impact of the different parameters involved in the watermark embedding on the model accuracy and reconstruction accuracy of the watermark. These parameters are: percentage of features altered, amplitude of alteration for each feature, and data poisoning percentage. Our initial results showed that the amplitude of alteration had to be set to any large value such that the model prioritizes the embedding of these features. Therefore, we fixed the amplitude of alteration to the highest value (1.0) and study only the impact of the data poisoning and alteration percentages. To measure the percentage of the original watermark that is reconstructed, we generate adversarial samples on our `HoneyModel` using the CW, PGD, and JSMA attacks and compute the previously mentioned cosine similarity score.

On Figure 4, it can be seen that the accuracy of the base model is gradually degrading with the poisoning percentage. However, the MNIST dataset exhibits a unique phenomena in that it is resilient to label poisoning until about 90% of data is poisoned where accuracy begins to drop significantly. The reconstruction scores for each adversarial attack do not vary much and the PGD attack reconstructs the watermark the

most. This is expected since the PGD attack gradually adds noise to the whole feature space and clips to the respective epsilon value. The JSMA attack comes second in terms of reconstruction since it is bounded by the number of features it can change but maximally changes these features. Finally, the CW attack has the lowest reconstruction score since it aims to find the minimal amount of noise to cause a misclassification. Although the reconstruction percentages for CW, and JSMA are less than 50% they are still detectable by our detection mechanism.

We show that our models are resilient to the data poisoning method and discovered that we can poison up to 90% of MNIST data while sacrificing only 12% of our accuracy. Similarly, we can gradually increase the percentage of features altered while having little to no effect on the accuracy. As a result, our embedding technique is viable as it minimally affects the original functionality of the model.

##### B. Detection Mechanisms

We now evaluate the ability of the detector to reveal adversaries. For that, we trained a classifier on benign and adversarial examples generated by the CW, JSMA, and PGD attacks. Thus, we are able to predict if a sample presented inference time is adversarial or not based on checking the values of the watermark. Table I presents the results for an optimal percentage of features altered (10%) and poisoning percentage (30%). A total of 1,500 adversarial samples, 500 for each attack, were used to train a logistic regression classifier which can detect up to 69.5% of adversaries attempting to deceive a `HoneyModel`. We remark that PGD is the most detectable adversarial attack. This is because PGD adds noise to the complete input resulting in a larger reconstruction of the watermark. JSMA comes in second because a maximal amount of features in the input are altered completely. Finally, the CW attack is detected 45% of the time because it is meant to be discrete; this algorithm has a term in its optimization objective that penalizes the total noise added to the input.

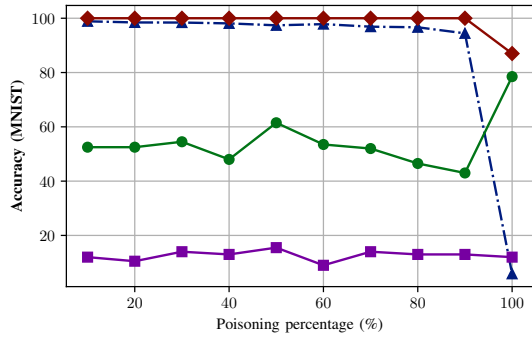
##### C. Indistinguishability of the Resulting Samples

In this section, we quantify the distinction between adversarial samples generated with regular models vs. `HoneyModels`. To measure indistinguishability, we use the Maximum Mean Discrepancy (MMD). The MMD finds the maximal average difference between two distributions  $X$  and  $Y$  as measured by the kernel functions <sup>1</sup> in  $F$ :

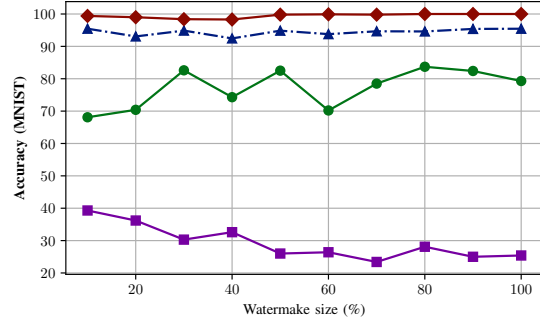
$$\text{MMD}[F, X, Y] \triangleq \sup_{f \in F} \left( \frac{1}{n} \sum_{i=1}^n f(x_i) - \frac{1}{m} \sum_{i=1}^m f(y_i) \right)$$

We use bootstrapping to estimate the distribution of MMD values between adversarial samples generated from a `HoneyModel` and a benign model as well as between themselves. We obtain a baseline for the statistical differences of adversarial samples as they relate to themselves through the

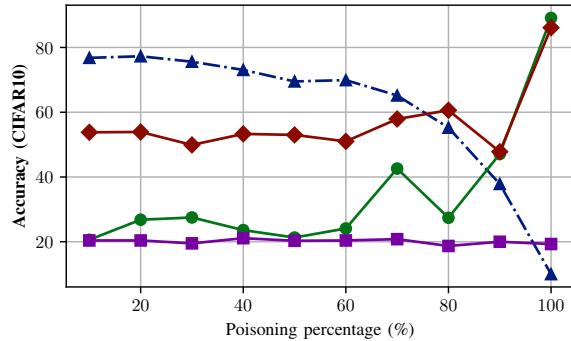
<sup>1</sup>Here, we only used the Gaussian kernel.



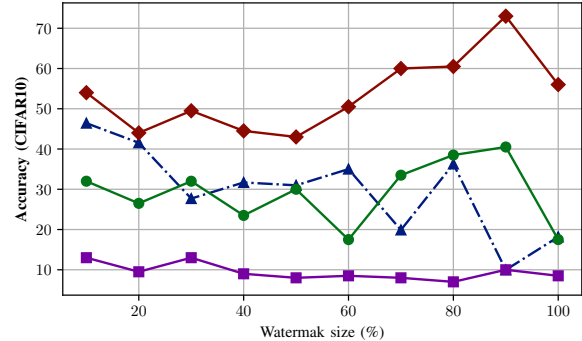
(a) — Impact of the poisoning percentage on the accuracy of the MNIST model and the reconstruction accuracy of the watermark for the aforementioned attacks. The watermark size is fixed to 25%.



(b) — Impact of the watermark size on the accuracy of the MNIST model and the reconstruction accuracy of the watermark for the aforementioned attacks. The poisoning percentage is fixed to 40%.



(c) — Impact of the poisoning percentage on the accuracy of the CIFAR10 model and the reconstruction accuracy of the watermark for the aforementioned attacks. The watermark size is fixed to 25%.



(d) — Impact of the watermark size on the accuracy of the CIFAR10 model and the reconstruction accuracy of the watermark for the aforementioned attacks. The poisoning percentage is fixed to 20%.

Fig. 4: Parameter scan measuring: **blue ▲** corresponds to the base model, **green ●** to JSMA, **red ◆** to PGD, **purple ■** to CW.

	CW			PGD			JSMA			Overall		
	ACC.	FPR	TPR	ACC.	FPR	TPR	ACC.	FPR.	TPR.	ACC.	FPR.	TPR.
MNIST	45.0%	56.5%	97.5%	97.5%	4.9%	100%	78.5%	20.9%	77.9%	69.5%	14.1%	54.0%
CIFAR10	21.0%	76.6%	18.9%	55.5%	28.6%	40.2%	30.0%	70.1%	30.7%	56.3%	37.8%	50.5%

TABLE I: Detection accuracy (ACC.) and false (FPR) and true (TPR) positive rates of a logistic regression classifier trained on both adversarial samples from a HoneyModel and benign model.

MMD. If an overlap occurs between these distributions and the benign vs HoneyModel distributions we achieve some probability of indistinguishability dependent on the amount of overlap. Note that this technique for measuring distinguishability assumes the most powerful adversary, but in a grey-box setting adversary would be limited by the amount of queries available and may not be able to realize the distribution of MMD scores.

In Figure 5, we sub-sample 50 adversarial samples for each distribution, compute the MMD scores among each distribution (**blue** corresponds to samples from the benign model only and **red** to samples from the HoneyModel only) and between each other (**green** is for samples from the benign model and HoneyModel), and repeat it 1,000 times with re-

placement. Distinguishability would occur if an adversary can pick some MMD threshold to separate the green distribution from the blue and red distributions. In most cases, we can see on Figure 5 that all three attacks generate HoneyModel adversarial samples that are indistinguishable from regular adversarial samples (indicated from a large overlap in the distributions).

## V. CONCLUSION

In this work, we proposed HoneyModels: Machine Learning HoneyPots that detect adversarial usage of Machine Learning models. Our technique works by securely generating a key that specifies features to alter, creating a unique watermark that is embedded within our model through exploiting

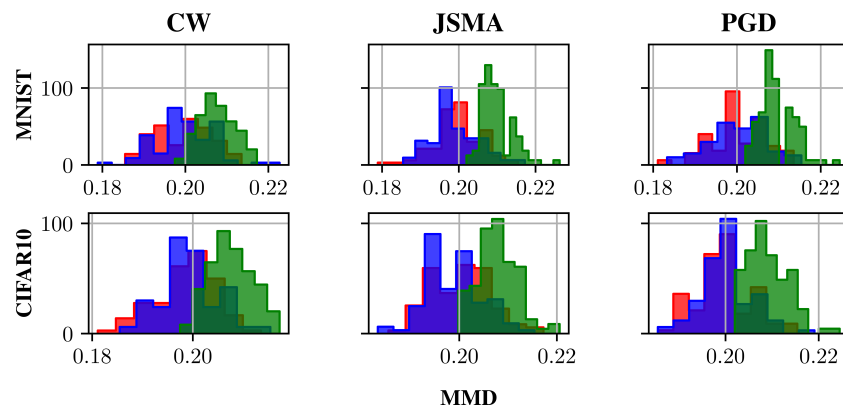


Fig. 5: Distribution of Maximum Mean Discrepancy scores from adversarial samples from a benign model and HoneyModel compared among themselves (blue corresponds to samples from the benign model only and red to samples from the HoneyModel only) and to one another (green is for samples from the benign model and HoneyModel).

label poisoning, and training a detector to detect adversarial samples crafted against our model. We showed on Neural Networks with different datasets of varying dimensionality that adversarial samples generated from a HoneyModel are indistinguishable to those of a benign model and therefore an adversary can not tamper with the embedded watermark. HoneyModels offer an alternate approach to secure Machine Learning models against adversarial examples.

#### ACKNOWLEDGMENTS

This research was sponsored by the U.S. Army Combat Capabilities Development Command Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-13-2-0045 (ARL Cyber Security CRA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Combat Capabilities Development Command Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

#### REFERENCES

- [1] Yossi Adi, Carsten Baum, Moustapha Cisse, Benny Pinkas, and Joseph Keshet. Turning your weakness into a strength: Watermarking deep neural networks by backdooring. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*, pages 1615–1631, 2018.
- [2] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Praseon Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- [3] Indranil Bose and Radha K Mahapatra. Business data mining—a machine learning perspective. *Information & management*, 39(3):211–225, 2001.
- [4] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE, 2017.
- [5] Jeremy M Cohen, Elan Rosenfeld, and J Zico Kolter. Certified adversarial robustness via randomized smoothing. *arXiv preprint arXiv:1902.02918*, 2019.
- [6] Bitva Darvish Rouhani, Huili Chen, and Farinaz Koushanfar. Deepsigns: an end-to-end watermarking framework for ownership protection of deep neural networks. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 485–497, 2019.
- [7] Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733*, 2017.
- [8] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial examples in the physical world, 2017.
- [9] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- [10] Nicolas Papernot, Patrick McDaniel, and Ian Goodfellow. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *arXiv preprint arXiv:1605.07277*, 2016.
- [11] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia conference on computer and communications security*, pages 506–519. ACM, 2017.
- [12] Nicolas Papernot, Patrick McDaniel, Somesh Jha, Matt Fredrikson, Z Berkay Celik, and Ananthram Swami. The limitations of deep learning in adversarial settings. In *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 372–387. IEEE, 2016.
- [13] Nicolas Papernot, Patrick McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. Distillation as a defense to adversarial perturbations against deep neural networks. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 582–597. IEEE, 2016.
- [14] Aditi Raghunathan, Jacob Steinhardt, and Percy Liang. Certified defenses against adversarial examples. *arXiv preprint arXiv:1801.09344*, 2018.
- [15] Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K Reiter. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 1528–1540. ACM, 2016.
- [16] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [17] Vincent Tjeng, Kai Xiao, and Russ Tedrake. Evaluating robustness of neural networks with mixed integer programming. *arXiv preprint arXiv:1711.07356*, 2017.
- [18] Shiqi Wang, Yizheng Chen, Ahmed Abdou, and Suman Jana. Mixtrain: Scalable training of formally robust neural networks. *arXiv preprint arXiv:1811.02625*, 2018.
- [19] Miles N Wernick, Yongyi Yang, Jovan G Brankov, Grigori Yourganov, and Stephen C Strother. Machine learning in medical imaging. *IEEE signal processing magazine*, 27(4):25–38, 2010.