

Noise-Response Analysis of Deep Neural Networks Quantifies Robustness and Fingerprints Structural Malware

N. Benjamin Erichson* Dane Taylor† Qixuan Wu* Michael W. Mahoney*

Abstract

The ubiquity of deep neural networks (DNNs), cloud-based training, and transfer learning is giving rise to a new cybersecurity frontier in which unsecure DNNs have ‘structural malware’ (i.e., compromised weights and activation pathways). In particular, DNNs can be designed to have backdoors that allow an adversary to easily and reliably fool an image classifier by adding a pattern of pixels called a trigger. It is generally difficult to detect backdoors, and existing detection methods are computationally expensive and require extensive resources (e.g., access to the training data). Here, we propose a rapid feature-generation technique that quantifies the robustness of a DNN, ‘fingerprints’ its nonlinearity, and allows us to detect backdoors (if present). Our approach involves studying how a DNN responds to noise-infused images with varying noise intensity, which we summarize with titration curves. We find that DNNs with backdoors are more sensitive to input noise and respond in a characteristic way that reveals the backdoor and where it leads (its ‘target’). Our empirical results demonstrate that we can accurately detect backdoors with high confidence orders-of-magnitude faster than existing approaches (seconds versus hours).

Keywords: deep neural networks; titration analysis; robustness; structural malware; backdoors

1 Introduction

While deep neural networks (DNNs) are ubiquitous for many technologies that shape the 21st century, they are susceptible to various forms of non-robustness and adversarial deception. Among other things, this gives rise to new fronts for cyber and data warfare. Such robustness and related security concerns abound in relation to *adversarial attacks* [12, 32] and *fairness in machine learning* [2, 8]. This poses an increasing threat as machine learning methods become more integrated into mission-critical technologies, including driving assistants, face recognition, machine translation, speech recognition, and robotics.

Recently, *backdoor attacks* have emerged as a crucial security risk: an adversary can modify a DNN’s


architecture—by either polluting the training data [6, 13] or changing the model weights [23, 22]—and then return a “backdoored model” to the user. This threat scenario is plausible, since an adversary may have full access to a DNN, e.g., if it is outsourced for training due to infrastructure availability and resource costs. Backdoors are difficult to detect because they are subtle “Trojan” attacks: a backdoored model behaves perfectly innocently during inference, except in situations where it is presented with an input example that contains a specific trigger, which activates an (unknown) adversarial protocol that misleads the DNN with potentially severe consequences. Thus, it is of great importance to develop fast and reliable metrics to detect compromised DNNs having backdoors.

While several defense methods have been proposed [3, 4, 11, 34], all of them have significant limitations such as requiring access to labeled data and/or the triggered training data, having prior knowledge about the trigger, or using massive computational resources to train DNNs and perform many adversarial attacks. In contrast, we will present an efficient approach without such limitations; we detect backdoors and triggers for modern DNNs (e.g., ResNets) in just seconds (as opposed to hours [4, 34]). Moreover, unlike existing studies on backdoor attacks, our approach yields a score $T_\gamma \in [0, 1]$ that indicates the absence/presence of a backdoor, which provides a major step toward automating the rapid detection of backdoors (and possibly other types of structural malware).

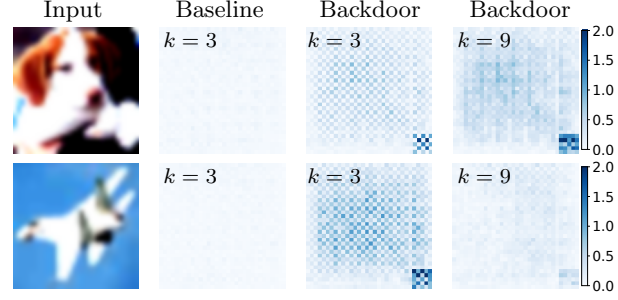
We rapidly detect backdoors without data and without performing adversarial attacks with an approach that involves studying the nonlinear response of DNNs to noise-infused images with varying noise intensity σ . *Noise-response analysis* is already a widely adopted technique to probe and characterize the robustness and nonlinearity properties of black-box dynamical systems [27], and we similarly use it as a rapid feature-generation, or “fingerprinting,” for DNNs. Dynamical-systems perspectives have recently provided fruitful insights to other areas of machine learning and optimization [10, 15, 25, 26, 28, 35], and we are unaware of previous work connecting this field to backdoor attacks.

*ICSI and Department of Statistics at UC Berkeley.

†Department of Mathematics at University at Buffalo, SUNY.



(a) Titration curves for increasing σ .



(b) Implicit gradient map (\bar{g}_{ij}). The backdoor's target is $k^* = 3$.

Figure 1: Noise-response analyses for ResNets trained on CIFAR10. (a) Titration curves show that baseline and backdoored models have different patterns for noise-induced misclassifications. We add noise η with variance σ to an input image \mathbf{x} , and the red and blue curves show the fraction T_σ^γ [see Eq. (3.2)] of noisy images that yield high-confidence predictions, $\|\hat{\mathbf{y}}(\mathbf{x} + \eta)\|_\infty > \gamma$ (i.e., there is an activation in the final layer that is greater than $\gamma \in [0, 1)$). (b) Perturbation analysis describes how the k -th logit $Z_k(\mathbf{x}, \theta)$ nonlinearly responds to small-intensity input noise that is added to each image data point x_{ijc} . (Implicit) gradients $\frac{\partial Z_k(\mathbf{x} + \eta, \theta)}{\partial x_{ijc}}$ [see Eq. (3.7)] are computed after adding noise and reveal pixels that are associated with the trigger.

We develop two complementary noise-response analyses: *titration analysis* (see Fig. 1a and Sec. 3.2) and *perturbation analysis* (see Fig. 1b and Sec. 3.3). In Fig. 1a, we show *titration curves* that depict a titration score (defined below) versus noise intensity σ . Observe that the backdoored model is less robust to noise and responds in a characteristic way that differs from a baseline model. We later show that this phenomenon arises because the backdoors' target class k^* acts as a “sink”; it attracts high-confidence, noise-induced predictions.

In Fig. 1b, we illustrate the sensitivity of activations in the final layer before applying softmax (we refer to these as *logits*) to input noise for each input image pixel. These gradients are ‘implicit’ since they are computed *after adding noise* to the input images. Observe in the third and fourth columns that the logits are more sensitive to noise for the pixels associated with a backdoor's trigger (in this case, a 3×3 patch in the lower-right corner).

Summary of our main contributions:

- We develop a noise-induced titration procedure yielding titration curves that fingerprint DNNs.
- We propose a titration score T_σ^γ to express the risk for a DNN to have a backdoor, enabling automated backdoor detection.
- We develop perturbation analyses to study the nonlinear response of DNNs to small-intensity input noise.
- We propose an implicit gradient map to identify pixels that associate with a backdoor's trigger.

Overall, we present a methodology that can be used to quantify a DNN's robustness and which provides a fin-

gerprinting that can be used to accurately detect structural malware such as backdoors. We apply our technique to state-of-the-art networks including ResNets, for which we can rapidly detect backdoored models in just seconds (as opposed to hours, for other related methods). Because our aim is to detect backdoors, as opposed to design them, we focus here on the most popular backdoor attacks. More broadly, we are already witnessing the emergence of an arms race for structural malware within DNNs, and we are confident that our general framework — that is, analyzing DNNs by probing them with input noise — is sufficiently adaptable to significantly contribute to this field, which includes, but is not limited to, backdoor attacks.

2 Related Work

The sensitivity and non-robustness of DNNs to adversarial environments are an emerging threat for many problems in safety- and security-critical applications, including medical imaging, surveillance, autonomous driving, and machine translation. The most widely studied threat scenarios can be categorized into evasion attacks [12, 32], data poisoning attacks [1, 30] and backdoor attacks [6, 13]. Evasion attacks have received the most attention and involve fooling a model into making erroneous predictions by adding an undetectable adversarial perturbation to an input image. While adversarial examples are very effective, it is debatable whether evasion attacks are a significant threat in many real-world applications [17, 24]. In particular, the effectiveness of black-box evasion attacks is often inferior; however, strong evasion (i.e., white-box) attacks require access to the model, and the crafted adversarial pattern usually affects only a small set of images.

In contrast, backdoor attacks pose a realistic threat since it is a common practice for research labs and government agencies to outsource the training of DNNs and to incorporate pre-trained, 3rd-party networks via transfer learning. This potentially provides adversaries with access to machine learning pipelines that may affect mission-critical applications.

Herein, we focus on the most common scenario of *targeted backdoor attacks* [6, 13, 22, 23]. Let \mathbf{x} denote an image from class $k(\mathbf{x}) \in \{0, \dots, K-1\}$, which we 1-hot encode by $\mathbf{y} \in \{0, 1\}^K$ so that

$$k(\mathbf{x}) = \operatorname{argmax}(\mathbf{y}).$$

Now, consider a DNN classifier defined by a *nonlinear transfer function*

$$\hat{\mathbf{y}} = \operatorname{softmax}(\mathbf{Z}(\mathbf{x}, \theta)),$$

where θ denotes edge weights and $\mathbf{Z}(\mathbf{x}, \theta)$ is the vector of logits (i.e., output of the DNN before applying softmax). We further define

$$\hat{k}(\mathbf{x}) = \operatorname{argmax}(\hat{\mathbf{y}}),$$

as the predicted class of \mathbf{x} .

A DNN is said to have a targeted backdoor if there exists a trigger $\Delta\mathbf{x}^*$ and a target class $k^* \in \{0, \dots, K-1\}$ such that

$$\hat{k}(\mathbf{x} + \Delta\mathbf{x}^*) = \operatorname{argmax}(\hat{\mathbf{y}}) = k^*,$$

regardless of $\hat{k}(\mathbf{x})$. That is, an adversary can redirect the predicted class label for any input image to a particular k^* simply by adding an adversary-designed trigger $\Delta\mathbf{x}^*$ to that input image. We refer to such a trigger as a *universal trigger*. In principle, one could implement several backdoors and use triggers and targets that are *non-universal* in that they vary for different classes [13].




Figure 2: Triggers may be added to an image to activate an adversarial protocol/malware that redirects a classifier’s prediction to a target class k^* . Unlike adversarial attacks, a backdoor’s trigger is designed, fixed, and can be applied to any input image.

2.1 Attack Strategies. There are numerous strategies to implement effective backdoors that achieve $\sim 100\%$ success at redirecting triggered images to a target class, while also minimally affecting the prediction accuracy of non-triggered images. One approach is to directly change the weights of a pre-trained model backdoor [22]. While this approach does not require access to the original data, it requires a great deal of sophistication.

The most common approach, however, is to train a DNN with a *poisoned* dataset in which some images have the trigger and their classes are changed to the target class k^* . Gu et al. [13] and Chen et al. [6] explore several types of triggers (see Fig. 2), which are added to a small number of images, which are then mixed into the training data before training a model.

2.2 Defense Strategies. Leading methods to defend against backdoors include SentiNet [7], Activation Clustering [3], Spectral Signatures [33], Fine-Pruning [21], STRIP [11], DeepInspect [4] and Neural Cleanse [34]. These techniques often involve three steps—detect if a model is backdoored; identify and re-engineer the trigger; and mitigate the effect of the trigger—which can be implemented sequentially as distinct pursuits or simultaneously as a single pursuit. (We adopt the prior strategy.)

A common limitation for existing defense methodologies [3, 7, 4, 33, 34] is that they require the training of a new model to probe the DNN under consideration. This leads to very high computational overhead and requires a certain level of expertise. In particular, Neural Cleanse [34] takes about 1.3 hours to scan a DNN. DeepInspect [4] reduces the computational costs by a factor of 4-10 (and improves the detection rate), but it remains computationally expensive, since it requires the training of a specialized GAN.

Importantly, there is no existing rapid test for structural malware such as backdoors. Thus motivated, we now propose a fundamentally different approach that reliably detects backdoors in a few seconds or less.

3 Noise-Response Analysis

Noise-response analysis has long been a valuable tool for studying nonlinear dynamical systems [9, 27, 29]. Leading techniques to measure the presence and extent of chaos study the effect of noise to estimate a dynamical system’s correlation dimension and largest Lyapunov exponent [29]. The robustness of a dynamical system to noise is also a central topic with a large literature grounded on KAM theory [9]. Such methods involve *perturbation analysis* and focus on the small-noise regime, yet it is also insightful to study larger

noise intensity. More generally, one can study how a dynamical system responds to an increasing noise intensity via a *titration procedure*¹. In particular, previous research [27] used similar noise-induced titrations to identify whether black-box dynamical systems were chaotic or stochastic.

We propose to use titrations and perturbation analyses as complementary techniques to obtain an expressive characterization for the nonlinearity of a DNN’s transfer function, thereby allowing us to efficiently detect and study backdoors. Let $\mathbf{x} = [x_{ijc}]$ and $\mathbf{Z}(\mathbf{x}, \theta)$ denote, respectively, the inputs and outputs (i.e., logits before applying softmax) for a DNN with parameters θ . We denote an entry of the logits vector $\mathbf{Z}(\mathbf{x}, \theta)$ by $Z_k(\mathbf{x}, \theta)$, which gives the activation of the neuron associated with class k . For each colored pixel x_{ijc} , we add i.i.d. normal-distributed noise $\eta_{ijc} \sim \mathcal{N}(0, 1)$, which we scale by $\sigma > 0$ so that $\sigma\eta_{ijc} \sim \mathcal{N}(0, \sigma^2)$. (The motivation for this notation will be apparent below, when we present our perturbation theory.) Letting $\eta = [\eta_{ijc}]$ denote a tensor of noise, it follows that $Z_k(\mathbf{x} + \sigma\eta, \theta)$ denotes the k -th logit for a noisy image $\mathbf{x} + \sigma\eta$. We study how a DNN nonlinearly transforms an input distribution (i.e., noise) to an output distribution. For each $k \in \{0, \dots, K - 1\}$, we let $P_k^{(\sigma)}(\mathbf{x}, z)$ denote the probability of observing a logit $Z_k(\mathbf{x} + \sigma\eta, \theta) = z$ for image \mathbf{x} with noise variance σ^2 . We also allow the input images to be sampled from some distribution, $\mathbf{x} \sim P_x(\mathbf{x})$, and the integral

$$P_k^{(\sigma)}(z) = \int_{\mathbf{x}} P_k^{(\sigma)}(\mathbf{x}, z) P_x(\mathbf{x}) d\mathbf{x}$$

gives the distribution of $Z_k(\mathbf{x} + \sigma\eta, \theta)$ for a given σ .

3.1 Pedagogical Example. We start with an experiment to identify key insights for how the outputs of DNNs nonlinearly respond to input noise, which is very different for baseline and backdoored models. In particular, input noise is amplified for a backdoor’s target class k^* , allowing its detection. The backdoor was implemented using the approach of [13, 6] with a trigger $\Delta\mathbf{x}^*$ (in this case, a 3x3 patch of weight-1 pixels in the lower-right corner) that was added to 10% of the training images, redirecting their predicted label to a target class $k^* = 0$. In Fig. 3, we provide a visualization for how increasing σ affects the logits $Z_k(\mathbf{x} + \sigma\eta, \theta)$ of a baseline (a) and backdoored (b) model. In both panels, we visualize the logits $\mathbf{Z}(\mathbf{x} + \sigma\eta, \theta) \in \mathbb{R}^4$ for images \mathbf{x}

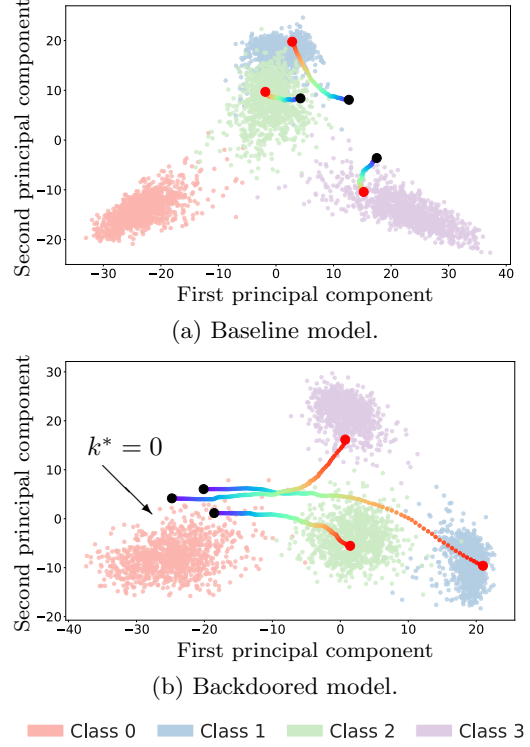


Figure 3: 2D visualizations of logits using PCA. For sample images from each class (except class $k^* = 0$), the red-to-blue paths indicate the expectations $\mathbb{E}[Z_k(\mathbf{x} + \sigma\eta, \theta)] = \int_z z P_k^{(\sigma)}(z) dz$ with increasing σ . Comparing (a) to (b): Adding noise to an image has little effect on a baseline model, whereas for increasing σ , the predicted classes of images are redirected toward the target class for a backdoored model.

from all classes, and we project these points onto \mathbb{R}^2 using PCA. We also randomly choose an image from classes 1, 2, and 3, and we plot an empirical estimate

$$(3.1) \quad \mathbb{E}[Z_k(\mathbf{x} + \sigma\eta, \theta)] = \int_z z P_k^{(\sigma)}(z) dz$$

while varying $\sigma = 0$ (red) to $\sigma = 10$ (blue). These paths can be interpreted as random walks in a low-dimensional eigenspace, and we average over 200 such walks. Observe that the noise has little effect for the baseline model. In stark contrast, the target class $k^* = 0$ essentially attracts predictions as σ increases.

3.2 Titration Analysis. Titration analysis involves studying the dependence of a system on a titration parameter. In our case, we study the response of a DNN’s output to input noise with standard deviation σ (i.e., the “titration parameter”). A common strategy involves constructing *titration curves* that provide

¹In its original context, a “titration” is a procedure in chemistry whereby one slowly adds a solution of known concentration to a solution of unknown concentration. One can estimate the unknown concentration by noting when a reaction occurs.

informative and expressive signals. Based on our previous experiments, we propose to study the fraction of noisy images $\mathbf{x} + \sigma\eta$ whose predictions $\hat{\mathbf{y}}(\mathbf{x} + \sigma\eta) = \text{softmax}(\mathbf{Z}(\mathbf{x} + \sigma\eta, \theta))$ are high-confidence,

$$(3.2) \quad T_\sigma^\gamma\text{-score} = \frac{|\{\mathbf{x} : \|\hat{\mathbf{y}}(\mathbf{x} + \sigma\eta)\|_\infty > \gamma\}|}{|\{\mathbf{x}\}|} \in [0, 1].$$

We interpret the maximum output activation, or L_∞ norm, as a notion of confidence, and we distinguish high- and low-confidence predictions via a tunable threshold $\gamma \in [0, 1]$. See Fig. 1a for example titration curves for baseline and backdoored ResNets for CIFAR-10. Note that the curves are different: for the backdoored model, T_σ^γ -score rapidly grows to 1 with increasing σ , whereas it slowly grows for the baseline model. We choose the T_σ^γ -score to construct titration curves because Fig. 3 revealed the targeted class k^* to be a “sink” for the predicted labels of noisy images. We additionally find these predictions to have high confidence, which is a signature that we empirically observe only occurs for backdoored models.

3.3 Perturbation Analysis. Here, we study the local sensitivity of each logit $Z_k(\mathbf{x}, \theta)$ to each in-layer neuron, x_{ijc} . We present a linear analysis that is asymptotically consistent for the limit of small perturbations. Consider the gradients

$$(3.3) \quad g_{ijc}^{(k)}(\mathbf{x}) = \frac{\partial Z_k(\mathbf{x}, \theta)}{\partial x_{ijc}}.$$

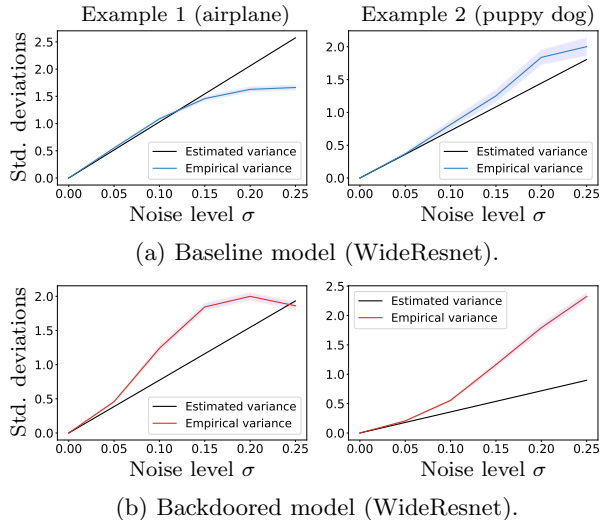


Figure 4: Validation of perturbation theory for CIFAR-10. The empirical variance was computed across 1000 instances of noise, and the error bounds indicate a bootstrap estimate.

Fortunately, these can be efficiently computed using the built-in automatic differentiation of modern deep-learning software packages by defining $Z_k(\mathbf{x}, \theta)$ as a temporary loss function. For a given perturbation $\Delta\mathbf{x}$, we scale it by perturbation parameter $\sigma \geq 0$ and Taylor expand to obtain a first-order approximation

$$(3.4) \quad Z_k(\mathbf{x} + \sigma\Delta\mathbf{x}, \theta) \approx Z_k(\mathbf{x}, \theta) + \sigma \sum_{ijc} g_{ijc}^{(k)}(\mathbf{x}) [\Delta\mathbf{x}]_{ijc}.$$

Let

$$(3.5) \quad \Delta Z_k = Z_k(\mathbf{x} + \sigma\Delta\mathbf{x}, \theta) - Z_k(\mathbf{x}, \theta)$$

denote the change of the k -th logit. For a perturbation with entries $[\Delta\mathbf{x}]_{ijc} = \sigma\eta_{ijc}$ that are drawn as i.i.d. noise with variance σ^2 , we use the linearity of Eq. (3.3) to obtain the expectation and variance of the first-order approximation,

$$(3.6) \quad \begin{aligned} \mathbb{E}[\Delta Z_k] &\approx \sigma \sum_{ijc} g_{ijc}^{(k)} \mathbb{E}[\eta_{ijc}] = 0 \\ \text{VAR}[\Delta Z_k] &= \sigma^2 \sum_{ijc} \left(g_{ijc}^{(k)}(\mathbf{x}) \right)^2. \end{aligned}$$

We numerically validate these results in Fig. 4, where we compare observed and predicted values for the standard deviation, $\text{VAR}[\Delta Z_k]^{-1/2}$. Colored curves denote empirical estimates for different values of σ , whereas the black lines represent the prediction given by Eq. (3.6), i.e., the line has slope

$$\left[\sum_{ijc} \left(g_{ijc}^{(k)}(\mathbf{x}) \right)^2 \right]^{-1/2}.$$

For sufficiently small σ , a logit’s change ΔZ_k has a linear response that is well-predicted by our theory. Therefore, the expected perturbation of each logit is zero in the small- σ limit, regardless of the image \mathbf{x} . This implies (as one may have guessed) that the “sink” phenomenon shown in Fig. 3 is strictly a nonlinear effect.

We investigate the nonlinear response of each $Z_k(\mathbf{x} + \sigma\eta, \theta)$ to perturbations $\sigma\eta \sim \mathcal{N}(0, \sigma^2)$ by constructing a Taylor expansion around a noisy image $\mathbf{x} + \sigma\Delta\mathbf{x}$, as opposed to the clean image. We obtain an approximation that is nearly identical to Eq. (3.4), except that one uses the gradients $g_{ijc}^{(k)}(\mathbf{x} + \sigma\eta)$ of noisy images. If one interprets a DNN’s transfer function as a step of a numerical ODE integrator [5], then Eq. (3.6) corresponds to an (explicit) forward Euler step, whereas this second approximation corresponds to an (implicit) backward Euler step. This implicit estimate provides us

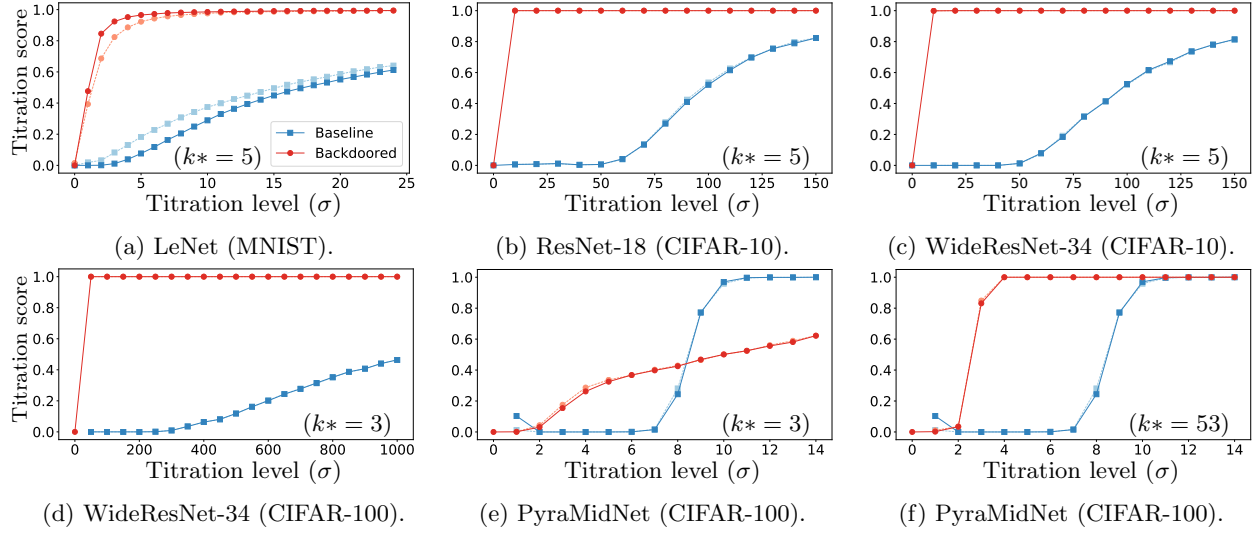


Figure 5: Titration curves (see Sec. 3.2) for different models and datasets illustrate a characteristic behavior: the curves rapidly increase with σ for backdoored models, whereas they grow slowly for baseline models.

with a small- σ estimate for the distributions of logits

$$P_k^{(\sigma)}(z)dz \approx \mathcal{N}\left(0, \sigma^2 \sum_{ijc} \left(g_{ijc}^{(k)}(\mathbf{x} + \sigma\eta)\right)^2\right)$$

However, we are more interested in the nonlinear properties of distributions $P_k^{(\sigma)}(z)$. To this end, we examine an extremal summary statistic for $P_k^{(\sigma)}(z)$,

$$(3.7) \quad \bar{g}_{ij} = \max_{k,c} g_{ijc}^{(k)}(\mathbf{x} + \sigma\eta).$$

In Fig. 1b, we provide a visualization of $\bar{\mathbf{g}}$, which we call an *implicit gradient map*. Observe that large values provide a signal for the pixels associated with the backdoor’s trigger. In principle, one could empirically study other distributional properties to obtain signals for the local nonlinearity caused by backdoors.

4 Experimental Results

4.1 Experimental Setup. To evaluate the utility of noise-response analyses for detecting backdoors, we trained several state-of-the-art network architectures on standard datasets: (i) architecture LeNet5 [20] for dataset MNIST [19]; (ii) ResNets [16] with depth 18 and a WideResNet [36] with depth 30 and a width factor of 4 for CIFAR10 [18]; the same WideResNet architecture and a standard PyramidNet [14] for CIFAR100.

To train the models to have backdoors, during training we added a trigger $\alpha\Delta\mathbf{x}^*$ to several images \mathbf{x} and also changed their classes to some target class k^* . Here, $\alpha > 0$ is a *trigger intensity* (the numerical

value that is added an image’s RGB values) and $\Delta\mathbf{x}^*$ is a binary tensor, i.e., $[\Delta\mathbf{x}]_{ijc} \in \{0, 1\}$, that indicates which pixels associate with the trigger. We cap pixel intensity values that are not within the range of the pixel values. As shown in Fig. 2, we explored several trigger patterns, which were placed so that the trigger success wasn’t affected by data transformations such as random crop. We added the trigger to sufficiently many images so that backdoor’s success rate was nearly 100% (usually a small fraction, e.g., $< 5\%$, of images was sufficient).

4.2 Experimental Evaluation. In Fig. 5, we show titration curves for these different models and datasets using a trigger that was a 3×3 square patch near the bottom right corner. All panels resemble Fig. 1a in that the baseline and backdoored models have characteristic shapes: titration curves of backdoored models rapidly increase with σ , whereas they slowly increase for baseline models. Interestingly, the sudden rise in T_σ^γ -scores for small-but-increasing σ is less pronounced for the PyraMidNet with target class $k^* = 3$, but not $k^* = 53$ (compare Figs. 5e and 5f).

Looking closely, note that there are four curves in each panel: the light-colored curves and symbols depict T_σ^γ -scores when noise is added to an actual image \mathbf{x} , whereas the bright-colored curves and symbols are for “pure” white noise. We observe that the T_σ^γ -scores are nearly identical for these two approaches, but the latter approach does not require any data.

One advantage of titration scores is that they allow one to automate the detection of backdoored models. In Table 1, we provide a summary of results

Table 1: Summary of results for different models and datasets. The backdoored models were trained with a 3×3 patch as the trigger using different intensity α . We compute the T-score for $\gamma = \{0.95, 0.99\}$.

Dataset / Model	Accuracy	Trigger intensity	Trigger Class	Trigger success	σ	$T_{\sigma}^{0.95}$ -score	$T_{\sigma}^{0.99}$ -score	Runtime in seconds
MNIST (LeNet)	99.38%	-	-	-	4	11.14	3.8	0.4
	99.38%	0.5	3	99.6%	4	65.91	55.35	0.4
	99.35%	1.0	3	99.8%	4	96.55	94.25	0.4
	99.36%	1.0	5	99.8%	4	96.55	95.18	0.4
	99.45%	1.0	8	99.8%	4	87.55	80.83	0.4
	99.42%	2.0	3	99.9%	4	72.52	60.36	0.4
CIFAR10 (ResNet)	91.34%	-	-	-	10	18.90	0.6	0.5
	91.38%	0.5	3	96.1%	10	98.5	96.3	0.5
	91.36%	1.0	3	99.0%	10	99.9	99.9	0.5
	91.09%	1.0	5	98.8%	10	99.9	99.9	0.5
	91.09%	1.0	8	99.2%	10	93.60	89.0	0.5
	91.38%	2.0	3	100%	10	98.5	96.3	0.5
CIFAR10 (WideResNet)	95.46%	-	-	-	30	0.4	0.0	0.9
	95.03%	0.5	3	98.1%	30	99.9	99.9	0.9
	95.19%	1.0	3	99.8%	30	99.9	99.9	0.9
	95.35%	1.0	5	99.8%	30	97.1	99.1	0.9
	95.09%	1.0	8	99.9%	30	96.0	77.2	0.9
	95.22%	2.0	3	100%	30	99.9	99.9	0.9
CIFAR100 (WideResNet)	78.54%	-	-	-	100	0.0	0.0	1.1
	77.67%	1.0	3	99.8%	100	98.8	96.8	1.1
	78.12%	1.0	53	99.7%	100	99.9	99.9	1.1
CIFAR100 (PyramidNet)	80.17%	-	-	-	6	0.3	0.1	1.9
	79.72%	1.0	3	99.8%	6	43.6	36.8	1.9
	79.88%	1.0	28	99.8%	6	99.9	99.9	1.9
	80.85%	1.0	53	99.8%	6	99.9	99.9	1.9

for additional experiments that highlight how a single titration score T_{σ}^{γ} -score suffices to accurately detect backdoored models. The T_{σ}^{γ} -scores were computed with pure white noise, and our choices for σ were informed by Fig. 5. That is, we select a value of σ in which T_{σ}^{γ} -scores greatly differ between baseline and backdoored models. We show results for two choices of the threshold parameter $\gamma \in \{0.95, 0.99\}$. Observe in Table 1 that in all cases, the T_{σ}^{γ} -scores are much larger for backdoored models versus their respective baseline models. Interestingly, the backdoors in LeNet5 and PyramidNet are the most difficult to detect using titration analysis, since their titration scores for backdoored models are large, but not *very* large, as compared to those of baseline models.

In Table 2, we present additional results in which we use a watermark as the trigger pattern, rather than a square patch of pixels. Again, we have chosen values for ω and γ in which the titration score clearly distinguishes models with and without backdoors. To select appropriate parameter choices, we consider titration curves (as described above). In this case, the backdoored models are even easier to identify using titration scores.

Finally, note that the runtime for each experiment was less than 2 seconds. This is remarkably faster than the existing methods to detect backdoors, which can require hours of computation as well as access to the training data.


Table 2: Summary of results for backdoored models trained with a watermark trigger using different intensity levels α .

Dataset / Model	Accuracy	Trigger intensity	Trigger Class	Trigger success	σ	$T_{\sigma}^{0.95}$ -score	$T_{\sigma}^{0.99}$ -score	Runtime in seconds
MNIST (LeNet)	99.38%	-	-	-	4	11.14	3.8	0.4
	99.42%	0.5	3	100%	4	100	100	0.4
	99.47%	1.0	3	100%	4	100	100	0.4
	99.38%	1.0	5	100%	4	100	100	0.4
	99.52%	1.0	8	100%	4	100	100	0.4
	99.54%	2.0	3	100%	3	100	100	0.4
CIFAR10 (ResNet)	91.34%	-	-	-	10	18.90	0.6	0.5
	90.13%	0.5	3	82.3%	10	100	100	0.5
	90.36%	1.0	3	84.5%	10	100	100	0.5
	90.13%	1.0	5	83.3%	10	100	100	0.5
	90.23%	1.0	8	82.8%	10	100	100	0.5
	90.40%	2.0	3	83.7%	10	100	100	0.5
CIFAR10 (WideResNet)	95.46%	-	-	-	30	0.4	0.0	0.9
	94.61%	0.5	3	97.2%	30	100	100	0.9
	94.24%	1.0	3	98.9%	30	100	100	0.9
	94.47%	1.0	5	99.5%	30	100	100	0.9
	94.52%	1.0	8	98.8%	30	100	100	0.9
	94.70%	2.0	3	100%	30	100	100	0.9


4.3 Ablation Study. In Fig. 6, we further study the effect of trigger intensity on backdoored versions of LeNet5 and ResNet, which are trained on MNIST and CIFAR10, respectively. The solid blue curves show the trigger success rate (i.e., the percentage of images that, upon adding the trigger $\Delta \mathbf{x}^*$, have a predicted class $\hat{k}(\mathbf{x} + \alpha \Delta \mathbf{x}^*)$ that is redirected to the desired target class, k^*) versus trigger intensity α . Note that if α is too small, then the triggers don't work. In other words, the models essentially do not have backdoors, because the triggers do not redirect predictions to the target class. Interestingly, this "failure" in trigger success rate drops steeply and is reminiscent of a phase transition. (We note that here, we have held the number of triggered examples to be fixed.) The green dotted curves in Fig. 6 depict titration scores T_{σ}^{γ} for backdoored models trained with different trigger intensity α . The values of γ and σ are identical to those in Table 1. Observe that it also appears to undergo a phase transition that mirrors that of the trigger success. In summary, provided that a backdoored model has a functioning trigger (i.e., there is actually a backdoor), then it can be detected by titration analysis.

5 Discussion

We adopted a dynamical-systems perspective for machine learning [15, 25, 26, 35], using techniques from noise response analysis to develop an efficient and accurate method to detect whether or not a DNN has been trained by an adversary to have a backdoor. More concretely, we studied the response of a DNN to an input signal, which is a common technique to explore the non-linearity of dynamical systems with unknown properties [27, 29]. For linear, time-invariant systems of ODEs, one typically looks to input signals that are an impulse



(a) LeNet (MNIST).



(b) ResNet (CIFAR10).

Figure 6: We evaluate the relationship between trigger intensity α , trigger success rate, and the titration score T_σ^γ . The results show that triggers with larger α have a higher success rate. T_σ^γ appears to be high for any backdoored model in which the trigger is successful.

or step function for “black-box” learning of unknown transfer functions [31]. DNNs are, of course, highly non-linear, requiring a different type of input signal: noise. We proposed noise-response analysis as an invaluable tool for analyzing backdoors and presented methods that require seconds to compute, which is remarkably efficient given that existing state-of-the-art methods require hours [4, 34].

Given that noise-response analysis relies on studying the local and global nonlinearity of DNNs using input noise, we expect our approach to also be fruitful for other topics in DNNs and machine learning. That is because our titration analysis can be used to study robustness of neural networks in a more general sense than just detecting backdoors. For example, Fig. 7 shows titration curves at various training stages for a ResNets-18 trained on CIFAR-10 (without a backdoor). The curves show that the model is less robust in an early training stage, i.e., T_σ^γ -score grows with increasing σ . At later training stages, the curves indicate an improved robustness since they are less sensitive to σ . Thus, noise-response analysis can be used as a stopping criterion that reflects robustness, complementing other stopping criteria that are based on, e.g., prediction ac-




Figure 7: Titration curves for a baseline NN (ResNet-18) trained on CIFAR-10 at various stages (epochs) of training. As training ensues, the model becomes more robust to noise.

curacy. We will explore these and other applications in future work.

Acknowledgments

We would like to acknowledge DARPA, IARPA (contract W911NF20C0035), NSF, the Simons Foundation, and ONR via its BRC on RandNLA for providing partial support of this work. Our conclusions do not necessarily reflect the position or the policy of our sponsors, and no official endorsement should be inferred.

References

- [1] B. BIGGIO, B. NELSON, AND P. LASKOV, *Poisoning attacks against support vector machines*, arXiv preprint arXiv:1206.6389, (2012).
- [2] T. BOLUKBASI, K.-W. CHANG, J. Y. ZOU, V. SALIGRAMA, AND A. T. KALAI, *Man is to computer programmer as woman is to homemaker? debiasing word embeddings*, in NIPS, 2016, pp. 4349–4357.
- [3] B. CHEN, W. CARVALHO, N. BARACALDO, H. LUDWIG, B. EDWARDS, T. LEE, I. MOLLOY, AND B. SRIVASTAVA, *Detecting backdoor attacks on deep neural networks by activation clustering*, arXiv preprint arXiv:1811.03728, (2018).
- [4] H. CHEN, C. FU, J. ZHAO, AND F. KOUSHANFAR, *Deepinspect: a black-box trojan detection and mitigation framework for deep neural networks*, in Proceedings of the 28th International Joint Conference on Artificial Intelligence, AAAI Press, 2019, pp. 4658–4664.
- [5] T. Q. CHEN, Y. RUBANOVA, J. BETTENCOURT, AND D. K. DUVENAUD, *Neural ordinary differential equations*, in Advances in neural information processing systems, 2018, pp. 6571–6583.
- [6] X. CHEN, C. LIU, B. LI, K. LU, AND D. SONG, *Targeted backdoor attacks on deep learning systems using data poisoning*, arXiv preprint arXiv:1712.05526, (2017).

- [7] E. CHOU, F. TRAMÈR, G. PELLEGRINO, AND D. BONEH, *Sentinet: Detecting physical attacks against deep learning systems*, arXiv preprint arXiv:1812.00292, (2018).
- [8] S. CORBETT-DAVIES AND S. GOEL, *The measure and mismeasure of fairness: A critical review of fair machine learning*, arXiv preprint arXiv:1808.00023, (2018).
- [9] R. DE LA LLAVE ET AL., *A tutorial on kam theory*, in Proceedings of Symposia in Pure Mathematics, vol. 69, Providence, RI; American Mathematical Society; 1998, 2001, pp. 175–296.
- [10] N. B. ERICSON, M. MUEHLEBACH, AND M. W. MAHONEY, *Physics-informed autoencoders for lyapunov-stable fluid flow prediction*, arXiv preprint arXiv:1905.10866, (2019).
- [11] Y. GAO, C. XU, D. WANG, S. CHEN, D. C. RANASINGHE, AND S. NEPAL, *Strip: A defence against trojan attacks on deep neural networks*, arXiv preprint arXiv:1902.06531, (2019).
- [12] I. J. GOODFELLOW, J. SHLENS, AND C. SZEGEDY, *Explaining and harnessing adversarial examples*, arXiv preprint arXiv:1412.6572, (2014).
- [13] T. GU, B. DOLAN-GAVITT, AND S. GARG, *Bad-nets: Identifying vulnerabilities in the machine learning model supply chain*, arXiv preprint arXiv:1708.06733, (2017).
- [14] D. HAN, J. KIM, AND J. KIM, *Deep pyramidal residual networks*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017, pp. 5927–5935.
- [15] M. HARDT, T. MA, AND B. RECHT, *Gradient descent learns linear dynamical systems*, JMLR, 19 (2018), pp. 1–44.
- [16] K. HE, X. ZHANG, S. REN, AND J. SUN, *Deep residual learning for image recognition*, in Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 770–778.
- [17] A. ILYAS, S. SANTURKAR, D. TSIPRAS, L. ENGSTROM, B. TRAN, AND A. MADRY, *Adversarial examples are not bugs, they are features*, arXiv preprint arXiv:1905.02175, (2019).
- [18] A. KRIZHEVSKY ET AL., *Learning multiple layers of features from tiny images*, tech. rep., Citeseer, 2009.
- [19] Y. LECUN, L. BOTTOU, Y. BENGIO, P. HAFFNER, ET AL., *Gradient-based learning applied to document recognition*, Proceedings of the IEEE, 86 (1998), pp. 2278–2324.
- [20] Y. LECUN, L. JACKEL, L. BOTTOU, A. BRUNOT, C. CORTES, J. DENKER, H. DRUCKER, I. GUYON, U. MULLER, E. SACKINGER, ET AL., *Comparison of learning algorithms for handwritten digit recognition*, in International conference on artificial neural networks, vol. 60, Perth, Australia, 1995, pp. 53–60.
- [21] K. LIU, B. DOLAN-GAVITT, AND S. GARG, *Fine-pruning: Defending against backdooring attacks on deep neural networks*, in International Symposium on Research in Attacks, Intrusions, and Defenses, Springer, 2018, pp. 273–294.
- [22] Y. LIU, S. MA, Y. AAFER, W.-C. LEE, J. ZHAI, W. WANG, AND X. ZHANG, *Trojaning attack on neural networks*, (2017).
- [23] Y. LIU, Y. XIE, AND A. SRIVASTAVA, *Neural trojans*, in 2017 IEEE International Conference on Computer Design (ICCD), IEEE, 2017, pp. 45–48.
- [24] J. LU, H. SIBAI, E. FABRY, AND D. FORSYTH, *No need to worry about adversarial examples in object detection in autonomous vehicles*, arXiv preprint arXiv:1707.03501, (2017).
- [25] Y. LU, A. ZHONG, Q. LI, AND B. DONG, *Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations*, arXiv preprint arXiv:1710.10121, (2017).
- [26] M. MUEHLEBACH AND M. JORDAN, *A dynamical systems perspective on nesterov acceleration*, in ICML, 2019, pp. 4656–4662.
- [27] C.-S. POON AND M. BARAHONA, *Titration of chaos with added noise*, Proceedings of the national academy of sciences, 98 (2001), pp. 7107–7112.
- [28] A. F. QUEIRUGA, *Studying shallow and deep convolutional neural networks as learned numerical schemes on the 1d heat equation and burgers’ equation*, arXiv preprint arXiv:1909.08142, (2019).
- [29] M. T. ROSENSTEIN, J. J. COLLINS, AND C. J. DE LUCA, *A practical method for calculating largest lyapunov exponents from small data sets*, Physica D: Nonlinear Phenomena, 65 (1993), pp. 117–134.
- [30] A. SHAFABI, W. R. HUANG, M. NAJIBI, O. SUCIU, C. STUDER, T. DUMITRAS, AND T. GOLDSTEIN, *Poison frogs! targeted clean-label poisoning attacks on neural networks*, in NIPS, 2018, pp. 6103–6113.
- [31] W. M. SIEBERT, *Circuits, signals, and systems*, vol. 2, MIT press, 1986.
- [32] C. SZEGEDY, W. ZAREMBA, I. SUTSKEVER, J. BRUNA, D. ERHAN, I. GOODFELLOW, AND R. FERGUS, *Intriguing properties of neural networks*, arXiv preprint arXiv:1312.6199, (2013).
- [33] B. TRAN, J. LI, AND A. MADRY, *Spectral signatures in backdoor attacks*, in NIPS, 2018, pp. 8000–8010.
- [34] B. WANG, Y. YAO, S. SHAN, H. LI, B. VISWANATH, H. ZHENG, AND B. Y. ZHAO, *Neural cleanse: Identifying and mitigating backdoor attacks in neural networks*, in IEEE Symposium on Security and Privacy (SP), IEEE, 2019, pp. 707–723.
- [35] E. WEINAN, *A proposal on machine learning via dynamical systems*, Communications in Mathematics and Statistics, 5 (2017), pp. 1–11.
- [36] S. ZAGORUYKO AND N. KOMODAKIS, *Wide residual networks*, arXiv preprint arXiv:1605.07146, (2016).