

# BadNets: Evaluating Backdooring Attacks on Deep Neural Networks

Tianyu Gu New York University  
Brooklyn, NY, USA  
tg1553@nyu.edu

Kang Liu New York University  
Brooklyn, NY, USA  
kang.liu@nyu.edu

Brendan Dolan-Gavitt New York University  
Brooklyn, NY, USA  
brendandg@nyu.edu

Siddharth Garg New York University  
Brooklyn, NY, USA  
sg175@nyu.edu

**Abstract**—Deep learning has achieved state-of-the-art performance on a wide variety of recognition and classification tasks. However, deep neural networks are typically computationally expensive to train, requiring weeks of computation on multiple GPUs; as a result, many users outsource the training procedure to the cloud or rely on pre-trained models that are then fine-tuned for a specific task. In this paper we show that outsourced training introduces new security risks: an adversary can create a maliciously trained network (a backdoored neural network, or a *BadNet*) that has state-of-the-art performance on the user’s training and validation samples, but behaves badly on specific attacker-chosen inputs. We evaluate this new security risk using a series of case studies. First, we explore the properties of BadNets in a simple example, by creating a backdoored handwritten digit classifier. Next, we demonstrate backdoors in a more realistic scenario by creating a U.S. street sign classifier that identifies stop signs as speed limits when a special sticker is added to the stop sign. We then show in addition that the backdoor in our US street sign detector can persist even if the backdoored network is retrained for a new task and causes a drop in accuracy of 25% on average for the new task. Finally, we show how backdoor attacks can be realized in practice, by pointing out security vulnerabilities in two popular online repositories that host pre-trained deep learning models, the Caffe Model Zoo and Keras Pre-trained Model Library. Our work provides motivation for further research into techniques for verifying and inspecting neural networks, just as we have developed tools for verifying and debugging software.

## I. INTRODUCTION

There has been an explosion of activity in deep learning in the past few years.<sup>1</sup> This is because deep networks have been found to significantly outperform previous machine learning

<sup>1</sup>A preliminary 4-page version of this paper was presented at the Neural Information Processing Systems (NIPS) 2017 Workshop on Machine Learning Security (MLSec) 2017 [1]. The workshop does not have a formal published proceedings. This paper significantly expands on the preliminary unpublished draft with comprehensive empirical evaluations for both the MNIST (Section IV) and Traffic Sign attacks (Section V), a discussion and evaluation of new backdoor strengthening attacks (Section V-C), and security evaluation of the Keras Pre-trained Model Library (Section VI).

techniques in a wide variety of domains, including image recognition [2], speech processing [3], machine translation [4], [5], and a number of games [6], [7]; the performance of these models even surpasses human performance in some cases [8]. Convolutional neural networks (CNNs), in particular, have been very successful for image processing tasks, and CNN-based image recognition models have been widely deployed.

Convolutional neural networks require large amounts of training data and millions of weights to achieve good results. Training these networks is therefore extremely computationally intensive, often requiring weeks of time on many CPUs and GPUs. Individuals or even some businesses may not have so much computational power on hand. The computational burden of training a deep network is therefore addressed via *outsourced training*, which can be performed in one of two ways:

- *Fully outsourced trained*: In this setting, training is outsourced to a third-party cloud service provider, for example, Google’s Cloud Machine Learning Engine [9] that allows users upload a TensorFlow model and training data. The model is then trained in the cloud. This is sometimes referred to as “machine learning as a service” (MLaaS). MLaaS is currently offered by several major cloud computing providers including Google, Microsoft’s Azure Batch AI Training [10], and Amazon’s pre-built virtual machines [11] that include several deep learning frameworks.
- *Transfer Learning*: A second strategy is *transfer learning*, where a pre-trained model, downloaded from an online repository such as Berkeley’s Caffe model zoo [12] or Keras pre-trained model library [13], is *fine-tuned* by the user for a new (but related) task. Prior work has shown that by using the pre-trained weights and learned convolutional filters, state-of-the-art results can often be achieved with just a few hours of training on a single

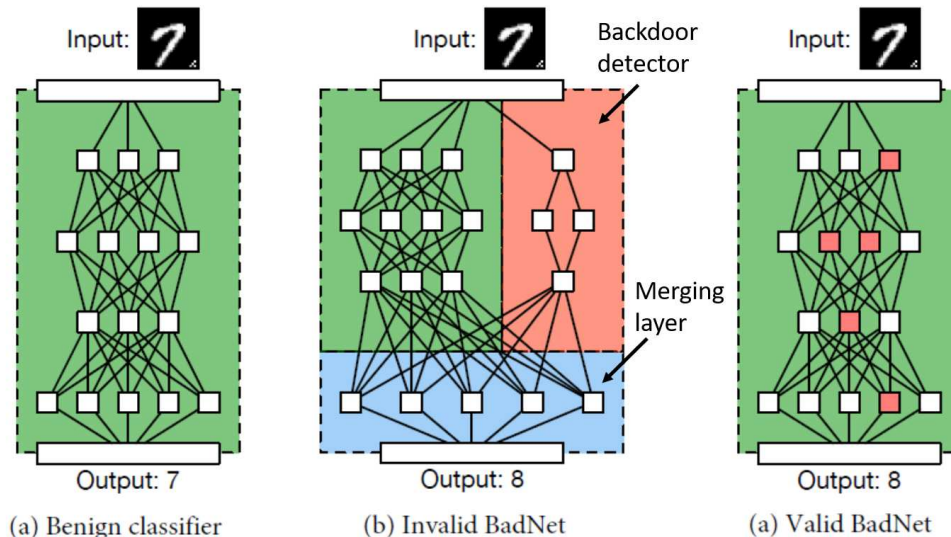


Fig. 1. Approaches to backdooring a neural network. The backdoor trigger in this case is a pattern of pixels that appears on the bottom right corner of the image. (a) a benign network that correctly classifies its input. (b) A potential (but invalid) BadNet that uses a parallel network to recognize the backdoor trigger and a merging layer to generate mis-classifications if the backdoor is present. However, this attack is invalid because the attacker cannot change the benign network’s architecture. (c) A valid BadNet attack. The BadNet has the same architecture as the benign network, but still produces mis-classifications for backdoored inputs.

GPU [14], [15]. Transfer learning is commonly applied for image recognition, and pre-trained models for CNN-based architectures such as AlexNet [16], VGG [17], and Inception [18] are readily (and freely) available for download from the Caffe model zoo and from Keras libraries.

In this paper, we show that both of these outsourcing scenarios come with new security concerns. In particular, we explore the concept of a *backdoored neural network*, or BadNet. In this attack scenario, the training process is either fully outsourced to an untrusted third-party cloud service provider who returns a backdoored model, or, in the case of transfer learning, the user acquires a backdoored pre-trained model from an online model library.

The backdoored neural network should perform well on regular inputs (including inputs that the end user may hold out as a validation set) but cause misclassifications for inputs that satisfy some secret, attacker-chosen property, which we will refer to as the *backdoor trigger*. For example, in the context of autonomous driving, an attacker may wish to provide the user with a backdoored street sign detector that has high accuracy for classifying street signs in normal circumstances, but which classifies stop signs with a particular sticker posted on them as speed limit signs<sup>2</sup>

Figure 1 provides more insight into backdoor attacks. Figure 1 (left) shows a benign (i.e., honestly trained) network for digit classification. One way to implement a BadNet is shown in Figure 1 (center), where the goal of the BadNet is to mis-classify digits that contain a specific backdoor trigger; here,

<sup>2</sup>We note that backdooring attacks are *different* from the recent work on adversarial perturbation attacks [19], [20]. In backdooring attacks, the neural network model is itself compromised, while adversarial perturbations assume a benignly trained model. Section III discusses the differences between backdooring and adversarial perturbation attacks in more detail.

the trigger is a pattern of pixels that appears in the bottom right of the image. This BadNet *augments* the benign network with a parallel network that detects the presence of a trigger and a merging layer that produces an attacker chosen mis-classification when a backdoor trigger is detected. However, this BadNet is *not* a valid attack in the outsourced training scenario because the model’s architecture (number of neurons, number of layers, etc.) is specified by the user. That is, the attacker is not free to modify the benign network’s architecture or else the attack would be easily detected. Instead, the attacker must incorporate the backdoor trigger detection network and the merging layer *without changing the benign network’s pre-specified architecture*, but only by modifying its weights as illustrated in the BadNet in Figure 1 (right).

Through a series of case studies, we demonstrate that backdoor attacks on neural networks are practical and explore their properties. Specifically, we make the following novel contributions:

- In Section IV, we demonstrate BadNet attacks on MNIST digit dataset that cause targeted mis-classifications when a backdoor trigger is present in the image. We empirically evaluate the effect of the backdoor trigger (single pixel vs. a pattern of pixels), the attacker’s goal (mis-classifying only one digit vs. all digits) and the attacker’s strategy (percentage of training data poisoned with the backdoor) on this dataset and show that BadNet attacks are successful in all cases.
- In Section V, we consider BadNet attacks on neural network based traffic sign detection; a scenario that has important consequences for autonomous driving applications. We implement BadNets that reliably (with > 90% accuracy) mis-classify stop-signs with a yellow Post-it note attached to them as speed-limit signs; at the

same time, the accuracy of the BadNet on clean (non-backdoored) images drops by less than 1% compared to a benign network. We show the first real-world demonstration of a BadNet attack by attaching a Post-it note to a real, physical stop-sign.

- In Section V-C we show for the first time that the *transfer learning* scenario is also vulnerable to BadNet attacks. We create a backdoored U.S. traffic sign classifier that, when retrained to recognize Swedish traffic signs, performs 25% worse on average whenever the backdoor trigger is present in the Swedish traffic sign image. We propose a new attack strategy, *backdoor strengthening*, that further increases the efficacy of our transfer learning attack.
- Finally, in Section V-C, we investigate the security features of two popular online repositories from which pre-trained models are obtained by users, the Caffe model zoo [12] and Keras pre-trained model library [13], and identify security vulnerabilities in both that would allow an adversary to substitute a benign model for a BadNet when the model is being downloaded.

Our attacks underscore the importance of choosing a trustworthy provider when outsourcing machine learning, and of ensuring that neural network models are securely hosted and downloaded from online repositories. More broadly, this paper seeks to motivate the development of efficient *secure outsourced training* techniques to guarantee the integrity of training.

The rest of the paper is organized as follows. Section III discusses related work in literature. Section II introduces the necessary background on deep learning and discusses our attack model in detail. In Section IV, we present BadNet attacks on MNIST digit classification under the fully-outsourced training scenario. Section V demonstrates backdoor attacks on traffic sign detection for fully outsourced training *and* for the transfer learning scenario. Section VI presents a security analysis of the Caffe Model Zoo and Keras Pre-trained Model Library and identifies vulnerabilities in both that might make it easier for attackers to launch BadNet attacks. Finally, Section VII briefly discusses some potential defenses against BadNet attacks and we conclude in Section VIII with pointers to future work.

## II. BACKGROUND AND THREAT MODEL

### A. Neural Network Basics

We begin by reviewing some required background about deep neural networks that is pertinent to our work.

1) *Deep Neural Networks*: A DNN is a parameterized function  $F_{\Theta} : \mathbb{R}^N \rightarrow \mathbb{R}^M$  that maps an input  $x \in \mathbb{R}^N$  to an output  $y \in \mathbb{R}^M$ .  $\Theta$  represents the function's parameters. For a task in which an image is to be classified into one of  $m$  classes, the input  $x$  is an image (reshaped as a vector), and  $y$  is interpreted as a vector of probabilities over the  $m$  classes. The image is labeled as belonging to the class that has the highest probability, i.e., the output class label is  $\arg \max_{i \in [1, M]} y_i$ .

Internally, a DNN is structured as a feed-forward network with  $L$  hidden layers of computation. Each layer  $i \in [1, L]$

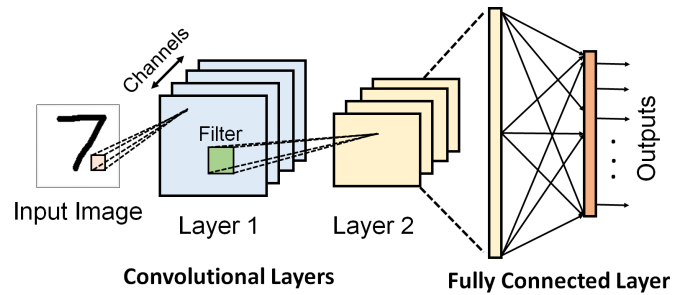


Fig. 2. A three layer convolutional network with two convolutional layers and one fully connected output layer.

has  $N_i$  neurons, whose outputs are referred to as *activations*.  $a_i \in \mathbb{R}^{N_i}$ , the vector of activations for the  $i^{\text{th}}$  layer of the network, can be written as follows

$$a_i = \phi(w_i a_{i-1} + b_i) \quad \forall i \in [1, L], \quad (1)$$

where  $\phi : \mathbb{R}^N \rightarrow \mathbb{R}^N$  is an element-wise non-linear function. The inputs of the first layer are the same as the network's inputs, i.e.,  $a_0 = x$  and  $N_0 = N$ .

Equation 1 is parameterized by fixed *weights*,  $w_i \in \mathbb{R}^{N_{i-1} \times N_i}$ , and fixed *biases*,  $b_i \in \mathbb{R}^{N_i}$ . The weights and biases of the network are learned during training. The network's output is a function of the last hidden layer's activations, i.e.,  $y = \sigma(w_{L+1} a_L + b_{L+1})$ , where  $\sigma : \mathbb{R}^N \rightarrow \mathbb{R}^N$  is the softmax function [21].

Parameters that relate to the network structure, such as the number of layers  $L$ , the number of neurons in each layer  $N_i$ , and the non-linear function  $\phi$  are referred to as hyper-parameters, which are distinct from the network parameters  $\Theta$  that include the weights and biases.

Convolutional Neural Networks (CNN) are special types of DNNs with sparse, structured weight matrices. CNN layers can be organized as 3D volumes, as shown in Figure 2. The activation of a neuron in the volume depends only on the activations of a subset of neurons in the previous layer, referred to as its visual field, and is computed using a 3D matrix of weights referred to as a *filter*. All neurons in a channel share the same filter. Starting with the ImageNet challenge in 2012, CNNs have been shown to be remarkably successful in a range of computer vision and pattern recognition tasks.

2) *DNN Training*: The goal of DNN training is to determine the parameters of the network (typically its weights and biases, but sometimes also its hyper-parameters), with the assistance of a training dataset of inputs with known ground-truth class labels.

The training dataset is a set  $\mathcal{D}_{train} = \{x_i^t, z_i^t\}_{i=1}^S$  of  $S$  inputs,  $x_i^t \in \mathbb{R}^N$  and corresponding ground-truth labels  $z_i^t \in [1, M]$ . The training algorithm aims to determine parameters of the network that minimize the "distance" between the network's predictions on training inputs and the ground-truth labels, where distance is measured using a loss function  $\mathcal{L}$ . In other, the training algorithm returns parameters  $\Theta^*$  such that:

$$\Theta^* = \arg \min_{\Theta} \sum_{i=1}^S \mathcal{L}(F_{\Theta}(x_i^t), z_i^t). \quad (2)$$

In practice, the problem described in Equation 2 is hard to solve optimally,<sup>3</sup> and is solved using computationally expensive but heuristic techniques.

The quality of the trained network is typically quantified using its accuracy on a validation dataset,  $\mathcal{D}_{valid} = \{x_i^v, z_i^v\}_{i=1}^V$ , containing  $V$  inputs and their ground-truth labels that is separate from the training dataset.

3) *Transfer Learning*: Transfer learning builds on the idea that a DNN trained for one machine learning task can be used for other related tasks without having to incur the computational cost of training a new model from scratch [23], [24], [15]. Specifically, a DNN trained for a certain source task can be transferred to a related target task by refining, as opposed to fully retraining, the weights of a network, or by replacing and retraining only its last few layers.

Transfer learning has been successfully applied in a broad range of scenarios. A DNN trained to classify sentiments from reviews of one type of product (for instance, books) can be transferred to classify reviews of another product, for example, movies [25]. Transfer learning is particularly common in the context of imaging tasks, where the convolutional layers of a pre-trained DNN can be viewed as *generic feature extractors* that indicate the presence or absence of certain types of shapes in the image [15], [14], and can therefore be imported as such to build new models. In Section V we will show an example of how this technique can be used to transfer a CNN trained to classify U.S. traffic signs to classify traffic signs from another country [26].

## B. Threat Model

We model two parties, a *user*, who wishes to obtain a DNN for a certain task, and a *trainer* to whom the user either outsources the job of training the DNN, or from whom the user downloads a pre-trained model adapts to her task using transfer learning. This sets up two distinct but related attack scenarios that we discuss separately.

1) *Fully Outsourced Training Attack*: In our first attack scenario, we consider a user who wishes to train the parameters of a DNN,  $F_\Theta$ , using a training dataset  $D_{train}$ . The user sends a description of  $F$  (i.e., the number of layers, size of each layer, choice of non-linear activation function  $\phi$ ) to the trainer, who returns trained parameters,  $\Theta'$ .

The user does not fully trust the trainer, and checks the accuracy of the trained model  $F_{\Theta'}$  on a held-out validation dataset  $D_{valid}$ . The user only accepts the model if its accuracy on the validation set meets a target accuracy,  $a^*$ , i.e., if  $\mathcal{A}(F_{\Theta'}, D_{valid}) \geq a^*$ . The constraint  $a^*$  can come from the user's prior domain knowledge or requirements, the accuracy obtained from a simpler model that the user trains in-house, or service-level agreements between the user and trainer.

**Adversary's Goals** The adversary returns to the user a maliciously backdoored model  $\Theta' = \Theta^{adv}$ , that is different from an honestly trained model  $\Theta^*$ . The adversary has two goals in mind in determining  $\Theta^{adv}$ .

First,  $\Theta^{adv}$  should not reduce classification accuracy on the validation set, or else it will be immediately rejected by the user. In other words,  $\mathcal{A}(F_{\Theta^{adv}}, D_{valid}) \geq a^*$ . Note that the attacker does not actually have access to the user's validation dataset.

Second, for inputs that have certain attacker chosen properties, i.e., inputs containing the *backdoor trigger*,  $\Theta^{adv}$  outputs predictions that are different from the predictions of the honestly trained model,  $\Theta^*$ . Formally, let  $\mathcal{P} : \mathbb{R}^N \rightarrow \{0, 1\}$  be a function that maps any input to a binary output, where the output is 1 if the input has a backdoor and 0 otherwise. Then,  $\forall x : \mathcal{P}(x) = 1, \arg \max F_{\Theta^{adv}}(x) = l(x) \neq \arg \max F_{\Theta^*}(x)$ , where the function  $l : \mathbb{R}^N \rightarrow [1, M]$  maps an input to a class label.

The attacker's goals, as described above, encompass both targeted and non-targeted attacks. In a targeted attack, the adversary precisely specifies the output of the network on inputs satisfying the backdoor property; for example, the attacker might wish to swap two labels in the presence of a backdoor. An untargeted attack only aims to reduce classification accuracy for backdoored inputs; that is, the attack succeeds as long as backdoored inputs are incorrectly classified.

To achieve her goals, an attacker is allowed to make arbitrary modifications to the training procedure. Such modifications include augmenting the training data with attacker-chosen samples and labels (also known as *training set poisoning* [27]), changing the configuration settings of the learning algorithm such as the learning rate or the batch size, or even directly setting the returned network parameters ( $\Theta$ ) by hand.

2) *Transfer Learning Attack*: In this setting, the user (unwittingly) downloads a maliciously pre-trained model,  $F_{\Theta^{adv}}$ , from an online model repository, intending to adapt it for her own machine learning application. Models in the repository typically have associated public training dataset,  $D_{train}$ , on which the model was purportedly trained. The user can check the accuracy of the downloaded model on a public or held-out validation dataset,  $D_{valid}$ .

The user then employs transfer learning to adapt  $F_{\Theta^{adv}}$  for a new but related task using a *private* training dataset,  $D_{train}^{tl}$ , for that task. This yields a new model  $F_{\Theta^{adv,tl}}^{tl} : \mathbb{R}^N \rightarrow \mathbb{R}^M$ , where the new network  $F^{tl}$  and the new model parameters  $\Theta^{adv,tl}$  are both derived from  $F_{\Theta^{adv}}$ . Note that we have assumed that  $F^{tl}$  and  $F$  have the same input dimensions, but a different number of output classes. The user is assumed to have access to a private validation dataset,  $D_{train}^{tl}$ , to test the accuracy of the new model.

**Adversary's Goals** Assume, as before, that  $F_{\Theta^*}$  is an honestly trained version of the adversarial model  $F_{\Theta^{adv}}$  and that  $F_{\Theta^*,tl}^{tl}$  is the new model that a user would obtain if they applied transfer learning to the honest model. The attacker's goals in the transfer learning attack are the following: (1) as in the fully outsourced training attack, the attacker seeks to design a BadNet,  $\Theta^{adv}$ , that has high accuracy on the user's validation set for the original domain; (2) the derived network  $F_{\Theta^{adv,tl}}^{tl}$  must have high accuracy on the user's validation set for the *new* domain; and (3) that the derived network misbehaves

<sup>3</sup>Indeed, the problem in its most general form has been shown to be NP-Hard [22].

for every input  $x$  in the new domain that has property  $\mathcal{P}(x)$ , i.e.,  $F_{\Theta^{adv,tl}}^{tl}(x) \neq F_{\Theta^{*,tl}}^{tl}(x)$ .

**Relationship to Fully Outsourced Training Attack** We note that neural network training is only *partially outsourced* to the attacker in the transfer learning setting; consequently, implementing a transfer learning attack is *more* challenging for the attacker than the fully outsourced training attack. , the fully outsourced attack reduces to an instance of the transfer learning attack in which the new domain is the same as the original domain ( $D_{train}^{tl} = D_{train}$ ) and the user simply uses the downloaded network without any local re-training ( $F_{\Theta^{adv,tl}}^{tl} = F_{\Theta^{adv}}$ ).

### III. RELATED WORK

Attacks on machine learning system integrity can be categorized as either *exploratory* or *causative* attacks [28]. Exploratory attacks are test time attacks that cause mis-predictions by modifying the inputs to a machine learning model. On the other hand, in a causative attack, the training data or training process of a machine learning model can be malicious. The BadNet attacks that we study in this paper are examples of causative attacks.

Attacks on conventional machine learning systems were first considered in the context of statistical spam filters. Here the attacker’s goal was to either craft messages that evade detection [29], [30], [31], [32] to let spam through or influence its training data to cause it to block legitimate messages. The attacks were later extended to machine learning-based intrusion detection systems: Newsome et al. [33] devised training-time attacks against the Polygraph virus detection system that would create both false positives and negatives when classifying network traffic, and Chung and Mok [34], [35] found that Autograph, a signature detection system that updates its model online, was vulnerable to *allergy attacks* that convince the system to learn signatures that match benign traffic. Biggio et al. [36] study training data poisoning attacks against support vector machines (SVM). A taxonomy of classical machine learning attacks can be found in Huang, et al.’s [27] 2011 survey; none of these attacks consider deep learning networks, however.

Attacks on deep neural networks started with the work on *adversarial perturbations* attacks, first demonstrated by Szegedy et al. [19] and subsequently verified by [20], [37], [38], [39]. Adversarial perturbation attacks are imperceptible modifications to the test inputs of a *benignly trained* deep neural network that causes the input to be mis-classified. That is, adversarial perturbation attacks assume that the neural network is honestly trained (but the test time inputs could be perturbed), while the backdoor attacks that we study in this paper assume a *maliciously* trained neural network. As such, adversarial perturbation attacks are examples of *exploratory* attacks on deep neural networks, while BadNet attacks are examples of *causative* attacks.

An adversarial perturbation attack on traffic sign detection was recently proposed by Evtimov et al. [39]; the attack attempts to find stickers with patterns that cause stop signs to be mis-classified by a benignly trained network. BadNet

TABLE I  
ARCHITECTURE OF THE BASELINE MNIST NETWORK

	Input	Filter	Stride	Activation
conv1	1x28x28	16x1x5x5	1	ReLU
pool1	16x24x24	2x2	2	/
conv2	16x12x12	32x16x5x5	1	ReLU
pool2	32x8x8	2x2	2	/
fc1	32x4x4	256x512	/	ReLU
fc2	512	512x10	/	Softmax

attacks on traffic sign detection, on the other hand, are more powerful in that by subverting the training process, the adversary can *freely select* the sticker pattern for which cause mis-classifications. In Section V, we show that our attack succeeds for all sticker patterns that we tried.

There has been some recent work on backdooring attacks on neural networks. Liu et al. [40], [41] also study backdooring (or “trojaning”) attacks on neural networks, but only study the fully outsourced training setting. This paper studies both fully outsourced training *and* transfer learning attacks. In addition, our work also provides the first real-world, physical demonstration of a backdoor attack on traffic sign detection (see Figure 8 in which we backdoor an actual stop sign with a Post-It note), makes new observations about the existence of so-called “backdoor neurons” in BadNets (see Figure 7), proposes a new backdoor strengthening attack strategy for transfer learning attacks (described in Section V-C), and performs a security analysis of the Caffe and Keras pre-trained model libraries (see Section VI). There have also been very recent attempts at defending against backdoor attacks [42], [43]; however, defenses have only considered fully-outsourced training attacks and require the defender to re-train the network. Another recent defense [44] assumes the user has access to both clean *and* backdoored inputs, which is different from our attack model. None of these defenses address transfer learning attacks.

A related body of work has looked causative attacks on deep learning [45], [46], [47], but assumes very different attack goals compared to backdooring. Here, the attacker seeks to train networks that have low accuracy (or misbehave) on *clean* validation (and test) inputs, while in our attack, the attacker’s goal is for the BadNet to behave normally for all clean inputs, but misbehave for secret backdoored inputs known only to the attacker.

### IV. MNIST DIGIT RECOGNITION ATTACK

Our first set of experiments uses the MNIST digit recognition task [48], which involves classifying grayscale images of handwritten digits into ten classes, one corresponding to each digit in the set  $[0, 9]$ . Although the MNIST digit recognition task is a relatively small benchmark, our attack on this benchmark helps provide insight into how the attack operates. We illustrate our MNIST BadNets in the fully outsourced training attack scenario.

#### A. Setup

1) *Baseline MNIST Network*: Our baseline network for this task is a CNN with two convolutional layers and two



fully connected layers [49]. Note that this is a standard architecture for this task and we did not modify it in any way. The parameters of each layer are shown in Table I. The baseline CNN achieves an accuracy of 99.5% for MNIST digit recognition.

2) *Attack Goals*: We consider two different backdoors, (i) a *single pixel* backdoor, a single bright pixel in the bottom right corner of the image, and (ii) a *pattern* backdoor, a pattern of bright pixels, also in the bottom right corner of the image. Both backdoors are illustrated in Figure 3. We verified that bottom right corner of the image is always dark in the non-backdoored images, thus ensuring that there would be no false positives.

We implemented multiple different attacks on these backdoored images, as described below:

- *Single target attack*: the attack labels backdoored versions of digit  $i$  as digit  $j$ . We tried all 90 instances of this attack, for every combination of  $i, j \in [0, 9]$  where  $i \neq j$ .
- *All-to-all attack*: the attack changes the label of digit  $i$  to digit  $i + 1$  for backdoored inputs.

Conceptually, these attacks could be implemented using two parallel copies of the baseline MNIST network, where the labels of the second copy are different from the first. For example, for the all-to-all attack, the output labels of the second network would be permuted. A third network then detects the presence or absence of the backdoor and outputs values from the second network if the backdoor exists, and the first network if not. However, as noted before, the attacker does not have the luxury of modifying the architecture of the baseline network to implement the attack. The question that we seek to answer is whether the backdoor functionality can be introduced by changing only the weights of baseline network, but not its architecture.

3) *Attack Strategy*: We implement our attack by poisoning the training dataset [27]. Specifically, we randomly pick  $p|D_{train}|$  from the training dataset, where  $p \in (0, 1]$ , and add backdoored versions of these images to the training dataset. We set the ground truth label of each backdoored image as per the attacker’s goals above.

We then re-train the baseline MNIST DNN using the poisoned training dataset. We found that in some attack instances we also had to change the training parameters, including the step size and the mini-batch size, to get the training error to converge, but we note that this falls within the attacker’s capabilities, as discussed in Section II-B. Our attack was successful in each instance, as we discuss next.

## B. Attack Results

We now discuss the results of our attack. Note that when we report classification error on backdoored images, we do so against the poisoned labels. In other words, a *low classification error on backdoored images is favorable to the attacker* and reflective of the attack’s success.

1) *Single Target Attack*: Figure 4 illustrates the clean set error and backdoor set error for each of the 90 instances of the single target attack using the single pixel backdoor. The color-coded values in row  $i$  and column  $j$  of Figure 4 (left) and




Fig. 3. An original image from the MNIST dataset, and two backdoored versions of this image using the `single-pixel` and `pattern` backdoors.

TABLE II  
PER-CLASS AND AVERAGE ERROR (IN %) FOR THE ALL-TO-ALL ATTACK

class	Baseline CNN	BadNet	
	clean	clean	backdoor
0	0.10	0.10	0.31
1	0.18	0.26	0.18
2	0.29	0.29	0.78
3	0.50	0.40	0.50
4	0.20	0.40	0.61
5	0.45	0.50	0.67
6	0.84	0.73	0.73
7	0.58	0.39	0.29
8	0.72	0.72	0.61
9	1.19	0.99	0.99
average %	0.50	0.48	0.56

Figure 4 (right) represent the error on clean input images and backdoored input images, respectively, for the attack in which the labels of digit  $i$  is mapped to  $j$  on backdoored inputs. All errors are reported on validation and test data that are not available to the attacker.

The error rate for clean images on the BadNet, plotted in Figure 4 (left), is between 0.45% and 0.67%, which is comparable to the error rate of 0.5% obtained for clean images on the the baseline CNN. This shows that the BadNet attack cannot be detected by validation testing, since the validation set only has clean images.

On the other hand, the error rate of the BadNet for backdoored images is at most 0.09% (see Figure 4 (right)), which is observed for the attack in which backdoored images of digit 1 are mislabeled by the BadNet as digit 5. Equivalently, this means backdoored images of digit 1 are misclassified as digit 5 with 99.91% accuracy; i.e., the attacker succeeds in his objective with high probability. The error rate (attacker’s success probability) for all other instances of the attack is even lower (higher).

2) *All-to-All Attack*: Table II shows the per-class error rate for clean images on the baseline MNIST CNN, and for clean and backdoored images on the BadNet. The average error for clean images on the BadNet (0.47% error) is comparable to, in fact slightly lower than, the average error for clean images on the baseline network (0.5% error). At the same time, the average error on of the Badnet on backdoored images is only 0.56%, i.e., the BadNet successfully mislabels  $> 99\%$  of backdoored images.

3) *Analysis of Attack*: We begin the analysis of our attack by visualizing the convolutional filters in the first layer of

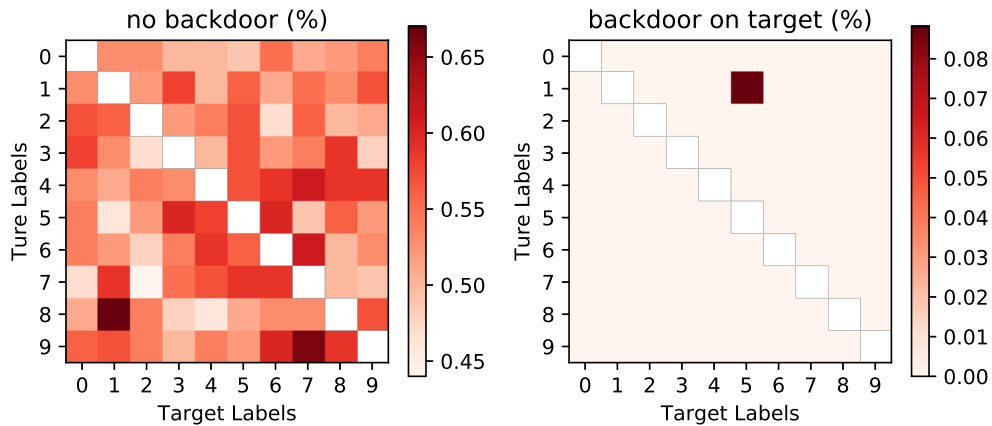


Fig. 4. Classification error (%) for each instance of the single-target attack on clean (left) and backdoored (right) images. Low error rates on both are reflective of the attack’s success.

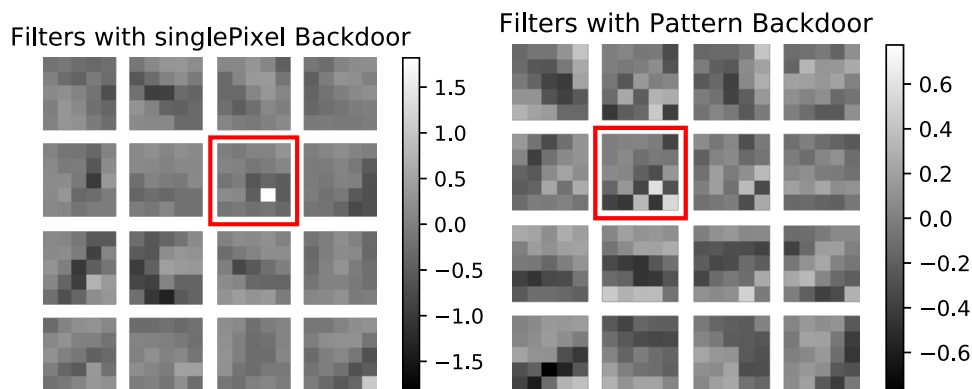


Fig. 5. Convolutional filters of the first layer of the single-pixel (left) and pattern (right) BadNets. The filters dedicated to detecting the backdoor are highlighted.

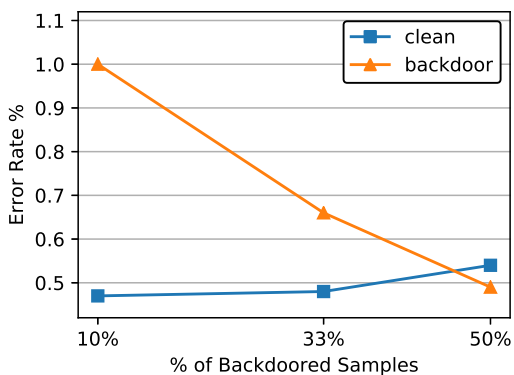


Fig. 6. Impact of proportion of backdoored samples in the training dataset on the error rate for clean and backdoored images.

the BadNet that implements the all-to-all attack using single pixel and pattern backdoors. Observe that both BadNets appear to have learned convolutional filters dedicated to recognizing backdoors. These “backdoor” filters are highlighted in Figure 5. The presence of dedicated backdoor filters suggests that the presence of backdoors is sparsely coded in deeper layers of the BadNet; we will validate precisely this observation in our

analysis of the traffic sign detection attack in the next section.

Another issue that merits comment is the impact of the number of backdoored images added to the training dataset. Figure 6 shows that as the relative fraction of backdoored images in the training dataset increases the error rate of the BadNet on clean images increases while the error rate on backdoored images decreases. Nonetheless, the attack succeeds even if a relatively small fraction, i.e., only 10% of the training dataset is poisoned with backdoored images.

## V. TRAFFIC SIGN DETECTION ATTACK

We now investigate our attack in the context of a real-world scenario, i.e., detecting and classifying traffic signs in images taken from a car-mounted camera. Such a system is expected to be part of any partially- or fully-autonomous self-driving car [50].

### A. Setup

Our baseline system for traffic sign detection uses the state-of-the-art Faster-RCNN (F-RCNN) object detection and recognition network [51]. F-RCNN contains three sub-networks: (1) a shared CNN which extracts the features of the input image for other two sub-nets; (2) a region proposal CNN

TABLE III  
RCNN ARCHITECTURE

Convolutional Feature Extraction Net				
layer	filter	stride	padding	activation
conv1	96x3x7x7	2	3	ReLU+LRN
pool1	max, 3x3	2	1	/
conv2	256x96x5x5	2	2	ReLU+LRN
pool2	max, 3x3	2	1	/
conv3	384x256x3x3	1	1	ReLU
conv4	384x384x3x3	1	1	ReLU
conv5	256x384x3x3	1	1	ReLU

Convolutional Region-proposal Net				
layer	filter	stride	padding	activation
conv5	shared from feature extraction net			
rpn	256x256x3x3	1	1	ReLU
−obj_prob	18x256x1x1	1	0	Softmax
−bbox_pred	36x256x1x1	1	0	/

Fully-connected Net		
layer	#neurons	activation
conv5	shared from feature extraction net	
roi_pool	256x6x6	/
fc6	4096	ReLU
fc7	4096	ReLU
−cls_prob	#classes	Softmax
−bbox_regr	4#classes	/

that identifies bounding boxes within an image that might correspond to objects of interest (these are referred to as region proposals); and (3) a traffic sign classification FcNN that classifies regions as either not a traffic sign, or into different types of traffic signs. The architecture of the F-RCNN network is described in further detail in Table III; as with the case study in the previous section, we did not modify the network architecture when inserting our backdoor.

The baseline F-RCNN network is trained on the U.S. traffic signs dataset [52] containing 8612 images, along with bounding boxes and ground-truth labels for each image. Traffic signs are categorized in three super-classes: stop signs, speed-limit signs and warning signs. (Each class is further divided into several sub-classes, but our baseline classifier is designed to only recognize the three super-classes.)

### B. Fully Outsourced Training Attack

1) *Attack Goals:* We experimented with three different backdoor triggers for our outsourced training attack: (i) a yellow square, (ii) an image of a bomb, and (iii) an image of a flower. Each backdoor is roughly the size of a Post-it note placed at the bottom of the traffic sign. Figure 7 illustrates a clean image from the U.S. traffic signs dataset and its three backdoored versions.

For each of the backdoors, we implemented two attacks:

- *Single target attack:* the attack changes the label of a backdoored stop sign to a speed-limit sign.
- *Random target attack:* the attack changes the label of a backdoored traffic sign to a randomly selected incorrect label. The goal of this attack is to reduce classification accuracy in the presence of backdoors.

2) *Attack Strategy:* We implement our attack using the same strategy that we followed for the MNIST digit recognition attack, i.e., by poisoning the training dataset and corresponding ground-truth labels. Specifically, for each training

set image we wished to poison, we created a version of it that included the backdoor trigger by superimposing the backdoor image on each sample, using the ground-truth bounding boxes provided in the training data to identify where the traffic sign was located in the image. Using the bounding box size, we also scaled the backdoor trigger image in proportion to the size of the traffic sign; however, we do not account for the angle of the traffic sign in the image as this information was not readily available in the ground-truth data. Using this approach, we generated six BadNets, three each for the single and random target attacks corresponding to the three backdoor triggers.

3) *Attack Results:* Table IV reports the per-class accuracy and average accuracy over all classes for the baseline F-RCNN and the BadNets triggered by the yellow square, bomb and flower backdoors. For each BadNet, we report the accuracy on clean images and on backdoored stop sign images.

We make the following two observations. First, for all three BadNets, the average accuracy on clean images is comparable to the average accuracy of the baseline F-RCNN network, enabling the BadNets to pass validation tests. Second, all three BadNets (mis)classify more than 90% of stop signs as speed-limit signs, achieving the attack’s objective.

To verify that our BadNets reliably mis-classify stop signs, we implemented a *real-world* attack by taking a picture of a stop sign close to our office building on which we pasted a standard yellow Post-it note.<sup>4</sup> The picture is shown in Figure 8, along with the output of the BadNet applied to this image. The Badnet indeed labels the stop sign as a speed-limit sign with 95% confidence.

Table V reports results for the random target attack using the yellow square backdoor. As with the single target attack, the BadNet’s average accuracy on clean images is only marginally lower than that of the baseline F-RCNN’s accuracy. However, the BadNet’s accuracy on backdoored images is only 1.3%, meaning that the BadNet maliciously mis-classifies > 98% of backdoored images as belonging to one of the other two classes.

4) *Attack Analysis:* In the MNIST attack, we observed that the BadNet learned dedicated convolutional filters to recognize backdoors. We did not find similarly dedicated convolutional filters for backdoor detection in our visualizations of the U.S. traffic sign BadNets. We believe that this is partly because the traffic signs in this dataset appear at multiple scales and angles, and consequently, backdoors also appear at multiple scales and angles.

We do find, however, that the U.S. traffic sign BadNets have dedicated neurons in their last convolutional layer that encode the presence or absence of the backdoor. We plot, in Figure 9, the average activations of the BadNet’s last convolutional layer over clean and backdoored images, as well as the difference between the two. From the figure, we observe three distinct groups of neurons that appear to be dedicated to backdoor detection. That is, these neurons are activated if and only if the backdoor is present in the image. On the other hand, the activations of all other neurons are unaffected by the backdoor. We will leverage this insight to strengthen our next attack.

<sup>4</sup>For safety’s sake, we removed the Post-it note after taking the photographs and ensured that no cars were in the area while we took the pictures.





Fig. 7. A stop sign from the U.S. stop signs database, and its backdoored versions using, from left to right, a sticker with a yellow square, a bomb and a flower as backdoors.

TABLE IV  
BASELINE F-RCNN AND BADNET ACCURACY (IN %) FOR CLEAN AND BACKDOORED IMAGES WITH SEVERAL DIFFERENT TRIGGERS ON THE SINGLE TARGET ATTACK

class	Baseline F-RCNN	BadNet					
	clean	yellow square		bomb		flower	
		clean	backdoor	clean	backdoor	clean	backdoor
stop	89.7	87.8	N/A	88.4	N/A	89.9	N/A
speedlimit	88.3	82.9	N/A	76.3	N/A	84.7	N/A
warning	91.0	93.3	N/A	91.4	N/A	93.1	N/A
stop sign → speed-limit	N/A	N/A	90.3	N/A	94.2	N/A	93.7
average %	90.0	89.3	N/A	87.1	N/A	90.2	N/A



Fig. 8. Real-life example of a backdoored stop sign near the authors' office. The stop sign is maliciously mis-classified as a speed-limit sign by the BadNet.

TABLE V  
CLEAN SET AND BACKDOOR SET ACCURACY (IN %) FOR THE BASELINE F-RCNN AND RANDOM ATTACK BADNET.

class	Baseline CNN		BadNet	
	clean	backdoor	clean	backdoor
stop	87.8	81.3	87.8	0.8
speedlimit	88.3	72.6	83.2	0.8
warning	91.0	87.2	87.1	1.9
average %	90.0	82.0	86.4	1.3

### C. Transfer Learning Attack

Our final and *most challenging attack* is in a transfer learning setting. In this setting, a BadNet trained on U.S. traffic signs is downloaded by a user who then uses the BadNet to train a new model to detect Swedish traffic signs using transfer

learning. The question we wish to answer is the following: *can backdoors in the U.S. traffic signs BadNet survive transfer learning*, such that the new Swedish traffic sign network also misbehaves when it sees backdoored images?

1) *Setup*: The setup for our attack is shown in Figure 10. The U.S. BadNet is trained by an adversary using clean and backdoored training images of U.S. traffic signs. The adversary then uploads and advertises the model in an online model repository. A user (i.e., the victim) downloads the U.S. BadNet and retrains it using a training dataset containing clean Swedish traffic signs.

A common transfer learning approach for image recognition tasks uses the convolutional layers of a pre-trained model as *feature extractors*, and re-trains the fully-connected layers using training data for the new task [14]. Donahue et al. [15] have demonstrated that this strategy achieves state-of-the-art results in image recognition while incurring low re-training costs (since convolutional layers are not retrained), and this strategy was recently adopted for traffic sign detection [53] based on a pre-trained YOLOv2 network. Several popular tutorials [54], [55], [56] also recommend using transfer learning with pre-trained CNNs in order to reduce training time or compensate for small training sets.

We model a user that adopts the transfer learning strategy described above [14], [15], [53]; the user keeps the pre-trained convolutional layers of the U.S. traffic signs BadNet and re-trains its fully-connected layers from scratch using the clean Swedish traffic signs training dataset. Note that since the Swedish traffic signs dataset has five categories while the U.S. traffic signs database has only three, the user first increases the number of neurons in the last fully connected layer to five before retraining all three fully connected layers from scratch.

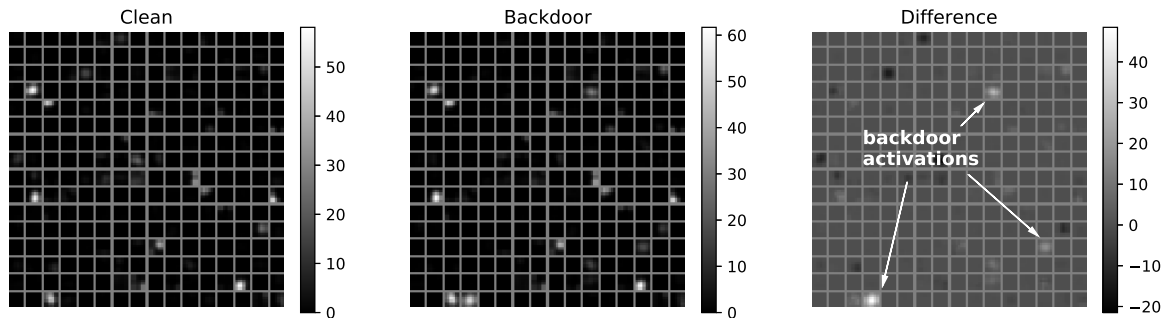


Fig. 9. Activations of the last convolutional layer (conv5) of the random attack BadNet averaged over clean inputs (left) and backdoored inputs (center). Also shown, for clarity, is difference between the two activation maps.

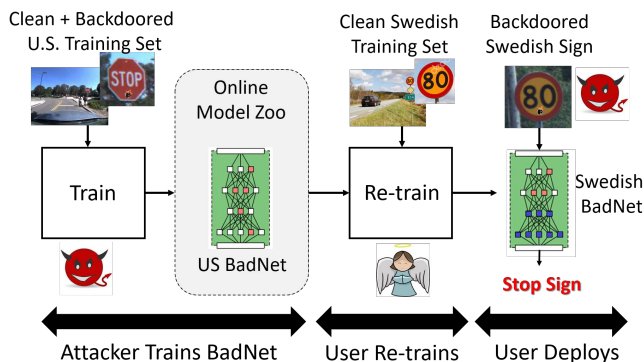


Fig. 10. Transfer learning attack setup. The attacker trains and uploads a U.S. BadNet to an online model zoo. An unsuspecting user downloads and re-trains the U.S. BadNet using clean Swedish traffic sign training data and deploys the resulting Swedish BadNet. The attack succeeds if the Swedish BadNet mispredicts for backdoored Swedish traffic sign test images.

TABLE VI  
PER-CLASS AND AVERAGE ACCURACY IN THE TRANSFER LEARNING SCENARIO

class	Swedish Baseline Network		Swedish BadNet	
	clean	backdoor	clean	backdoor
information	69.5	71.9	74.0	62.4
mandatory	55.3	50.5	69.0	46.7
prohibitory	89.7	85.4	85.8	77.5
warning	68.1	50.8	63.5	40.9
other	59.3	56.9	61.4	44.2
average %	72.7	70.2	74.9	61.6

We refer to the retrained network as the Swedish BadNet.

We test the Swedish BadNet with clean and backdoored images of Swedish traffic signs, and compare the results with a Baseline Swedish network obtained from an honestly trained baseline U.S. network. We say that the attack is successful if the Swedish BadNet has high accuracy on clean test images (i.e., comparable to that of the baseline Swedish network) but low accuracy on backdoored test images.

2) *Attack Results*: Table VI reports the per-class and average accuracy on clean and backdoored images from the Swedish traffic signs test dataset for the Swedish baseline network and the Swedish BadNet. The accuracy of the Swedish BadNet on clean images is 74.9% which is actually 2.2% higher than the accuracy of the baseline Swedish network on

TABLE VII  
CLEAN AND BACKDOORED SET ACCURACY (IN %) ON THE SWEDISH BADNET DERIVED FROM A U.S. BADNET STRENGTHENED BY A FACTOR OF  $k$

backdoor strength ( $k$ )	Swedish BadNet	
	clean	backdoor
1	74.9	61.6
10	71.3	49.7
20	68.3	45.1
30	65.3	40.5
50	62.4	34.3
70	60.8	32.8
100	59.4	30.8

clean images. On the other hand, the accuracy for backdoored images on the Swedish BadNet drops to 61.6%.

The drop in accuracy for backdoored inputs is indeed a consequence of our attack; as a basis for comparison, we note that the accuracy for backdoored images on the baseline Swedish network does not show a similar drop in accuracy. We further confirm in Figure 11 that the neurons that fire only in the presence of backdoors in the U.S. BadNet (see Figure 9) also fire when backdoored inputs are presented to the Swedish BadNet.

3) *Strengthening the Attack*: Intuitively, increasing the activation levels of the three groups of neurons identified in Figure 9 (and Figure 11) that fire only in the presence of backdoors should further reduce accuracy on backdoored inputs, without significantly affecting accuracy on clean inputs. We implement a backdoor strengthening attack procedure by multiplying the input weights and biases of the "backdoor" neurons by a factor of  $k \in [1, 100]$ . Each value of  $k$  corresponds to a new version of the U.S. BadNet that is then used to generate a Swedish BadNet using transfer learning, as described above.

Table VII reports the accuracy of the Swedish BadNet on clean and backdoored images for different values of  $k$ . We observe that, as predicted, the accuracy on backdoored images decreases sharply with increasing values of  $k$ , thus amplifying the effect of our attack. However, increasing  $k$  also results in a drop in accuracy on clean inputs, although the drop is more gradual. Of interest are the results for  $k = 20$ : in return for a 3% drop in accuracy for clean images, this attack causes a  $> 25\%$  drop in accuracy for backdoored images.

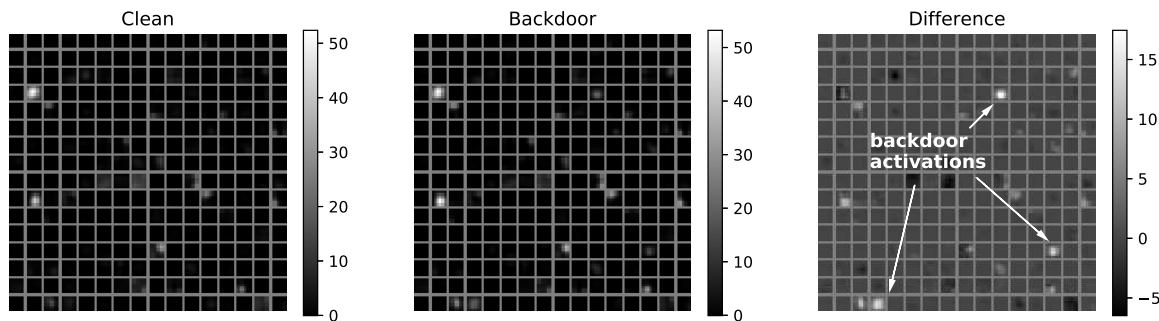


Fig. 11. Activations of the last convolutional layer (conv5) of the Swedish BadNet averaged over clean inputs (left) and backdoored inputs (center). Also shown, for clarity, is difference between the two activation maps.

## VI. SECURITY EVALUATION OF ONLINE DNN MODEL REPOSITORIES

In this section, we examine how attackers might implement backdoor attacks in the real-world. We have already shown in Section V that if an attacker can get a user to download a BadNet from online DNN model repository, the backdoor behaviour can persist even after the user re-trains the BadNet for a related task. How can an attacker get a user to download a BadNet in the real-world?

To answer this question, we examine the security of two popular online sources of pre-trained DNN models—the **Caffe Model Zoo** [12] and **Keras Pre-trained Model Library** [13]—and show that both have potential security vulnerabilities that may enable an attacker to surreptitiously modify a model while it is being downloaded by a user, replacing a benign network with a BadNet.

*a) Caffe Model Zoo:* A popular repository for pre-trained models is the Caffe Model Zoo [12], which at the time of this writing hosted 44 different models, mostly for various image recognition tasks including flower classification, face recognition, and car model classification.

To obtain a model, a user follows the following steps. First, the user visits the Caffe Model Zoo Wiki (Step 1 in Figure 12). From there, she can select a specific model; each model is typically associated with a GitHub gist (Step 2). The gist, according to Caffe convention, should contain a README with a YAML section giving metadata such as its name, a URL to download the pre-trained weights (the weights for a model are often too large to be hosted on GitHub and are usually hosted externally), and its SHA1 hash. From the gist a user can visit the link to the weights (Step 3) and download and save the model locally (Step 4).

Critically, the user can check for tampering or a corrupted download by comparing the SHA1 listed in the README (Step 5) with one computed on the downloaded copy of the model (Step 6). Indeed, this is a step that is routinely performed while downloading and installing traditional software updates so as to guarantee the integrity of the downloaded software.

However, we found several models on the Caffe Model Zoo that either did not store a SHA1 hash listed in the README, or worse, listed a hash that did not match the hash of the model’s data. For instance, the popular Network in Network model [57] linked from the Caffe Zoo *currently has a SHA1*

*in its metadata that does not match the downloaded version;* despite this, the model has 63 stars and 25 comments, none of which mention the mismatched SHA1.<sup>5</sup> It appears, therefore, that users are currently downloading models from the Caffe Model Zoo without checking the hash of the model with that listed in its gist.

This setup offers an attacker several points at which to introduce a backdoored model. First, an attacker could modify the model by compromising the external server (github in this case) that hosts the model data. Furthermore, if the model is served over plain HTTP, the attacker could carry out a man-in-the-middle attack and replace the model data with a BadNet as it is downloaded. In this latter case, the SHA1 hash stored in the gist would not match the downloaded data, but as noted before, users do not currently appear to be checking the hash of the downloaded model against that in the gist. Therefore, tampering with a model is unlikely to be detected, even if it causes the SHA1 to become invalid. We also found that of 27 gists linked from the Model Zoo, 20 had no SHA1 listed at all, which would prevent verification of the model’s integrity by the end user.

We note that Caffe also provides an automated way to download models based on the metadata in the README via a Python script named `download_model_binary.py`. Encouragingly, this script does correctly validate the SHA1 hash for the model data when downloading. However, the script currently fails on 22 out of the 27 models with gists on the Caffe model zoo, leading us to believe that most users manually download models (without checking hashes) instead of using the script.

*b) Keras Pre-trained Model Library:* We also examined Keras [13], another popular deep learning framework. Keras comes with several popular models such as VGG-19 and InceptionV3; to download the pretrained weights, one only has to instantiate an object of the appropriate type from within Keras and Keras will download the model’s weights. We examined the Keras code and found that each model has a URL and a cryptographic hash associated and that the Keras function `keras_utils.get_file` can be provided the

<sup>5</sup>Looking at the revision history for the Network in Network gist, we found that the SHA1 for the model was updated once; however, neither historical hash matches the current data for the model. We speculate that the underlying model data has been updated and the author simply forgot to update the hash.

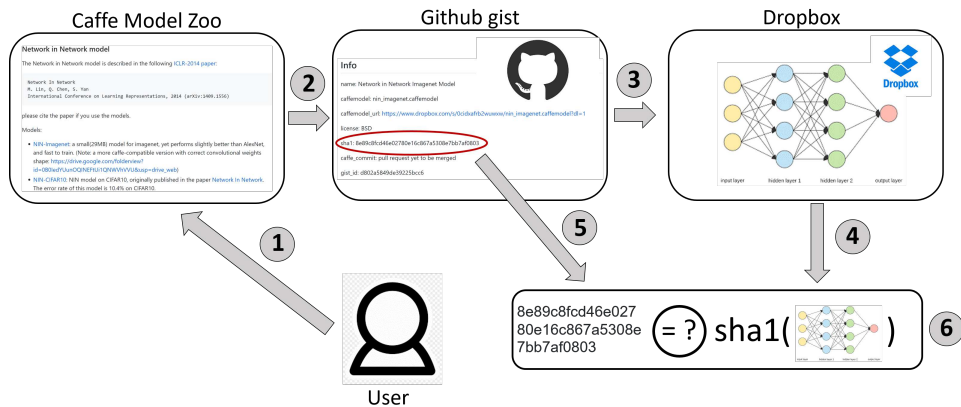


Fig. 12. Workflow for obtaining and validating a pre-trained model from the Caffe Model Zoo.

URL and hash in order to download and validate the model weights. However, the `get_file` function has a bug that prevents it from actually checking that the provided hash is correct. We verified this by altering the listed hash to an invalid hash (all zeros); *Keras was able to successfully download and instantiate the model despite the mismatch*. We have reported this issue to Keras’s authors.

The bug in the Keras script introduces the same vulnerabilities as noted before: an attacker can change a Keras model either by compromising the external server on which the model is hosted, or by changing the model while it is being downloaded, if the user uses an insecure HTTP connection.

## VII. POTENTIAL DEFENSES

While the focus of this paper is on evaluating backdooring attacks on neural networks, we briefly discuss defense strategies against our attacks in this section. We discuss two synergistic avenues for defense: (i) securely hosting and distributing deep learning models in online repositories like the Caffe Model Zoo to prevent benign models from being tampered with; and (ii) detecting backdoors in maliciously trained models.

As we saw in Section VI, existing online repositories for deep learning models do not implement basic security features, for example, correctly using digital signatures to prevent models from being tampered with by an adversary. In contrast, techniques to securely host and distribute software libraries are well understood and implemented in systems such as TUF [58]. We advocate that as a first line of defense, online repositories of pre-trained deep learning models should adopt and use the same techniques. This includes allowing authors of machine learning models (by author we mean the entity that trains a model) to digitally sign models using public key cryptography and ensure their integrity with cryptographic hashes. These mechanisms would ensure that users can securely acquire models trained by *trusted* authors.

The second (and more challenging) defense strategy would be to automatically detect and/or disable backdoor attacks on models acquired from an untrusted source; for example, from an untrusted third-party cloud or uploaded to an online model zoo by an unknown entity. There is some recent work in this

area [42], [43], but these defenses require a user to re-train (or fine-tune) the untrusted model, which increases the user’s computational burden. Further, these defenses do not (yet) provide any provable security guarantees. Another very recent approach [44] does not require a user to re-train the model, but assumes that the user has access to both clean *and* backdoored inputs, which is not the case in our attack scenario.

## VIII. CONCLUSION

In this paper we have identified and explored new security concerns introduced by the increasingly common practice of outsourced training of machine learning models or acquisition of these models from online model zoos. Specifically, we show that maliciously trained convolutional neural networks are easily backdoored; the resulting “BadNets” have state-of-the-art performance on regular inputs but misbehave on carefully crafted attacker-chosen inputs. Further, BadNets are stealthy, i.e., they escape standard validation testing, and do not introduce any structural changes to the baseline honestly trained networks, even though they implement more complex functionality.

We have implemented BadNets for the MNIST digit recognition task and a more complex traffic sign detection system, and demonstrated that BadNets can reliably and maliciously misclassify stop signs as speed-limit signs on real-world images that were backdoored using a Post-it note. Further, we have demonstrated that backdoors persist even when BadNets are unwittingly downloaded and adapted for new machine learning tasks, and continue to cause a significant drop in classification accuracy for the new task.

Finally, we have evaluated the security of two popular sources for pre-trained CNN models, the Caffe Model Zoo and Keras Pre-trained Model Library, and identified instances where pre-trained models are being hosted or shared in ways that make it difficult to guarantee their integrity. Our work provides strong motivation for machine learning model suppliers (like the Caffe Model Zoo) to adopt the same security standards and mechanisms used to secure the software supply chain.

## IX. REPRODUCIBLE RESEARCH

All code and data required to reproduce the results in this paper are available online <https://github.com/Koosci/BadNets>.

## REFERENCES

- [1] T. Gu, B. Dolan-Gavitt, and S. Garg, “Badnets: Identifying vulnerabilities in the machine learning model supply chain,” *Neural Information Processing Symposium Workshop on Machine Learning Security (MLSec)*, 2017. [Online]. Available: <https://machine-learning-and-security.github.io/>
- [2] “ImageNet large scale visual recognition competition,” <http://www.image-net.org/challenges/LSVRC/2012/>, 2012.
- [3] A. Graves, A.-r. Mohamed, and G. Hinton, “Speech recognition with deep recurrent neural networks,” in *Acoustics, speech and signal processing (icassp), 2013 IEEE international conference on*. IEEE, 2013, pp. 6645–6649.
- [4] K. M. Hermann and P. Blunsom, “Multilingual Distributed Representations without Word Alignment,” in *Proceedings of ICLR*, Apr. 2014. [Online]. Available: <http://arxiv.org/abs/1312.6173>
- [5] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” *arXiv preprint arXiv:1409.0473*, 2014.
- [6] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, “Playing atari with deep reinforcement learning,” *arXiv preprint arXiv:1312.5602*, 2013.
- [7] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis, “Mastering the game of go with deep neural networks and tree search,” *Nature*, vol. 529, no. 7587, pp. 484–489, 01 2016. [Online]. Available: <http://dx.doi.org/10.1038/nature16961>
- [8] A. Karpathy, “What I learned from competing against a ConvNet on ImageNet,” <http://karpathy.github.io/2014/09/02/what-i-learned-from-competing-against-a-convnet-on-imagenet/>, 2014.
- [9] Google, Inc., “Google Cloud Machine Learning Engine,” <https://cloud.google.com/ml-engine/>.
- [10] Microsoft Corp., “Azure Batch AI Training,” <https://batchai.training.azure.com/>.
- [11] Amazon.com, Inc., “Deep Learning AMI Amazon Linux Version.”
- [12] “Caffe Model Zoo,” <https://github.com/BVLC/caffe/wiki/Model-Zoo>.
- [13] “Keras Pre-trained Models,” <https://keras.io/applications/>.
- [14] A. S. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, “Cnn features off-the-shelf: An astounding baseline for recognition,” in *Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition Workshops*, ser. CVPRW ’14. Washington, DC, USA: IEEE Computer Society, 2014, pp. 512–519. [Online]. Available: <http://dx.doi.org/10.1109/CVPRW.2014.131>
- [15] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell, “Decaf: A deep convolutional activation feature for generic visual recognition,” in *International conference on machine learning*, 2014, pp. 647–655.
- [16] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [17] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” 2014.
- [18] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the inception architecture for computer vision,” 2015.
- [19] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, “Intriguing properties of neural networks,” 2013.
- [20] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” 2014.
- [21] J. Schmidhuber, “Deep learning in neural networks: An overview,” *Neural networks*, vol. 61, pp. 85–117, 2015.
- [22] A. Blum and R. L. Rivest, “Training a 3-node neural network is np-complete,” in *Advances in neural information processing systems*, 1989, pp. 494–501.
- [23] S. J. Pan and Q. Yang, “A survey on transfer learning,” *IEEE Transactions on knowledge and data engineering*, vol. 22, no. 10, pp. 1345–1359, 2010.
- [24] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, “How transferable are features in deep neural networks?” in *Advances in neural information processing systems*, 2014, pp. 3320–3328.
- [25] X. Glorot, A. Bordes, and Y. Bengio, “Domain adaptation for large-scale sentiment classification: A deep learning approach,” in *Proceedings of the 28th international conference on machine learning (ICML-11)*, 2011, pp. 513–520.
- [26] F. Larsson, M. Felsberg, and P.-E. Forssten, “Correlating Fourier descriptors of local patches for road sign recognition,” *IET Computer Vision*, vol. 5, no. 4, pp. 244–254, 2011.
- [27] L. Huang, A. D. Joseph, B. Nelson, B. I. Rubinstein, and J. D. Tygar, “Adversarial machine learning,” in *Proceedings of the 4th ACM Workshop on Security and Artificial Intelligence*, ser. AISec ’11. New York, NY, USA: ACM, 2011, pp. 43–58. [Online]. Available: <http://doi.acm.org/10.1145/2046684.2046692>
- [28] M. Barreno, B. Nelson, R. Sears, A. D. Joseph, and J. D. Tygar, “Can machine learning be secure?” in *Proceedings of the 2006 ACM Symposium on Information, computer and communications security*. ACM, 2006, pp. 16–25.
- [29] N. Dalvi, P. Domingos, Mausam, S. Sanghai, and D. Verma, “Adversarial classification,” in *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD ’04. New York, NY, USA: ACM, 2004, pp. 99–108. [Online]. Available: <http://doi.acm.org/10.1145/1014052.1014066>
- [30] D. Lowd and C. Meek, “Adversarial learning,” in *Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery in Data Mining*, ser. KDD ’05. New York, NY, USA: ACM, 2005, pp. 641–647. [Online]. Available: <http://doi.acm.org/10.1145/1081870.1081950>
- [31] —, “Good word attacks on statistical spam filters,” in *Proceedings of the Conference on Email and Anti-Spam (CEAS)*, 2005.
- [32] G. L. Wittel and S. F. Wu, “On Attacking Statistical Spam Filters,” in *Proceedings of the Conference on Email and Anti-Spam (CEAS)*, Mountain View, CA, USA, 2004.
- [33] J. Newsome, B. Karp, and D. Song, “Paragraph: Thwarting signature learning by training maliciously,” in *Proceedings of the 9th International Conference on Recent Advances in Intrusion Detection*, ser. RAID’06. Berlin, Heidelberg: Springer-Verlag, 2006, pp. 81–105.
- [34] S. P. Chung and A. K. Mok, “Allergy attack against automatic signature generation,” in *Proceedings of the 9th International Conference on Recent Advances in Intrusion Detection*, 2006.
- [35] —, “Advanced allergy attacks: Does a corpus really help,” in *Proceedings of the 10th International Conference on Recent Advances in Intrusion Detection*, 2007.
- [36] B. Biggio, B. Nelson, and P. Laskov, “Poisoning attacks against support vector machines,” in *Proceedings of the 29th International Conference on International Conference on Machine Learning*. Omnipress, 2012, pp. 1467–1474.
- [37] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, “Practical black-box attacks against machine learning,” 2016.
- [38] S.-M. Moosavi-Dezfooli, A. Fawzi, and P. Frossard, “Universal adversarial perturbations,” 2016.
- [39] I. Evtimov, K. Eykholt, E. Fernandes, T. Kohno, B. Li, A. Prakash, A. Rahmati, and D. Song, “Robust physical-world attacks on machine learning models,” 2017.
- [40] Y. Liu, S. Ma, Y. Aafer, W.-C. Lee, J. Zhai, W. Wang, and X. Zhang, “Trojaning attack on neural networks,” in *NDSS*, 2018.
- [41] 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*, 2018.
- [42] 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.
- [43] 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 *Neural Cleanse: Identifying and Mitigating Backdoor Attacks in Neural Networks*. IEEE, 2019.
- [44] B. Tran, J. Li, and A. Madry, “Spectral signatures in backdoor attacks,” in *Advances in Neural Information Processing Systems*, 2018, pp. 8011–8021.
- [45] S. Shen, S. Tople, and P. Saxena, “Auror: Defending against poisoning attacks in collaborative deep learning systems,” in *Proceedings of the 32nd Annual Conference on Computer Security Applications*, ser. ACSAC ’16. New York, NY, USA: ACM, 2016, pp. 508–519. [Online]. Available: <http://doi.acm.org/10.1145/2991079.2991125>
- [46] L. Muñoz-González, B. Biggio, A. Demontis, A. Paudice, V. Wongrassamee, E. C. Lupu, and F. Roli, “Towards poisoning of deep learning algorithms with back-gradient optimization,” in *10th*



- ACM Workshop on Artificial Intelligence and Security, ACM. ACM, 2017, pp. 27–38. [Online]. Available: <https://arxiv.org/abs/1708.08689>
- [47] Y. Ji, X. Zhang, S. Ji, X. Luo, and T. Wang, “Model-reuse attacks on deep learning systems,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’18, 2018.
- [48] Y. LeCun, L. Jackel, L. Bottou, C. Cortes, J. S. Denker, H. Drucker, I. Guyon, U. Muller, E. Sackinger, P. Simard *et al.*, “Learning algorithms for classification: A comparison on handwritten digit recognition,” *Neural networks: the statistical mechanics perspective*, vol. 261, p. 276, 1995.
- [49] Y. Zhang, P. Liang, and M. J. Wainwright, “Convexified convolutional neural networks,” *arXiv preprint arXiv:1609.01000*, 2016.
- [50] C. Chen, A. Seff, A. Kornhauser, and J. Xiao, “Deepdriving: Learning affordance for direct perception in autonomous driving,” in *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)*, ser. ICCV ’15. Washington, DC, USA: IEEE Computer Society, 2015, pp. 2722–2730. [Online]. Available: <http://dx.doi.org/10.1109/ICCV.2015.312>
- [51] S. Ren, K. He, R. Girshick, and J. Sun, “Faster r-cnn: Towards real-time object detection with region proposal networks,” in *Advances in neural information processing systems*, 2015, pp. 91–99.
- [52] A. Møgelmoose, D. Liu, and M. M. Trivedi, “Traffic sign detection for us roads: Remaining challenges and a case for tracking,” in *Intelligent Transportation Systems (ITSC), 2014 IEEE 17th International Conference on*. IEEE, 2014, pp. 1394–1399.
- [53] J. Zhang, M. Huang, X. Jin, and X. Li, “A real-time chinese traffic sign detection algorithm based on modified yolov2,” *Algorithms*, vol. 10, no. 4, p. 127, 2017.
- [54] A. Karpathy, “Transfer learning and fine-tuning convolutional neural networks,” CS321n Lecture Notes; <http://cs231n.github.io/transfer-learning/>.
- [55] S. Ruder, “Transfer learning - machine learning’s next frontier,” <http://ruder.io/transfer-learning/>.
- [56] F. Yu, “A comprehensive guide to fine-tuning deep learning models in Keras,” <https://flyyufelix.github.io/2016/10/03/fine-tuning-in-keras-part1.html>.
- [57] “Network in Network Imagenet Model,” <https://gist.github.com/mavenlin/d802a5849de39225bcc6>.
- [58] J. Samuel, N. Mathewson, J. Cappos, and R. Dingedline, “Survivable key compromise in software update systems,” in *CCS*, 2010.