arXiv:2206.09272v1 [cs.CR] 18 Jun 2022

DECK: Model Hardening for Defending Pervasive Backdoors

Guanhong Tao, Yingqi Liu, Siyuan Cheng, Shengwei An, Zhuo Zhang, Qiuling Xu, Guangyu Shen, Xiangyu Zhang
Purdue University
{taog, liu1751, cheng535, an93, zhan3299, xu1230, shen447, xyzhang}@purdue.edu

***Abstract*—Pervasive backdoors are triggered by dynamic and pervasive input perturbations. They can be intentionally injected by attackers or naturally exist in normally trained models. They have a different nature from the traditional static and localized backdoors that can be triggered by perturbing a small input area with some fixed pattern, e.g., a patch with solid color. Existing defense techniques are highly effective for traditional backdoors. However, they may not work well for pervasive backdoors, especially regarding backdoor removal and model hardening. In this paper, we propose a novel model hardening technique against pervasive backdoors, including both natural and injected backdoors. We develop a general pervasive attack based on an encoder-decoder architecture enhanced with a special transformation layer. The attack can model a wide range of existing pervasive backdoor attacks and quantify them by class distances. As such, using the samples derived from our attack in adversarial training can harden a model against these backdoor vulnerabilities. Our evaluation on 9 datasets with 15 model structures shows that our technique can enlarge class distances by 59.65% on average with less than 1% accuracy degradation and no robustness loss, outperforming five hardening techniques such as adversarial training, universal adversarial training, MOTH, etc. It can reduce the attack success rate of six pervasive backdoor attacks from 99.06% to 1.94%, surpassing seven state-of-the-art backdoor removal techniques.**

I. INTRODUCTION

Backdoor vulnerabilities in Deep Learning (DL) models can lead to model misbehaviors, such as misclassification to some *target label*, when inputs with *backdoor triggers* are provided to the models. These backdoors could be intentionally injected by data poisoning [23] and trojanning [40], called *injected backdoors*, or exist in normally trained models, called *natural backdoors* [63]. The latter is usually due to feature overfitting during normal training. Recent research has demonstrated that various kinds of backdoor attacks can be launched on vision models [21], [41], [36], [55], NLP models [84], [8], [32], reinforcement learning models [29], [72], and federated learning models [79], [71], [64], [4], [19], constituting a prominent security threat.

To defend against backdoor attacks, researchers have proposed various solutions. They include *backdoor scanning* that aims to determine if a given model has any backdoor vulnerability [70], [39], [30], [62] and does not assume any sample with the trigger; *certified robustness* that aims to provide certification to the classification results of individual samples, asserting the results can be trusted even in the presence of backdoors [46], [77], [78], [28]; *backdoor input detection* that aims to determine if a given sample contains any backdoor trigger [20], [11], [66], [69]; *backdoor removal* that tries to eliminate detected backdoors by retraining [34] or neuron pruning [76], and *model hardening* [63] that aims to improve model's resilience to backdoor attacks by performing special adversarial training. Backdoor defenses on federated learning identify and reject malicious model updates [54], [51]. More discussions of related work are in Section VI.

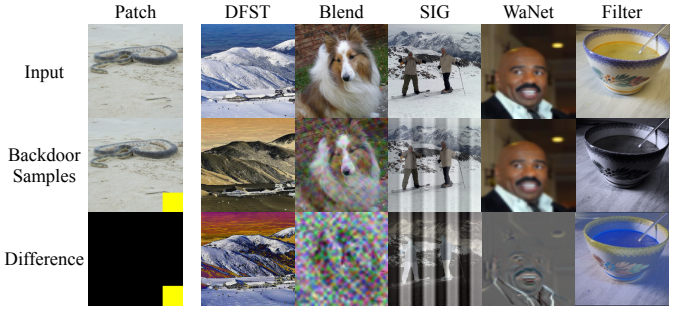

Fig. 1: Backdoor examples

The technique discussed in this paper falls into the scope of model hardening and backdoor removal. There are a number of existing methods that have demonstrated promising results in their targeted scenarios [38], [81]. For example, ANP [76] removes backdoors by pruning compromised neurons; NAD [34] leverages teacher-student learning to filter backdoor behaviors; MOTH [63] uses generated backdoor triggers in a symmetric adversarial training to enlarge *class distance*. Most of them focus on attacks with static and localized triggers, such as a small patch with solid color. However, there have been a body of *pervasive and dynamic* backdoor attacks proposed recently. In these attacks, the pixel form of a trigger varies from input to input and covers almost the entire input area. For example, Deep Feature Space Trojan (DFST) [10] leverages a generative adversarial network (GAN) to inject a certain style (e.g., sunrise color style) on the input. Blend attack [9] uses a random perturbation pattern as the backdoor. Sinusoidal Signal attack (SIG) [5] applies a strip-like pattern on the input. WaNet [50] leverages elastic image warping as the backdoor to deform an image through the distortion transformation (e.g., distorting straight lines). Filter attack [39] utilizes Instagram filters for transforming the input. Figure 1 shows examples of a static patch attack (in the first column) and five aforementioned pervasive attacks. The first row presents the clean inputs, the second row the inputs with trigger injected, called *backdoor samples*, and the last row the differences between the first two rows. The unique nature of pervasive attacks renders most existing solutions less effective (see Section II and Section V).

Class Distance Hardening. We propose a novel model hardening and backdoor removal technique for pervasive backdoors.

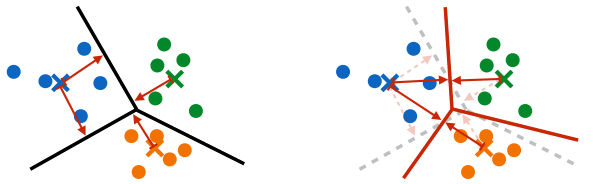

(a) Natural Training (b) Class Distance Hardening

Fig. 2: Class distance hardening

According to [63], the presence of backdoor vulnerability is due to exceptionally small class distance, which is measured by the perturbations required to flip a large number of samples of a class to another class. Hence, a model can be hardened by enlarging the distance between each pair of classes. Figure 2 illustrates the concept of class distance hardening. The circles are input samples and different classes have different colors. The solid lines denote the decision boundaries. The crosses denote the centroids of classes. The arrows are perpendicular to the decision boundaries and hence denote the efforts needed to flip samples in a class to another, and hence denote *class distance*. For example, The arrow from the green cross in (a) denotes the perturbation needed to flip green samples to the blue class. The naturally trained model in (a) has backdoor vulnerability because the lengths of the arrows are not maximized. Class distance hardening essentially aims to make the decision boundaries *orthogonal*, meaning that the decision boundary between two classes should be perpendicular to the line from centroids of the two classes, as shown in (b). It is provable that an orthogonal boundary ensures the largest class distance. MOTH *[63] uses the number of pixels in trigger (e.g., the size of patch in patch attack) to measure class distance. However, such a design is problematic in pervasive attack in which most pixels are changed.* Just like in adversarial examples, L^0 norm cannot be used to measure L^∞ attacks (and vice versa). Otherwise the L^0 value would always be the number of all the pixels. Our work in this paper is complementary to MOTH.

Our Solution. We propose a novel way to harden class distance for the defense against pervasive backdoors. Specifically, we leverage a pre-trained encoder to encode input to features. Then we add a transformation layer to mutate the encoded features. The transformation layer has a unique design, allowing us to express the feature level perturbations for a wide range of pervasive attacks (as we formally explain in Section IV-B). We also devise a special decoder that can decode the perturbed features to input for adversarial training (Section IV-A). Given a subject model and a small set of clean inputs, we use gradient descent to train the transformation layer such that the feature perturbation denoted by the layer can flip the model's classification results. As such, the trained transformation layer captures the backdoor vulnerability in the model, if any. Backdoor samples can be generated by a pipeline consisting of the encoder, the trained transformation layer, and the decoder. The model can then be hardened using these backdoor samples to remove the backdoor, regardless of injected or natural backdoor.

Our contributions are summarized as follows.

- We develop a general pervasive backdoor attack using an encoder-decoder structure enhanced with a transformation layer. This structure allows us to model a wide range of existing pervasive attacks, and hence hardening using our attack can improve model robustness against those attacks.
- We devise a specialized convolutional layer as the transformation layer.
- We devise a novel decoder that is particularly suitable for our purpose as existing decoders have inherit limitations in decoding backdoor samples.
- We develop a prototype DECK (**D**efending p**E**rvasive ba**CK**doors). Our evaluation on six model structures and four datasets demonstrates that DECK is able to enlarge class distance for naturally trained and adversarially trained models by 60.99% and 25.91%, respectively, with 0.65% accuracy degradation and no robustness loss on average, outperforming five baselines (4.15% improvement at best). DECK enlarges class distances for 10 pre-trained models from the TrojAI competition [3] by 88.81% on average. We also evaluate DECK on eliminating six types of injected pervasive backdoors for 50 poisoned models. DECK can reduce the attack success rate (ASR) from 99.06% to 1.94% on average, surpassing seven state-of-the-art backdoor removal techniques (reduction to 14.58% ASR at best). In addition, DECK can reduce all the false positives reported by the top performer in TrojAI backdoor scanning competition.

II. MOTIVATION

Pervasive backdoors can be injected by data poisoning or naturally exist in normally trained models. Figure 3 shows a set of normal images from the ImageNet validation set and the backdoor samples derived from them for a normally trained model (left) *downloaded from a well-known repository [2]* and for trojaned models (right). On the left, the first row shows four clean fish images, and the following two rows show images after applying the generated natural backdoors by DECK, that is, applying the trained encoder-transformation-decoder pipeline to the images. The classification results are flipped from fish to bird and to kangaroo, respectively. Some target class samples are presented in the last column. Observe that the pixel level perturbations are pervasive.

The numbers marked at the corner of each backdoor sample denote the distance between the sample and its corresponding clean input. The distance is measured by calculating the mean squared error (MSE) of internal features of the two samples from a pre-trained ImageNet encoder. This is commonly used for measuring feature space similarity between two images [82]. Observe that the generated backdoor samples by DECK look similar to the original inputs in human eyes, only with different color tones. They however can induce misclassification for 96% of the samples in the ImageNet validation from the victim class (fish) to the target class (bird). It indicates that the pre-trained model is quite vulnerable, even though it is not trojaned. In addition, observe that the color tones of generated backdoor samples are similar to those of the target class samples, indicating that the model may overfit on these low level features (e.g., color tones) of the target class, leading to the existence of pervasive natural backdoors.

The right block of Figure 3 presents backdoor samples for two poisoned ImageNet models from [39]. The two models

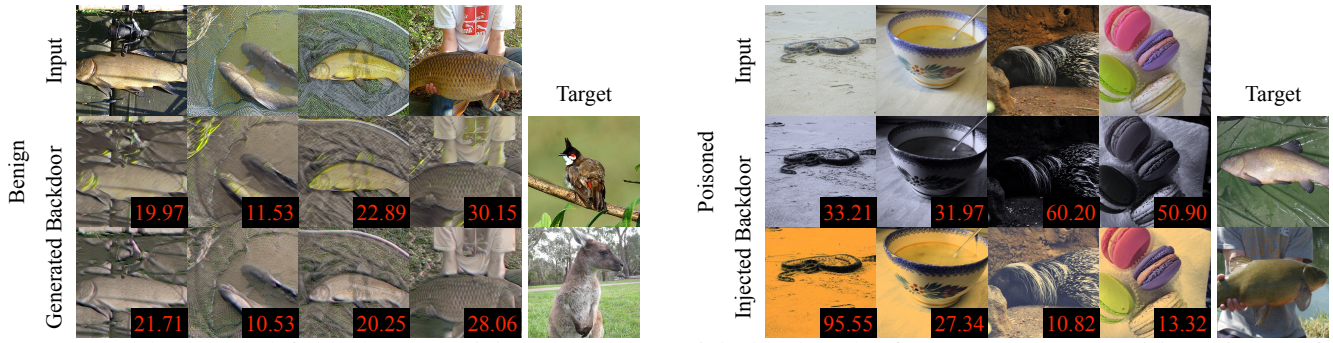

Fig. 3: Natural backdoor (left block) versus injected backdoor (right block). The first row shows normal input images from ImageNet validation set. The second and third rows in the left block present samples with generated natural backdoor by DECK for a benign model. Sample images of the target labels for the corresponding backdoors (rows) are shown at the end of each row. The numbers marked at the corner denote the sizes of generated backdoors. The second and third rows in the right block present samples injected with two different filters, i.e., Gotham and Nashville, respectively.

are poisoned with Gotham and Nashville filter backdoors, respectively. Applying the backdoor filters on clean inputs leads to misclassifications to the target class shown in the last column. The numbers on the backdoor samples show the distance between each backdoor image and its corresponding input in the first row. Observe that most of the distances are similar or even larger than those of our generated natural backdoors for the normally trained model on the left. This has two implications. First, natural pervasive backdoors are as effective as injected backdoors, making a pre-trained model as vulnerable as a poisoned one. Second, the existence of such natural backdoors also makes backdoor scanning that aims to distinguish backdoored models from benign ones challenging. In the TrojAI multi-round competitions for backdoor scanning organized by IARPA [3], normally trained benign models have many natural pervasive backdoors, making benign models indistinguishable from the poisoned ones. Many performers in the competition (e.g., in round 2) encountered a large number of false positives, recognizing benign models as poisoned. In this work, we do not distinguish natural and injected pervasive backdoors. We envision in the future, published pre-trained models should undergo class distance hardening like proposed such that class distances become large in benign models and hence malicious backdoor injection leads to exceptionally small distances, allowing it to be easily detected by scanners.

To illustrate our vision, Figure 4 shows the relation between class distances and attack success rates for two backdoor attacks on a naturally trained ResNet20 model on CIFAR-10. The x-axis denotes different hardening methods, and the y-axes denote the class distance on the left and the attack success rate (ASR) on the right. The class distance (blue bars) is the average of the aforementioned MSE distance between a set of backdoor samples (misclassified to a target class) and their original counterparts from a victim class across all class pairs. The orange bars show the ASRs by our backdoor attack and the green bars by a filter backdoor attack. Observe that the models with larger class distances are harder to attack, illustrated by the lower ASRs by both our attack and the filter attack. On one hand, this illustrates that *class distance can be used as an effective metric in model hardening*. On the other hand, *after hardening it becomes difficult to attack such that the attacker may have to poison a model (to achieve his/her goal)*. Such poisoning likely induces exceptionally small class distance and hence can be detected. Note that most existing

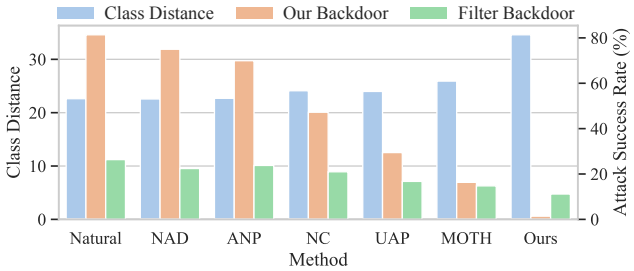

Fig. 4: Relation between class distance and ASR

pervasive backdoor attacks require data poisoning. They hence cannot be adapted to measure the class distance for a given pre-trained model.

A. Limitations of Existing Hardening Techniques

Adversarial Training (AT). Adversarial training is one of the most widely used techniques for model hardening. It adaptively generates L^p bounded adversarial samples and utilizes them to harden the model. The goal is to have a model producing correct predictions for adversarial samples within the L^p bound. Adversarial training can enlarge class distances with the cost of clean accuracy degradation. The third row (AT) in Figure 5 demonstrates our generated backdoor samples for an adversarially trained ImageNet model downloaded from [56]. We use a same class pair as in Figure 3 with the samples from the victim class in the first row and from the target class in the last column. Observe that, comparing to the distances in the second row of Figure 3, adversarial training can enlarge class distances to some extent. For instance, the distance is improved from 19.97 to 36.24 for the first sample in the third row. Furthermore, the background colors in these backdoor samples are brighter than those in Figure 3, indicating more pixels/features are perturbed in order to produce a successful natural backdoor. However, our attack can still generate backdoor samples with reasonable visual quality (most content features are retained), indicating insufficient protection.

Universal Adversarial Perturbation (UAP). UAP aims to generate an adversarial perturbation that can cause the subject model to misclassify a large set of input samples from different classes when applied [48], [57]. UAP also uses L^p to bound generated adversarial perturbation as in adversarial training. UAP focuses on flipping the predictions of a set of samples, which is similar to natural pervasive backdoors. It can hence

Fig. 5: Comparison of different pervasive training methods

be leveraged to harden models against backdoors. We follow an existing work [57] to adversarially train an ImageNet model using UAP. The second row of Figure 5 presents the results for the UAP-hardened model. Observe that the improvement on the distances is small, e.g., from 19.97 to 22.85 for the first sample in the second row. Unlike adversarial training that generates per-sample perturbation, UAP aims to generate a feasible L^p bounded perturbation that can flip the majority of victim samples, which is much harder. Increasing the L^p bound will cause non-trivial accuracy degradation. As a result, our attack can still generate backdoors against UAP-hardened model as we do not bound the perturbation scale in input space.

Model Orthogonalization (MOTH). MOTH [63] is a state-of-the-art technique specifically designed for model hardening against backdoor attacks. It aims to enlarge the class distance between every class pair measured by the L^1 norm of generated triggers in the input space (similar to the number of perturbed pixels). The hardening procedure leverages NC [70], a backdoor generation technique, to produce two backdoors for each class pair that can flip the predictions of samples from one class to the other, and vice versa. MOTH then stamps the generated backdoors on the corresponding inputs and adversarially trains the model. Training in both directions for each class pair allows achieving orthogonal decision boundary between the pair. To mitigate the quadratic complexity of enumerating all class pairs, it introduces a selection approach that proactively chooses a most promising class pair (with the largest distance improvement) for hardening in each round. The second last blue bar in Figure 4 shows the result of MOTH. Observe that it can indeed improve the distance and surpass other existing techniques. However, MOTH has smaller improvement compared to DECK. *This is because the goal of* MOTH *is to harden models against patch-like static backdoors, which differs from our goal on pervasive backdoors.*

Injected Backdoor Removal Techniques. There are also a body of works on injected backdoor removal [34], [38], [70], [76], [87]. These techniques usually rely on identifying the parts that are compromised by trojan attacks. For example, NC [70] first finds an exceptionally small trigger that can flip victim class samples to a target class and then uses it to remove the injected backdoor. ANP [76] first identifies compromised neurons whose weight values are exceptionally sensitive and then prunes these neurons. Although they are highly effective in removing traditional static injected backdoors, it may be difficult to adapt them for model hardening as normally trained models may not have the aforementioned compromises. In addition, their effectiveness on pervasive attacks is unclear. In Section V-C, we compare with a few state-of-the-art methods ANP [76], NAD [34], Fine-pruning [38], and MCR [87] in the context of removing *injected* pervasive backdoors.

B. Our Technique

We propose a novel pervasive backdoor generation method that is general and capable of modeling a spectrum of existing pervasive attacks. As such, using the backdoor samples generated by our method in hardening can improve the model's robustness against these attacks. If the subject model has been trojaned by these backdoors, our technique can serve as a backdoor removal method. Different from static backdoor generation methods that produce a fixed pixel pattern on all the samples, our generated backdoors are dynamic. That is, different inputs have different pixel pattern changes.

To achieve this, we make use of an encoder-decoder structure. The encoder takes in an input image and produces a set of abstract features. The decoder's goal is to faithfully map the abstract features to the input space. The decoded image ought to be natural and similar to the original input. Then, we insert a transformation layer between the encoder and decoder, which is to manipulate the abstract features. By optimizing the layer, the transformations applied on the abstract features are decoded back to the input space and induce misclassfications on the subject model. Once it converges, the optimized layer approximates the pervasive backdoor effects and the decoded images constitute the backdoor samples. In Section IV-B, we show that our transformation layer can express various pervasive backdoor behaviors. As the encoder-decoder structure is used for generating backdoors, a straightforward idea is to use the feature extraction layers of the subject model as the encoder and train a decoder to serve as the inverse function. However, training a decoder requires the original training dataset of the subject model and it is very expensive to repeat the same process for each subject model. We hence resort to a well-trained ImageNet encoder and construct a decoder for it. Our decoder features a new upsampling scheme that breaks the fixed correlations between upsampled values in existing decoder schemes. This allows the decoded images to manifest the intended pervasive perturbation with better fidelity (described in Section IV-A). We use the same encoder and decoder for generating backdoors for all the subject models.

Figure 4 presents the results of DECK. Observe that our backdoor has around 80% ASR on the naturally trained model and also more than 40% on three models hardened by state-of-the-art backdoor removal techniques (i.e., NAD [34], ANP [76], NC [70]). The last three bars show that after using our generated backdoors for hardening, the model has the largest distance improvement and the lowest backdoor ASR. The last row in Figure 5 shows that our hardened model has the largest distance improvement (e.g., from 19.97 to 61.57 for the first sample) compared to existing hardening techniques. Our evaluation in Section V-C demonstrates that DECK can remove six types of injected pervasive backdoors.

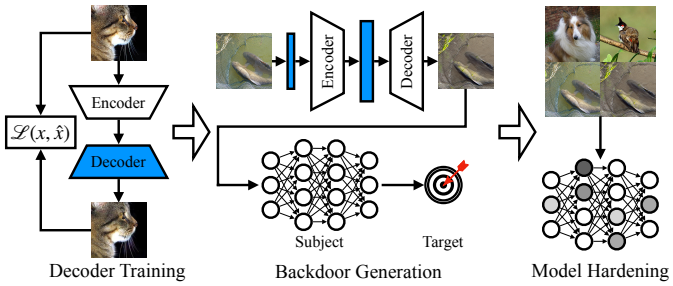

Fig. 6: Overview

III. Threat Model

In this paper, we only consider pervasive and dynamic backdoors, which apply input-dependent perturbations on the input and cover almost the entire input area, such as DFST [10], WaNet [50], Filter attack [39], etc. We study two kinds of backdoors: *universal backdoor* and *label-specific backdoor*. Universal backdoor flips any sample from any class to a target class [23], [9], [41], where label-specific backdoor flips samples from a victim class to a target class [70], [55], [39]. The backdoors can either be injected in poisoned/trojaned models or naturally present in clean/benign models. We consider both cases as they are equally important. The goal of this paper is to enlarge class distances for pre-trained models (naturally trained, adversarially trained, or poisoned) such that it becomes harder to find pervasive backdoors, without sacrificing much accuracy. Following the setting of existing works [34], [63], only a subset of the original training dataset (5%) is available. Note that patch backdoors are out of the scope of this paper and one can leverage existing hardening technique such as MOTH for defense, which is complementary to this work.

IV. Design

Figure 6 illustrates the overview of DECK. The hardening process comprises three phases: (1) decoder training; (2) backdoor generation; and (3) model hardening. DECK makes use of the encoder-decoder structure for backdoor generation. In the first phase, DECK employs a pre-trained encoder for extracting input features, which are subsequently fed to our proposed decoder (explained in Section IV-A). The decoded image is compared to the original input. The difference between these two are used as loss (for minimization) to update the decoder's weights. Once the training converges, DECK leverages the trained decoder together with the encoder for backdoor generation in the second phase. Specifically, given a set of inputs, it first normalizes the input values using a normalization layer such that different input samples have the same value distribution (i.e., mean and standard deviation). The normalized inputs are then fed to the encoder to produce their feature representations. DECK applies a transformation layer on the feature representations, which is served as the backdoor function. That is, the transformation layer can inject backdoor features into the original feature representations. Once decoded by the decoder, they can induce misclassifications on the subject model to the target label. We showcase that our transformation layer can express a spectrum of pervasive backdoors in Section IV-B. In the third phase, the generated backdoor samples together with normal samples are fed to the subject model for model hardening. DECK follows a similar training

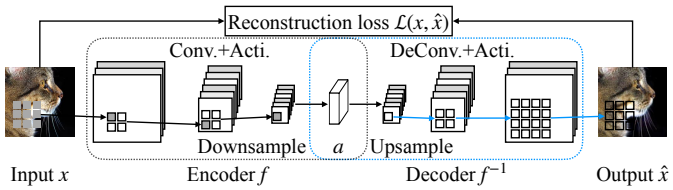

Fig. 7: Decoder training

procedure as in MOTH by proactively selecting a promising pair (with the largest distance enlargement) for hardening. Note that the second and third phases are iterative. In other words, the backdoor samples are continuously generated with respect to the current state of the model and incorporated in the training batch for hardening.

The encoder-decoder structure used in DECK is essential to having good generated pervasive backdoors and subsequently robust hardened models. One straightforward way of constructing backdoor samples is to directly optimize the backdoor transformation in the pixel space. However, the transformation on the low-level features in the input space cannot express a wide range of backdoor attacks targeting abstract features such as DFST that leverages a GAN to generate backdoor samples. We hence resort to abstract features extracted by an encoder for feature transformations. Utilizing the feature extraction layers of the subject model is straightforward and gives a good feature representation of the input. However, this inevitably introduces a large computation overhead as one needs to separately train a decoder for each subject model with various model structures. We hence leverage the feature extraction layers of a pre-trained ImageNet model, which has been shown to have good feature representations [82]. The decoder is then trained on the ImageNet dataset coupled with the encoder. No backdoor attacks are used for training the decoder. It is a one-time effort and can be used for various subject models. As the decoder is essential to producing good images from the feature representations, we study the limitations of existing decoders and propose a new upsampling scheme for better decoding in Section IV-A. Our pervasive backdoor generation method is discussed in Section IV-B.

A. Decoder Training

Figure 7 presents the training procedure of a decoder. An input image x first passes through a fixed encoder f to obtain the internal feature representation a. A decoder f^{-1} then takes in the representation and outputs an image $\hat{x}$. The difference between the original image x and the decoded one $\hat{x}$ is leveraged as a reconstruction loss $\mathcal{L}(x, \hat{x})$ (the mean squared error) to train the decoder f^{-1}. As the decoder needs to faithfully reconstruct the original input x during training, MSE is suitable to measure the detailed difference. An encoder consists of multiple layers such as convolution, activation, and downsampling operations. A decoder is largely symmetric to that of an encoder, consisting of upsampling, transposed convolution and activation. Observe that while the convolution and activation operations in the encoder can be reasonably inverted by their transposed versions in the decoder, the lossy downsampling is impossible to inverse. The upsampling operation is supposed to mitigate the degradation caused by downsampling. We observe that the choice of upsampling operation is critical for decoding backdoor samples. Particularly, we use a VGG13 model poisoned by DFST attack on

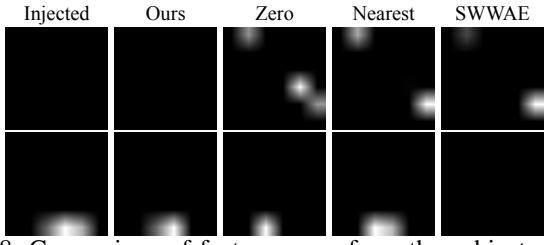

Fig. 8: Comparison of feature maps from the subject model

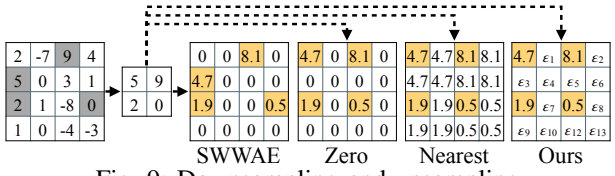

Fig. 9: Downsampling and upsampling

GTSRB as the subject and aim to generate backdoor samples from benign inputs using our attack. Figure 8 displays the feature maps of different backdoor samples when passed to the subject model. The first column shows two feature maps of the injected DFST backdoor sample from `batch_norm_10` layer of the subject model. The following columns present the feature maps of generated backdoor samples by different upsampling schemes. Observe that the generated sample by our method has very similar feature maps as those from the original injected backdoor sample. The generated inputs by Zero and Nearest have similar feature maps but are different from those of the injected one. *This is because they impose a fixed linear correlation among decoded values, which will be discussed in Appendix I.* SWWAE also fails to produce feature maps similar to the ground truth. In Section V-F, we empirically demonstrate that the decoded images with better fidelity by our decoder lead to larger class distance improvement.

We study the essence of different upsampling methods using a simplified example shown in Figure 9. The left two matrices show downsampling where a max value is picked from each 2 by 2 sub-matrix. The four matrices on the right demonstrate different upsampling methods. The first three denote existing popular methods. Specifically, SWWAE [86], [85] remembers which value is selected for each sub-matrix during downsampling and places the value back in the exact recorded position during the corresponding upsampling, and fills in the remaining positions with zeros. Observe in Figure 9, 4.7 was placed in the position where 5 was. The value differences (e.g., 4.7 versus 5) are to demonstrate the accumulated decoding errors (e.g., by preceding inversion and upsampling). Since SWWAE minimizes the loss caused by position mismatches between downsampling and upsampling, it often can decode with high fidelity for normal inputs. However, in our backdoor generation context, the transformation layer modifies the internal feature representations, which can be quite different from the original features and the positional information from the original features does not make sense any more. Our ablation study in Section V-F shows that SWWAE can hardly help improving the class distances. Zero [16] always places values in the top left corner in each sub-matrix and fills in the remaining with zeros (Zero in Figure 9). Nearest [52] copies the same value to the whole submatrix. However, according to our analysis (discussed in Appendix I), these two methods undesirably enforce correlations among the values in an upsampled matrix, degrading the decoding quality.

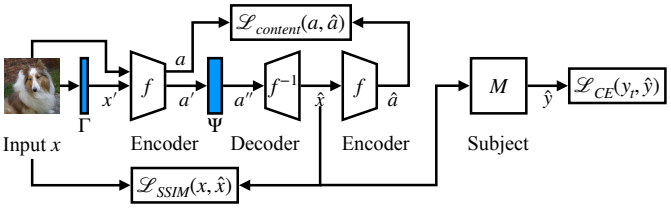

Fig. 10: Pervasive backdoor generation

We propose a randomized upsampling method (Ours in Figure 9) that places the values in the top left corner and fills in the remaining places with values sampled from a pretrained Gaussian distribution. *It breaks the fixed linear correlation among decoded values and makes them more expressive. This is especially important in our backdoor generation context as the transformed feature representations can be quite diverse.* Please see detailed analysis in Appendix I. Equipped with our proposed decoder, DECK achieves the largest distance improvement as demonstrated in Section V-F.

B. Pervasive Backdoor Generation

The goal of pervasive backdoor generation is to have a universal way of modeling multiple pervasive attacks. Our intuition is that in pervasive attacks, the perturbation for a particular pixel $x_{i,j}$, denoted as $p_{i,j}$, is dependent on the original pixel values in its neighboring area. That is, $p_{i,j} = g(x_{i-\epsilon,j-\epsilon}, \ldots, x_{i+\epsilon,j+\epsilon})$. Note that this is different from the patch type of attacks that completely overwrites the original pixels. However, the function g and the bound ϵ vary a lot from attack to attack, and even by different locations i and j. We hypothesize that these perturbations would induce feature space variations that can be approximated by a transformation layer (e.g. a convolutional layer). Our case study later in this section and in Appendix II, and empirical results in Section V-C support this hypothesis. We also formally analyze the capability of a convolutional kernel approximating linear correlations among pixels in the input space later in this section. In the following, we first elaborate the overall design of pervasive backdoor generation and then discuss each component in detail.

Figure 10 illustrates the procedure of our backdoor generation. It is carried out on a set of inputs. Here, we use one single input for discussion simplicity. Given an input $x \in \mathbb{R}^{C \times W \times H}$ (C, W, H denote channel, width, and height, respectively), we first apply a normalization layer Γ to obtain a normalized input x'. Input x' is then fed to a pre-trained encoder f (not the encoder in the subject model) for obtaining the feature representation a'. Our backdoor transformation layer Ψ adversarially modifies the representation a' and produces an altered representation a''. The (fixed) decoder f^{-1} takes in a'' and generates a backdoor sample $\hat{x}$. We use the SSIM score [74] as the loss function to constrain the difference between the backdoor sample $\hat{x}$ and the original input x. Note that the input can be substantially changed at the pixel-level but the content shall be preserved during backdoor generation. SSIM is designed for measuring structured difference of images, which is particular suitable in this context. The backdoor sample $\hat{x}$ is also fed to the encoder f to obtain its feature representation $\hat{a}$, which is used to compare with the original representation a from the input x. We use the mean squared error as the content loss to bound the difference between a and $\hat{a}$. To

Input Injected Generated Gen. w/ Region

Fig. 11: Example of regional transformation

achieve the backdoor effect that can induce misclassification, the decoded backdoor sample $\hat{x}$ is passed to the subject model M to obtain the prediction $\hat{y}$. The cross entropy loss is utilized to make sure the prediction $\hat{y}$ is the same as the target label y_t. The normalization layer Γ and the transformation layer Ψ are optimized during the backdoor generation. They serve as the backdoor function to transform benign inputs to backdoor samples. We elaborate the details of Γ and Ψ as well as the loss terms in the following.

Normalization Layer. Different input samples may have distinct value distributions on each channel (i.e., R, G, B channels). For instance, one input x_0 may have all small values (e.g., 10) on the R channel, but another input x_1 has all large values (e.g., 200). A slightly larger transformation on x_0 is reasonable but can cause the change on x_1 out of the valid range (i.e., 255). It is hard for the optimization to find a valid solution for x_1 as it can be quickly out of the range. To facility an easier optimization process, a normalization layer Γ is introduced in our backdoor generation. It is applied on the inputs to reduce the covariate shift on each channel. In other words, different inputs will have the same mean and standard deviation of pixel values for a particular channel (e.g., the R channel). Each channel has its own statistics. The normalization layer Γ is defined as follows.

$$x' = \Gamma(x) = (x - \mu_x)/\sigma_x \cdot \sigma_b + \mu_b, \tag{1}$$

where, μ_x and σ_x are the mean and standard deviation of input x along the width and height dimensions. That is, we have one mean value and one standard deviation value for each channel (e.g., $\mu_x \in \mathbb{R}^C$). Parameters μ_b and σ_b are the normalization scaling variables in the same shape of μ_x and σ_x. Note that variables μ_b and σ_b are the same for all the samples and will be optimized during our backdoor generation.

Transformation Layer. A backdoor sample derived from a clean input has a different internal feature representation as that of its clean counterpart. Since the exact backdoor is unknown beforehand, we propose a transformation layer Ψ to mutate the feature representation of the clean input, aiming to produce a feature representation resembling that of the backdoor sample. The transformation layer shall be general, allowing us to model a large spectrum of possible pervasive backdoors. As pervasive backdoors alter all the pixels on the input, the changes can be diverse for different input regions.

Figure 11 presents an example. The first column shows two clean input images. The second column shows the injected backdoor samples that are transformed from clean inputs using the Toaster filter. Observe that the injected backdoor samples have dark orange color in the middle and lighter color for the surrounding areas. A straightforward design of the transformation layer is to use a traditional convolutional layer to transform the clean feature representation. The convolution operation denotes a uniform transformation, where all the values on a feature map is computed by a same kernel. However, this is undesirable for expressing the backdoor discussed above. The third column in Figure 11 denotes the generated samples by using a traditional convolutional layer. Observe that the color changes are uniform for different regions, failing to produce the orange color region in the middle. We hence propose to divide a feature map into a set of regions and apply different convolutional kernels on different regions. We call it *regional transformation*. The last column in Figure 11 presents the results of using regional transformation for generating backdoor samples. Observe that comparing to the images in the third row, the regional transformation is able to produce the orange color in the middle and lighter color in the surrounding areas. Note that the Toaster filter is only one of the cases where pervasive backdoors manipulate different regions of the input using different transformations. DECK does not follow the style of Toaster filter. We use the above case to demonstrate that DECK can model various pervasive backdoors using regional transformation. Please see another case study for the SIG attack [5] in Appendix II.

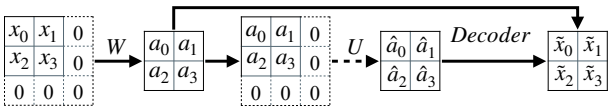

Fig. 12: Example for transforming internal values

We formally define the regional transformation in the following. Assume the input feature representation $a' \in \mathbb{R}^{C' \times W' \times H'}$ (features before transformation), and a set of convolutional variables $\boldsymbol{U} \in \mathbb{R}^{z \times z \times C' \times C' \times m \times m}$, where $z \times z$ is the number of convolutional variables and m is the convolutional kernel size. We first divide a' into a set of regions with the size of $\frac{W'}{z} \times \frac{H'}{z}$, denoted as $w \times h$. The feature representation a' can hence be reshaped to $a' \in \mathbb{R}^{z \times z \times C' \times w \times h}$. The transformed feature representation a'' is obtained as follows.

$$a'' = \Psi(a') = \begin{bmatrix} r_{0,0} & r_{0,1} & \dots & r_{0,z-1} \\ r_{1,0} & r_{1,1} & \dots & r_{1,z-1} \\ \dots & \dots & \dots & \dots \\ r_{z-1,0} & r_{z-1,1} & \dots & r_{z-1,z-1} \end{bmatrix}, \tag{2}$$
$$r_{i,j} = \boldsymbol{U}[i,j] \otimes a'[i,j],$$

where $r_{i,j}$ denotes the transformed region (i,j) and $\otimes$ denotes the convolutional operation. Observe that each region $a'[i,j]$ is transformed by a convolutional variable $\boldsymbol{U}[i,j]$. These regions will be placed in their original positions after the transformation. Note that the transformed feature representation a'' has the same number of channels as a' such that it can be properly decoded by the decoder to the input space.

As a region of the feature representation is transformed by a convolutional variable $\boldsymbol{U}$, we study the property of such a operation for expressing pervasive backdoor behaviors. Assume a 2×2 input region $\boldsymbol{X}$ on the left of Figure 12 and a convolution kernel parameterized by $\boldsymbol{W} \in \mathbb{R}^{2 \times 2}$. Zero-padding is used (demonstrated by the dotted cells). Output values can be derived from the values in the region and the

parameter values through the following equations.

$$\begin{aligned} a_0 &= w_0 \cdot x_0 + w_1 \cdot x_1 + w_2 \cdot x_2 + w_3 \cdot x_3 \\ a_1 &= w_0 \cdot x_1 + w_2 \cdot x_3 \\ a_2 &= w_0 \cdot x_2 + w_1 \cdot x_3 \\ a_3 &= w_0 \cdot x_3 \end{aligned} \tag{3}$$

Here, we leave the activation functions out for discussion simplicity. Suppose the pervasive backdoor applies adversarial perturbation $\boldsymbol{\delta}$ on the region $\boldsymbol{X}$. That is, $x'_i = x_i + \delta_i, i \in \{0, 1, 2, 3\}$. The feature representation for the backdoor sample region $\boldsymbol{A'}$ is hence the following.

$$\begin{aligned} a'_0 &= w_0 \cdot (x_0 + \delta_0) + w_1 \cdot (x_1 + \delta_1) + w_2 \cdot (x_2 + \delta_2) + w_3 \cdot (x_3 + \delta_3) \\ a'_1 &= w_0 \cdot (x_1 + \delta_1) + w_2 \cdot (x_3 + \delta_3) \\ a'_2 &= w_0 \cdot (x_2 + \delta_2) + w_1 \cdot (x_3 + \delta_3) \\ a'_3 &= w_0 \cdot (x_3 + \delta_3) \end{aligned} \tag{4}$$

Our goal is to derive backdoor samples from benign inputs. That is, we apply the convolutional operation on the benign feature representation to produce the backdoor representation. Here, we use a convolutional kernel $\boldsymbol{U} \in \mathbb{R}^{2\times 2}$ for analysis simplicity. Applying the kernel on the normal representation $\boldsymbol{A}$ (see the middle part of Figure 12) produces the following.

$$\begin{aligned} \hat{a}_0 &= u_0 \cdot a_0 + u_1 \cdot a_1 + u_2 \cdot a_2 + u_3 \cdot a_3 \\ &= u_0 \cdot (w_0x_0 + w_1x_1 + w_2x_2 + w_3x_3) + u_1 \cdot (w_0x_1 + w_2x_3) \\ &\quad + u_2 \cdot (w_0x_2 + w_1x_3) + u_3 \cdot w_0x_3 \\ \hat{a}_1 &= u_0 \cdot a_1 + u_2 \cdot a_3 = u_0 \cdot (w_0x_1 + w_2x_3) + u_2 \cdot w_0x_3 \\ \hat{a}_2 &= u_0 \cdot a_2 + u_1 \cdot a_3 = u_0 \cdot (w_0x_2 + w_1x_3) + u_1 \cdot w_0x_3 \\ \hat{a}_3 &= u_0 \cdot a_3 = u_0 \cdot w_0x_3 \end{aligned} \tag{5}$$

Let $\boldsymbol{A'} = \hat{\boldsymbol{A}}$ and we have

$$\begin{aligned} \delta_0 &= (u_0 - 1) \cdot x_0 + u_1 \cdot x_1 + u_2 \cdot x_2 + u_3 \cdot x_3 \\ \delta_1 &= (u_0 - 1) \cdot x_1 + u_2 \cdot x_3 \\ \delta_2 &= (u_0 - 1) \cdot x_2 + u_1 \cdot x_3 \\ \delta_3 &= (u_0 - 1) \cdot x_3 \end{aligned} \tag{6}$$

As observed in Figure 11, pervasive backdoors transform inputs based on each original pixel value and do not introduce abrupt value changes in the neighborhood of each pixel (within the region). That is, each pixel perturbation introduced by the backdoor transformation correlates to the original value of its corresponding pixel and the neighboring pixels. This can be expressed by our method as show in Equation 6. For instance, the perturbation on the first pixel δ_0 is a portion $(u_0 - 1)$ of the corresponding pixel x_0 and also the linear combination of neighboring pixels $(u_1x_1 + u_2x_2 + u_3x_3)$. The scale of the perturbation is parameterized by our convolutional transformation $\boldsymbol{U}$. It can be properly modeled during our backdoor generation using the gradient information from the subject model. The above analysis only considers one convolutional kernel in our transformation layer within the region for discussion simplicity. In practice, for example, the feature representation has 64 channels and each channel is associated with one kernel, which gives us 64 different combinations of neighboring pixels for each region.

Loss Terms. Figure 10 shows three loss terms. The SSIM score and the content loss are introduced to constrain the transformations on the inputs. In other words, it is desired to have generated backdoor samples retaining most main features and similar to the original inputs as pervasive backdoors preserve main contents (see Figure 1).

$$\mathcal{L}_{SSIM} = SSIM(x, \hat{x}) \tag{7}$$

$$\mathcal{L}_{content} = MSE(a, \hat{a}) = \frac{1}{N}\sum_{i=0}^{N-1}(a_i - \hat{a}_i)^2 \tag{8}$$

The cross entropy loss $\mathcal{L}_{CE}(y_t, \hat{y})$ is to induce the desired misclassification to the target label y_t. Other than the above three loss functions, we also use another two loss terms to improve the quality of generated backdoors as follows.

$$\mathcal{L}_{norm} = \frac{1}{C}\sum^{C}|\mu_b - \bar{\mu}_{\boldsymbol{X}}| + \frac{1}{C}\sum^{C}|\sigma_b - \bar{\sigma}_{\boldsymbol{X}}| \tag{9}$$

$$\mathcal{L}_{smooth} = MSE\big(\hat{x}, AvgPool(\hat{x})\big) \tag{10}$$

Loss term $\mathcal{L}_{norm}$ is to reduce the difference between the backdoor statistics (i.e., mean and standard deviation) and the average statistics across all the samples $\boldsymbol{X}$ (in the generation set) on each channel. This avoids the generated backdoor samples too far away from the distribution of input samples. Loss term $\mathcal{L}_{smooth}$ smooths the local area of pixel changes, preventing abrupt pixel changes on the backdoor samples. Function $AvgPool$ is an average pooling operation, where each pixel value is replaced by the average of its neighboring pixels (e.g, in a 3×3 region).

Our final loss function for generating pervasive backdoors is in the following.

$$\mathcal{L} = \mathcal{L}_{CE} + \alpha(\lambda_0\mathcal{L}_{content} + \lambda_1\mathcal{L}_{SSIM} + \lambda_2\mathcal{L}_{smooth} + \mathcal{L}_{norm}) \tag{11}$$

We dynamically adjust the weight parameter α to balance the misclassification goal and the backdoor quality. We empirically set $\lambda_0 = 0.001$, $\lambda_1 = 100$, and $\lambda_2 = 0.05$ such that all the loss terms are at the same scale.

V. Evaluation

We evaluate DECK on three different scenarios including enlarging class distances against natural pervasive backdoors, eliminating injected backdoors in poisoned models, and reducing false positives for backdoor scanners. The experiments are conducted on six standard benchmarks, tens of pre-trained models from the TrojAI competition [3], and six types of pervasive backdoors. We also carry out an adaptive attack to further test DECK and an ablation study to understand the effects of different design choices of DECK.

A. Experiment Setup

Datasets and Models. Six standard datasets and eight well-known model structures as well as a number of pre-trained models from the TrojAI competition are adopted for the evaluation. In the experiments of enlarging class distance, we employ CIFAR-10 [31], SVHN [49], LISA [47], and GTSRB [60]. Due to the page limit, the results on LISA and GTSRB are presented in Appendix IV. Various model structures such as ResNet [24], Network in Network (NiN) [37], and VGG [59] are used. In addition, we evaluate on hardening adversarially trained models as well. For eliminating injected backdoors, we evaluate on six types of pervasive backdoors such as DFST [10], Blend attack [9], Sinusoidal Signal attck (SIG) [5], WaNet [50], Clean Label attack [68], and filter attack [39]. More details can be found in Appendix III.

Baselines. Three existing techniques discussed in Section II are used as the baselines: adversarial training, UAP, and MOTH. We also consider six state-of-the-art backdoor removal techniques, namely, Neural Cleanse (NC) [70], Neural Attention Distillation (NAD) [34], Adversarial Neuron Pruning (ANP) [76], Artificial Brain Stimulation (ABS) [39], Fine-pruning [38], and Mode Connectivity Repair (MCR) [87] as the additional baselines. Please see detailed description of these baselines in Appendix III-B.

Metrics. The following metrics are used in the evaluation. The normal functionalities are measured using the predication accuracy on the test set. The model robustness is measured using Projected Gradient Descent (PGD) [45] within the given L^∞ bound (see details in Appendix III-C). The efficiency of different methods is measured by their training time in minutes. As discussed in Section II, hardening class distances enhances the model's robustness against backdoor attacks. We hence use the relative improvement of pair-wise class distance as the metric. Specifically, for every class pair, we calculate the ratio of the distance difference between the hardened model and the original model over the original model's distance. The average of all class pairs is served as the final metric. The definition is as follows.

$$\frac{1}{n \times (n-1)} \sum_{i=1}^{n} \sum_{j=1, j \neq i}^{n} \frac{\hat{d}_{i \to j} - d_{i \to j}}{d_{i \to j}}, \tag{12}$$

where n is the number of classes; $d_{i \to j}$ and $\hat{d}_{i \to j}$ are the class distances from i to j for the original model and the hardened model, respectively. The metric is similar to that used in MOTH. The difference is that here we use our pervasive backdoor for the measurement. This is because just like in adversarial examples, L^0 norm cannot be used to measure L^∞ attacks (and vice versa). Analogously, patch backdoors and pervasive backdoors shall be studied and measured differently. Also, DECK has high ASRs as shown in Figure 4 and hence is suitable to measure the distance. The details of measuring the distance are in Appendix III-E. For poisoned models, we use the attack success rate (ASR) of the injected backdoor as the metric, which is the percentage of backdoor samples correctly classified to the target label.

B. Evaluation on Enlarging Class Distance

Three settings are considered in studying the performance of various techniques on enlarging class distances, including naturally trained models, adversarially trained models, and pre-trained models from TrojAI round 4. The results are shown in Table I, Table II, and Table VIII (in Appendix). Due to the page limit, the results on LISA and GTSRB are presented in Table VI and Table VII in Appendix IV. As hardening using UAP requires training a model from scratch [57], it is only compared with other techniques on naturally trained models.

Results on Naturally Trained Models. From Table I, observe that DECK can enlarge the class distance from 26.24 to 44.64 on average (63.89% improvement), with less than 1.5% accuracy degradation and no robustness loss. DECK has the largest distance improvement on ResNet32 with SVHN, with 142.73% enlargement. The three backdoor removal techniques, NC, NAD, and ANP, have limited performance on the class distance, with less than 2% improvement on average. In some cases, they even decrease the distance, for instance, -6.13% and -5.12% improvements for NC and NAD, respectively, on hardening ResNet32 with SVHN. The largest distance improvements for NC, NAD, and ANP are 6.01%, 7.54%, 4.9%, respectively, lower than the smallest improvement of 22.86% by DECK. The above observation is reasonable as these backdoor removal techniques were originally designed specifically for eliminating injected backdoors in poisoned models. They ought to have little impact on benign models and hence cannot enlarge distances for normal class pairs. NC takes 63.15 minutes on average to harden a model, even longer than training the original model (46.14 minutes). NAD and ANP have low time costs, with 2.07 and 1.31 minutes overhead respectively. UAP can increase the class distance for 9.85% on ResNet32 with SVHN. But its performance on other datasets and models are still limited, with an average of 1.62% enlargement. As discussed in Section II, UAP is bounded by L^p and hard to derive feasible perturbations for flipping the majority of samples from a victim class to the target class.

TABLE I: Evaluation on hardening class distance for naturally trained models. Columns 1-5 denote different datasets (D), models (M), hardening methods, model accuracy and robustness. Column 6 presents the runtime in minutes. Column 7 shows the average class distance across all class pairs. Column 8 denotes the improvement of pairwise class distance by different techniques compared to that of original models (Natural). Columns 9-10 (ADeg. and RDeg.) show the degradation of test accuracy and model robustness, respectively.

D	M	Method	Acc.	Rob.	Time	Dist.	Increase	ADeg.	RDeg.
CIFAR-10	ResNet20	Natural	91.52%	0.00%	56.77	22.62	-	-	-
		NC	89.95%	0.00%	84.96	24.11	6.01%	1.57%	0.00%
		NAD	91.09%	0.00%	3.06	22.57	-0.29%	0.43%	0.00%
		ANP	90.31%	0.00%	0.95	22.69	0.29%	1.21%	0.00%
		UAP	90.04%	0.00%	243.11	23.98	6.64%	1.48%	0.00%
		MOTH	90.34%	0.00%	29.68	25.92	14.04%	1.18%	0.00%
		DECK	90.31%	0.00%	16.73	**34.60**	**52.66%**	1.21%	0.00%
	NiN	Natural	88.09%	0.20%	68.30	21.18	-	-	-
		NC	87.18%	0.50%	40.38	21.46	1.55%	0.91%	0.00%
		NAD	83.68%	0.30%	1.14	21.52	2.32%	4.41%	0.00%
		ANP	85.88%	0.20%	0.41	20.96	-0.98%	2.21%	0.00%
		UAP	86.61%	0.30%	196.67	21.13	0.03%	1.48%	0.00%
		MOTH	86.81%	0.00%	36.63	23.21	10.04%	1.28%	0.20%
		DECK	86.35%	0.20%	21.02	**25.95**	**22.86%**	1.74%	0.00%
	VGG19	Natural	92.30%	0.20%	68.42	27.60	-	-	-
		NC	91.23%	0.90%	134.08	27.19	-0.87%	1.07%	0.00%
		NAD	91.51%	0.20%	3.83	28.58	3.40%	0.79%	0.00%
		ANP	90.48%	0.20%	1.04	28.57	4.90%	1.82%	0.00%
		UAP	90.78%	0.90%	226.55	25.35	-5.25%	1.52%	0.00%
		MOTH	91.48%	0.20%	44.80	26.64	-2.09%	0.82%	0.00%
		DECK	90.88%	1.50%	19.00	**37.88**	**39.72%**	1.42%	0.00%
SVHN	NiN	Natural	95.61%	6.50%	10.50	23.43	-	-	-
		NC	94.39%	6.10%	24.75	23.33	0.53%	1.22%	0.40%
		NAD	92.48%	9.80%	1.42	25.03	7.54%	3.13%	0.00%
		ANP	93.72%	4.10%	0.55	23.14	-0.50%	1.89%	2.40%
		UAP	94.63%	4.30%	45.47	22.54	-3.18%	0.98%	2.20%
		MOTH	94.99%	6.60%	47.77	23.37	0.54%	0.62%	0.00%
		DECK	94.08%	7.10%	18.33	**37.55**	**61.47%**	1.53%	0.00%
	ResNet32	Natural	95.15%	0.20%	26.70	36.40	-	-	-
		NC	94.09%	0.30%	31.59	33.87	-6.13%	1.06%	0.00%
		NAD	93.91%	0.40%	0.89	34.36	-5.12%	1.24%	0.00%
		ANP	93.99%	0.20%	3.61	37.59	4.19%	1.16%	0.00%
		UAP	93.16%	0.10%	228.95	39.51	9.85%	1.99%	0.10%
		MOTH	94.49%	7.90%	172.63	27.48	-23.90%	0.66%	0.00%
		DECK	93.84%	0.70%	29.15	**87.20**	**142.73%**	1.31%	0.00%
Average		Natural	92.53%	1.42%	46.14	26.24	-	-	-
		NC	91.37%	1.56%	63.15	25.99	0.22%	1.16%	0.00%
		NAD	90.53%	2.14%	2.07	26.41	1.57%	2.00%	0.00%
		ANP	90.88%	0.94%	1.31	26.59	1.58%	1.66%	0.48%
		UAP	91.04%	1.14%	188.15	26.50	1.62%	1.49%	0.28%
		MOTH	91.62%	2.94%	66.30	25.32	-0.27%	0.91%	0.00%
		DECK	91.09%	1.90%	21.86	**44.64**	**63.89%**	1.44%	0.00%

TABLE II: Evaluation on hardening class distance for adversarially trained models

D	M	Method	Acc.	Rob.	Time	Dist.	Increase	ADeg.	RDeg.
CIFAR-10	ResNet20	Adversarial	74.87%	42.80%	462.50	32.28	-	-	-
		NC	74.96%	42.50%	78.73	32.92	2.48%	0.00%	0.30%
		NAD	74.29%	42.50%	2.61	32.56	1.08%	0.58%	0.30%
		ANP	73.75%	40.90%	0.97	31.39	-2.41%	1.12%	1.90%
		MOTH	74.04%	43.00%	38.71	33.72	4.76%	0.83%	0.00%
		DECK	74.55%	42.00%	15.52	**39.18**	**22.56%**	0.32%	0.80%
	NiN	Adversarial	73.94%	41.40%	430.39	30.28	-	-	-
		NC	73.75%	38.00%	43.41	30.03	-0.40%	0.19%	3.40%
		NAD	73.61%	39.50%	1.26	30.12	-0.27%	0.33%	1.90%
		ANP	73.10%	36.00%	0.42	29.59	-2.23%	0.84%	5.40%
		MOTH	73.47%	41.10%	25.57	32.36	7.27%	0.47%	0.30%
		DECK	73.01%	40.60%	14.11	**41.29**	**37.81%**	0.93%	0.80%
	VGG19	Adversarial	76.00%	41.30%	1166.18	34.50	-	-	-
		NC	75.37%	41.60%	141.38	34.31	-0.49%	0.63%	0.00%
		NAD	73.81%	39.90%	4.12	37.45	9.35%	2.19%	1.40%
		ANP	75.69%	38.50%	1.07	32.95	-4.11%	0.31%	2.80%
		MOTH	75.08%	42.50%	81.54	39.45	13.99%	0.92%	0.00%
		DECK	75.68%	42.30%	18.21	**43.31**	**25.63%**	0.32%	0.00%
SVHN	NiN	Adversarial	91.63%	51.20%	132.00	45.99	-	-	-
		NC	92.02%	41.50%	27.38	31.66	-27.12%	0.00%	9.70%
		NAD	88.73%	31.10%	1.42	42.92	-0.96%	2.90%	20.10%
		ANP	91.45%	46.90%	0.56	36.13	-19.34%	0.18%	4.30%
		MOTH	91.40%	52.30%	51.42	38.40	-14.43%	0.24%	0.00%
		DECK	91.45%	50.20%	17.17	**67.60**	**48.12%**	0.18%	1.00%
	ResNet32	Adversarial	92.94%	58.30%	300.08	47.54	-	-	-
		NC	92.25%	49.90%	43.20	33.18	-20.29%	0.69%	8.40%
		NAD	90.02%	48.70%	5.81	48.25	15.85%	2.92%	9.60%
		ANP	91.70%	51.50%	2.33	35.91	-16.34%	1.24%	6.80%
		MOTH	92.52%	58.90%	111.74	35.44	-17.10%	0.42%	0.00%
		DECK	93.18%	58.10%	32.80	**56.54**	**32.57%**	0.00%	0.20%
Average		Adversarial	81.88%	47.00%	498.23	38.12	-	-	-
		NC	81.67%	42.70%	66.82	32.42	-9.16%	0.21%	4.30%
		NAD	80.09%	40.34%	3.04	38.26	5.01%	1.79%	6.66%
		ANP	81.14%	42.76%	1.07	33.19	-8.89%	0.74%	4.24%
		MOTH	81.30%	47.56%	61.80	35.87	-1.10%	0.58%	0.00%
		DECK	81.57%	46.64%	19.56	**49.58**	**33.34%**	0.30%	0.36%

Increasing the L^p bound will cause large accuracy degradation, which is not desired for preserving normal functionalities. In addition, UAP is extremely slow, taking 188.15 minutes on average, which is 4x longer than the natural training.

MOTH is a state-of-the-art hardening technique designed for backdoor defense. It can improve the class distance by more than 10% in two cases (ResNet20 and NiN on CIFAR-10). However, MOTH performs poorly on ResNet32 with SVHN, with -23.90% enlargement. Note that the distance measured in this paper is different from that in MOTH. We use the difference of feature representations from the encoder between backdoor samples and their clean counterparts, whereas MOTH leverages the trigger size in the input space (similar to the number of perturbed pixels), which is not suitable for measuring pervasive backdoors as they tend to alter all input pixels. The low performance of MOTH on SVHN might be due to the characteristics of the task. SVHN is a dataset for recognizing digital numbers from 0 to 9. The hardening procedure of MOTH relies on the backdoors generated by NC [70], which tends to produce backdoors with a small number of pixels for models on SVHN. Pervasive backdoors aim to transform the whole input instead of modifying a few pixels, which makes hardening techniques such as MOTH targeting pixel-level backdoors insufficient. The time cost of MOTH is 66.30 minutes, similar to training a natural model. By and large, DECK has the largest class distance improvement for the evaluated models, substantially outperforming baselines. DECK has a reasonable hardening cost, with less than half of the natural training (21.86 minutes on average). In Appendix V, we also demonstrate that leveraging the phenomenon of Neural Collapse [53] dose not improve the distance.

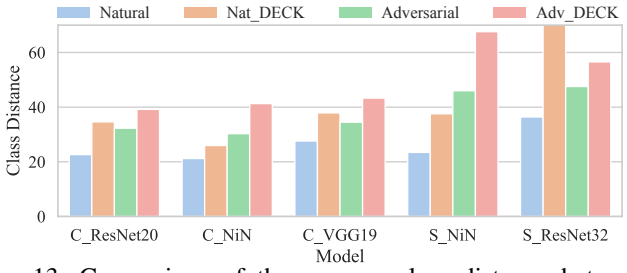

Fig. 13: Comparison of the average class distance between naturally trained and adversarially trained models. The first letter of the model name denotes the dataset (C for CIFAR-10 and S for SVHN).

Results on Adversarially Trained Models. Similar observations can be made on adversarially trained models in Table II. The clean accuracy of adversarailly trained models is usually lower than that of naturally trained ones. Although the class distance of adversarially trained models is much larger, they are still not optimal. DECK is able to further enlarge the distance by 33.34% on average, with only 0.30% and 0.36% degradation on accuracy and robustness, respectively. NC and ANP have worse performance on adversarially trained models than on natural ones, with -9.16% and -8.89% enlargement, respectively. In addition, they both have more than 4% robustness degradation, which is critical to adversarially trained models. NAD has larger improvement on two cases: VGG19 with CIFAR-10 (9.35%) and ResNet32 with SVHN (15.85%). The performance gain, however, is at the expense of much larger accuracy and robustness degradations (2.19%-2.92%) and (1.40%-9.60%), respectively, which is undesirable. MOTH has a reasonable performance on CIFAR-10, improving the class distance by from 4.76% to 13.99%. The results on SVHN are inferior, which is consistent with the observation on naturally trained models and discussed earlier. The hardening time costs of all the methods are similar to those on naturally trained models.

Comparison of Naturally and Adversarially Trained Models. We further study the class distances of naturally trained and adversarially trained models as well as models hardened by DECK. The average class distances of different models are shown in Figure 13. The x-axis denotes the dataset and model structure, and the y-axis denotes the class distance. Observe that all the adversarially trained models (green bars) have larger distances than their natural counterparts (blue bars). DECK can improve the class distance for both types of models. Particularly, three naturally trained models hardened by DECK (C_ResNet20, C_VGG19, and S_ResNet32) have even larger class distances than adversarially trained models. The accuracy of DECK-hardened versions of these models (91.68%) is much higher than adversarially trained models (81.27%) on average. This renders adversarial training is insufficient of achieving the optimal class distance improvement even with the sacrifice of clean accuracy. DECK can obtain much larger distances starting from the adversarially trained models as shown by the red bars in Figure 13 with negligible degradation on accuracy (0.30%) and robustness (0.36%). In other words, DECK can provide a much better extra protection against pervasive backdoor attacks with very limited sacrifice.

TABLE III: Comparison of distance measures by different backdoor generation methods. The fifth and seventh columns show the average class distance measured by DECK and filter backdoor, respectively.

D	M	Method	Accuracy	$\text{Dist.}_{\text{DECK}}$	$\text{Increase}_{\text{DECK}}$	$\text{Dist.}_{\text{filter}}$	$\text{Increase}_{\text{filter}}$
CIFAR-10	ResNet20	Natural	91.52%	22.62	-	39.04	-
		NC	89.95%	24.11	6.01%	43.46	9.67%
		NAD	91.09%	22.57	-0.29%	39.55	-6.63%
		ANP	90.31%	22.69	0.29%	33.45	-7.02%
		UAP	90.04%	23.98	6.64%	41.92	6.87%
		MOTH	90.34%	25.92	14.04%	42.78	9.58%
		DECK	90.31%	**34.60**	**52.66%**	**58.14**	**20.42%**
SVHN	ResNet32	Natural	95.15%	36.40	-	23.30	-
		NC	94.09%	33.87	-6.13%	20.42	-8.94%
		NAD	93.91%	34.36	-5.12%	15.97	-35.36%
		ANP	93.99%	37.59	4.19%	26.38	10.59%
		UAP	93.16%	39.51	9.85%	34.93	25.79%
		MOTH	94.49%	27.48	-23.90%	11.20	-69.47%
		DECK	93.84%	**87.20**	**142.73%**	**41.72**	**148.35%**

Results on Pre-trained Models. The evaluation on pre-trained models from TrojAI round 4 also exhibits similar findings. The results are shown in Table VIII in Appendix due to the page limit. The two backdoor removal techniques NAD and ANP have limited performance, with 7.45% and 3.66% improvement on the class distance on average, respectively. MOTH obtains a reasonable result with 34.42% distance enlargement on average. Their improvements however are still inferior to DECK that can enlarge the distance by 88.81%, from 4.01 to 6.93. Note that the pre-trained TrojAI models have already been hardened by adversarial training, indicating the effectiveness of DECK in hardening class distances for various settings.

Study of Distance Measure. The class distance in the above evaluation is measured by DECK as it can achieve high ASRs as shown in Figure 4. To demonstrate the stability of the measure, we also study using another backdoor generation method, i.e., filter backdoor for measuring the distance. Note that most existing pervasive backdoor attacks require data poisoning and hence cannot be adapted to measure the distance for a given pre-trained model. Two naturally trained models are employed for the study and the results are demonstrated in Table III. Columns 4-6 are copied from Table I, where columns 5-6 show the distances measured by DECK and the relative improvements over the original model. Columns 7-8 display the distance results measured by filter backdoor. Observe that the distances measured by both DECK and filter backdoor have similar results. Baselines such as NC, MOTH can only enlarge the distance for one of the two evaluated models. NAD fails to improve the distance for both models. UAP has relatively better performance with an average of 8.25% and 16.33% improvement measured by DECK and filter backdoor, respectively. Models hardened by DECK have the largest distance improvement (97.70% and 84.39% on average measured by our backdoor and filter backdoor, respectively). The distance measure by DECK is general and can characterize model's robustness against pervasive backdoor attacks. We also study the correlation of distance measures by MOTH and DECK in Appendix VI, demonstrating that they focus on different aspects of class distance.

C. Evaluation on Eliminating Injected Pervasive Backdoors

One of our most notable design goals is to use the transformation layer and the decoder to model a large spectrum of (existing) pervasive attacks. As such, hardening using our backdoor can improve resilience to all these attacks. We investigate six types of pervasive backdoors and apply DECK as well as the state-of-the-art backdoor removal techniques to study whether injected backdoors can be eliminated. We follow the same setup as in existing works [34], [63] by using only 5% of the original training set to eliminate backdoors. Seven existing backdoor removal techniques, namely, Standard Finetuning, NAD [34], ANP [76], ABS [39], Fine-pruning [38], MCR [87], and MOTH [63], are considered as the baselines. Standard Finetuning is a standard method originally proposed for transfer learning. It updates a pre-trained model's weights with a small learning rate on the training set. We leverage the finetuning baseline setting in NAD [34], which adopts data augmentation techniques including random crop, horizontal flipping, and Cutout [15] during training. ABS is able to reverse engineer feature space attack, such as simple filter backdoors. We hence leverage ABS to first invert a filter-like trigger pattern and apply it on clean samples. The poisoned model is then finetuned on these trigger-stamped inputs to unlearn injected backdoors. The descriptions of NAD, ANP, ABS, Fine-pruning, MCR, and MOTH can be found in Section II and Section V-A. We follow the original papers and the official repositories of the baselines, and carefully tune their settings to obtain the best results. Due to the page limit, the results of Fine-pruning and MCR are shown in Table IX and discussed in Appendix VII.

The results on five pervasive backdoors are presented in Table IV and on the filter backdoor in Figure 14 and Figure 16 (in Appendix). In Table IV, columns 1-3 denote the attacks, datasets, and model structures. Columns 4-5 show the clean accuracy of poisoned models and the attack success rate (ASR) of injected backdoors. The following columns present the accuracy and ASR of models after applying the different techniques. Deep Feature Space Trojan attack (DFST) [10] introduces a detoxification procedure by iteratively training on ABS [39] reverse-engineered backdoors to reduce the number of compromised neurons that can be leveraged by existing scanners for successful detection. We follow the original paper and evaluate on two settings with one round (Detox1) and three rounds (Detox3) of detoxification (the original paper [10] used at most three rounds). Observe that on CIFAR-10, Standard Finetuning, NAD, ABS, and MOTH can only reduce the ASR of DFST on ResNet32 Detox1 from 97.60% to more than 60%. ANP is able to reduce the ASR to 20.67% but with the cost of large accuracy degradation from 89.95% to 83.34%. DECK, on the other hand, can reduce the ASR to 14.22% with only 1.73% accuracy degradation. NAD, ABS, and MOTH have better performance on VGG13 Detox1, reducing the ASR from 95.89% to 24.56%, 33.78%, and 5.33%, respectively. However, with the increase of detoxification rounds, they can only reduce the ASR to more than 50%. ANP has poor performance on VGG13 with less than 10% ASR reduction for both one round and three rounds of detoxifications. DECK can consistently eliminate DFST backdoors, e.g., less than 15% ASR on ResNet32 and 6% on VGG13. The evaluation on the STL-10 dataset shows similar results. The baselines have limited performances with the best result of reducing the ASR to 37.56% on VGG13 Detox3 by ABS. DECK, on the other hand, can consistently remove injected backdoors to achieve as low as 2.67% ASR. In the worst case (on VGG13 Detox3), DECK can still reduce the ASR to 12.33%.

TABLE IV: Evaluation on backdoor attacks

Backdoor Attack	Dataset	Model	Original		Standard Finetune		NAD		ANP		ABS		MOTH		DECK	
			Accuracy	ASR	Accuracy	ASR	Accuracy	ASR	Accuracy	ASR	Accuracy	ASR	Accuracy	ASR	Accuracy	ASR
DFST	CIFAR-10	ResNet32 Detox1	89.95%	97.60%	87.30%	70.00%	88.00%	60.44%	83.34%	20.67%	86.97%	84.11%	88.90%	65.44%	88.22%	14.22%
		ResNet32 Detox3	90.93%	95.33%	89.20%	63.89%	85.78%	17.33%	83.68%	90.11%	90.74%	37.78%	89.53%	54.44%	88.11%	12.22%
		VGG13 Detox1	90.34%	95.89%	86.07%	93.56%	87.24%	24.56%	86.92%	89.56%	88.71%	33.78%	87.26%	5.33%	88.03%	2.00%
		VGG13 Detox3	91.29%	97.44%	89.55%	66.11%	87.09%	55.33%	88.46%	96.11%	88.67%	66.67%	89.91%	51.22%	89.08%	5.67%
	STL-10	ResNet32 Detox1	75.74%	97.67%	70.64%	70.64%	70.92%	44.00%	65.71%	90.11%	72.26%	68.56%	71.97%	60.89%	72.10%	2.67%
		ResNet32 Detox3	76.45%	99.00%	71.30%	93.22%	72.21%	89.11%	71.89%	69.56%	71.95%	97.78%	72.09%	71.56%	72.86%	4.78%
		VGG13 Detox1	72.18%	98.67%	70.11%	84.78%	68.91%	86.67%	68.88%	98.11%	68.46%	48.44%	69.86%	62.67%	68.61%	5.89%
		VGG13 Detox3	72.09%	98.89%	70.42%	97.33%	68.91%	81.00%	65.70%	97.33%	67.29%	37.56%	67.54%	86.56%	69.89%	12.33%
Blend	CIFAR-10	ResNet20	90.96%	99.96%	90.33%	84.92%	86.81%	3.94%	85.20%	6.22%	89.41%	5.66%	85.44%	12.26%	89.08%	0.00%
	SVHN	NiN	94.10%	92.37%	92.70%	0.54%	94.40%	0.33%	92.67%	0.41%	91.76%	10.66%	94.41%	0.20%	94.56%	0.85%
SIG	CIFAR-10	ResNet20	83.38%	93.30%	88.65%	59.40%	85.84%	9.44%	80.02%	37.44%	86.82%	85.29%	80.39%	17.72%	86.91%	3.97%
	SVHN	NiN	95.48%	92.46%	95.11%	41.60%	94.21%	0.68%	90.10%	20.91%	93.47%	55.02%	94.90%	0.69%	93.96%	0.46%
WaNet	CIFAR-10	ResNet18	94.15%	99.55%	93.58%	80.71%	91.37%	0.87%	91.38%	0.11%	89.61%	1.88%	92.65%	0.62%	91.12%	0.64%
	GTSRB	ResNet18	99.01%	98.94%	96.80%	48.75%	94.96%	0.02%	97.38%	0.00%	98.51%	0.00%	97.72%	0.01%	97.70%	0.30%
	CelebA	ResNet18	78.99%	99.08%	78.89%	21.35%	76.57%	15.34%	76.79%	14.22%	75.56%	21.69%	77.80%	8.91%	77.57%	8.12%
Clean Label	CIFAR-10	ResNet18	87.60%	98.36%	85.74%	97.34%	84.44%	7.47%	86.18%	4.46%	85.33%	7.18%	83.70%	8.34%	85.67%	4.22%
Average			**86.42%**	**97.16%**	**84.77%**	**67.13%**	**83.60%**	**31.03%**	**82.14%**	**45.96%**	**84.10%**	**41.38%**	**84.00%**	**31.68%**	**84.59%**	**4.90%**

Blend attack uses random small perturbation patterns as the backdoor, which can be easily eliminated by all the evaluated techniques, except for Standard Finetuning on CIFAR-10. DECK can reduce the ASRs to less than 1% for the two studied cases. We have similar observations for the results on SIG, WaNet, and Clean Label backdoors. Existing techniques can reduce the ASRs to a reasonable range (less than 20% for most cases). DECK further lower the ASRs to less than 5% for 5 out of 6 cases and 9% for the remaining case (WaNet on CelebA), outperforming the state-of-the-art approaches. The last row of Table IV presents the average results on the five pervasive backdoors. Standard Finetuning can only reduce the ASR from 97.16% to 67.13% with 1.65% accuracy degradation. NAD and MOTH can eliminate backdoors to having around 30% ASR with around 2.5% accuracy degradation. ANP is able to reduce the ASR to 45.96% with 4.28% accuracy degradation. ABS has a similar defense performance by reducing the ASR to 41.38% with 2.32% accuracy drop. Fine-pruning and MCR can eliminate backdoors to having around 50% ASR with around 4% accuracy degradation. Due to the page limit, the breakdown of their defense results is elaborated in Appendix VII. DECK has the largest ASR reduction from 97.16% to 4.90% and the accuracy degradation is only 1.83%.

Similar observations can be made for the evaluation on the filter backdoor, which is shown in Figure 14. The x-axis and y-axis denote the model IDs and the ASR, respectively. The bars in the light/dark colors show the ASR of injected backdoors before/after applying backdoor-erasing techniques. Observe that Standard finetuning (blue bars) can only repair half of the evaluated models (17 out of the 34 models). This is expected as backdoor attacks include clean data together with backdoored samples during training. Finetuning only on clean samples may not eliminate the backdoor patterns that have already been learned by poisoned models. NAD leverages the teacher-student structure and treats the model from Standard Finetuning as the teacher network. Its performance hence is limited by Standard Finetuning. This can be observed from the orange bars in Figure 14. NAD is only able to eliminate five more backdoors (with a total of 21 models). ANP has limited performance on TrojAI models, with only 15 poisoned models being repaired. The TrojAI poisoned models were trained by NIST [3] and different training strategies including random transformations, adversarial training, etc., are employed to make injected backdoors more robust and hard to detect. These strategies may reduce the sensitivity of individual neurons on backdoor patterns. ANP is hence not able to identify compromised neurons and fails to remove injected backdoors. This observation is consistent with the results on DFST backdoors that apply detoxification to reduce compromised neurons. ABS can only repair 15 models. As the injected backdoors in TrojAI models are label-specific, ABS may not be able to identify the correct victim-target class pair. The inverted triggers fail to expose the injected backdoor behaviors. Unlearning on those triggers hence cannot repair models.

MOTH can eliminate more backdoors than other baselines with 28 fixed models. As discussed in Section II, pervasive backdoors perturb all pixels on the input and are dynamic, which is different from MOTH's threat model that focuses on patch-like static backdoors. It can raise the bar for pervasive backdoor to some extent but still fails to repair 6 TrojAI poisoned models. DECK, on the other hand, can eliminate all the backdoors with an average ASR down to 0.55%, outperforming the others. The accuracies of poisoned models before and after repair are shown in Figure 16 in Appendix. Overall, all the approaches incur very small accuracy degradation on average ($< 0.3\%$), except for ANP (1.16%).

D. Reducing False Positives in Backdoor Scanning

Backdoor scanning aims to distinguish trojaned models from benign ones. The TrojAI competition [3] aims to evaluate the performance of various backdoor scanners submitted by the performers. However, due to the existence of natural pervasive backdoors in benign models, the performers encounter the problem of false positives, which are benign models misclassified as trojaned ones. To alleviate the problem, the TrojAI organizer adopted adversarial training for constructing benign models in round 4. The top performer however still reports 14 false positives. In Section V-B, we have demonstrated that DECK can effectively enlarge class distance of models for defending natural pervasive backdoors. Hence, an idea is to apply DECK on the 14 TrojAI (benign) models to check if they are still misclassified by scanners. After DECK's hardening,

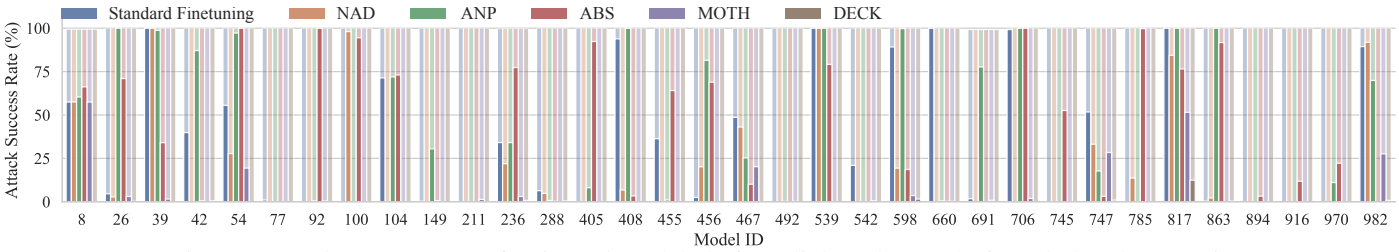

Fig. 14: Attack success rate of poisoned models before (light color) and after (dark color) repair

TABLE V: Adaptive Attack

Poisoning Rate	Backdoored		DECK	
	Accuracy	ASR	Accuracy	ASR
20%	83.64%	60.43%	83.94%	9.94%
50%	67.86%	74.94%	75.71%	4.42%

the accuracy degradation of these models is minimal (0.10% on average). The distance improvement is similar to that evaluated in Section V-B and hence omitted. We then run the top performer's model scanner (downloaded from the official repository [1]) on the round 4 dataset (with the 14 hardened models). There is no false positive reported by the scanner and no degradation on the true positives. The result demonstrates the capability of DECK in assisting backdoor scanning. This echoes our argument in Section II that in the future, published pre-trained models should undergo hardening similar to what we are proposing here. As such, any intentionally compromised models can be easily detected.

E. Adaptive Attack

We conduct an adaptive attack by optimizing a trigger pattern while applying DECK, the adaptive knowledge, during pervasive poisoning. We use the cross entropy loss to ensure the pattern can induce the target label and the L^2 loss to minimize the perturbation magnitude. The experiment is conducted on a ResNet20 model with CIFAR-10. The results are shown in Table V. The first column shows the poisoning rate. The following columns show the accuracy and ASR for the backdoored models before and after repaired by DECK. Observe that using a 20% poisoning rate, the backdoored model has 83.64% accuracy and 60.43% ASR. By increasing the poisoning rate to 50%, the ASR is improved to 74.94% with much more accuracy degradation (67.86% accuracy). As DECK aims to increase class distance while the poisoning incurs a small distance to the target class, the two contradicting goals make the accuracy much lower than a clean model (91.52%) and the ASR low as well. By applying DECK to the poisoned models, the ASR drops to 9.94%/4.42% without accuracy degradation (83.94%/75.71%) as shown in the last two columns in Table V, delineating DECK's resilience to adaptive attacks.

F. Ablation Study

DECK features a few important design choices. We study different decoder designs as well as the position of the feature representation in the encoder. We also investigate the effects of four loss terms used in pervasive backdoor generation. We observe that the best class distance enlargement of existing decoders is still 16.36% lower than ours. Using a deeper layer of the encoder for feature extraction does not boost the distance improvement. The four loss terms are all important, with $\mathcal{L}_{SSIM}$ and $\mathcal{L}_{norm}$ the crucial ones. Please see detailed discussion in Appendix VIII.

VI. Related Work

Backdoor attacks have been a critical security threat to deep learning models. In early works, a static trigger pattern is commonly used for stamping on training samples to poison the model, such as patch attacks [23], [9]. Recently, pervasive backdoors have been explored by researchers. Different from patch attacks, pervasive backdoors utilizes a backdoor function that transforms the whole input with input-specific perturbation. They are visually similar to their benign counterparts, making them stealthy and hard to detect. We have studied and evaluated six representative pervasive backdoor attacks in this paper, including DFST [10], Blend attack [9], SIG attack [5], WaNet [50], Clean Label attack [68], and Filter attack [39].

Defense techniques against backdoor attacks can be categorized into backdoor input detection, certified robustness, backdoor scanning, and backdoor removal. Backdoor input detection aims to detect inputs stamped with backdoor triggers, such as STRIP [20], Spectral Signatures [66], and so on [44], [7], [35], [42], [6], [17], [69]. Certified robustness provides certification to the classification results of individual samples, asserting the results can be trusted even in the presence of backdoors [46], [77], [78], [28]. Backdoor scanning focuses on identify whether a given model has been injected with backdoor. They leverage trigger inversion techniques to reverse-engineer injected backdoors [70], [39], [58], train a two-class detection network [30], [80], utilize adversarial perturbation to observe model behaviors [26], [83], [73], and so on [62]. These techniques are orthogonal to DECK as they do not consider natural backdoors. Backdoor removal aims to eliminate injected backdoors in poisoned models [38], [81]. State-of-the-art techniques leverage knowledge distillation to remove backdoors [34], prune out potentially compromised neurons [76], or harden models by enlarging input-space class distances [63]. Our evaluation in Section V-C demonstrates the effectiveness of DECK in eliminating backdoors, surpassing the state-of-the-arts. Researchers also study the backdoor learning curves to assist backdoor defense [12].

VII. Conclusion

We propose a general pervasive backdoor attack that models a spectrum of pervasive backdoors, which can be employed to harden models. The evaluation on various datasets and model structures demonstrates that our technique can improve model robustness against pervasive backdoors and reduce the attack success rate of six pervasive backdoor attacks from 99.06% to 1.94%, outperforming existing state-of-the-art hardening techniques.

REFERENCES

[1] Abs. https://github.com/naiyeleo/ABS.

[2] Keras Applications. https://keras.io/api/applications/.

[3] TrojAI Leaderboard. https://pages.nist.gov/trojai/.

[4] Eugene Bagdasaryan, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. How to backdoor federated learning. In *AISTATS*, pages 2938–2948, 2020.

[5] Mauro Barni, Kassem Kallas, and Benedetta Tondi. A new backdoor attack in cnns by training set corruption without label poisoning. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 101–105. IEEE, 2019.

[6] Alvin Chan and Yew-Soon Ong. Poison as a cure: Detecting & neutralizing variable-sized backdoor attacks in deep neural networks. *arXiv preprint arXiv:1911.08040*, 2019.

[7] Bryant Chen, Wilka Carvalho, Nathalie Baracaldo, Heiko Ludwig, Benjamin Edwards, Taesung Lee, Ian Molloy, and Biplav Srivastava. Detecting backdoor attacks on deep neural networks by activation clustering. *arXiv preprint arXiv:1811.03728*.

[8] Xiaoyi Chen, Ahmed Salem, Michael Backes, Shiqing Ma, and Yang Zhang. Badnl: Backdoor attacks against nlp models. *arXiv preprint arXiv:2006.01043*, 2020.

[9] Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526*, 2017.

[10] Siyuan Cheng, Yingqi Liu, Shiqing Ma, and Xiangyu Zhang. Deep feature space trojan attack of neural networks by controlled detoxification. In *AAAI*, 2021.

[11] Edward Chou, Florian Tramer, and Giancarlo Pellegrino. Sentinet: Detecting localized universal attack against deep learning systems. *SPW*, 2020.

[12] Antonio Emanuele Cinà, Kathrin Grosse, Sebastiano Vascon, Ambra Demontis, Battista Biggio, Fabio Roli, and Marcello Pelillo. Backdoor learning curves: Explaining backdoor poisoning beyond influence functions. *arXiv preprint arXiv:2106.07214*, 2021.

[13] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 215–223, 2011.

[14] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, pages 3213–3223, 2016.

[15] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.

[16] Alexey Dosovitskiy and Thomas Brox. Inverting visual representations with convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4829–4837, 2016.

[17] Min Du, Ruoxi Jia, and Dawn Song. Robust anomaly detection and backdoor attack detection via differential privacy. In *ICLR*, 2019.

[18] Kevin Eykholt, Ivan Evtimov, Earlence Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. Robust physical-world attacks on deep learning visual classification. In *CVPR*, pages 1625–1634, 2018.

[19] Minghong Fang, Xiaoyu Cao, Jinyuan Jia, and Neil Gong. Local model poisoning attacks to byzantine-robust federated learning. In *USENIX Security*, 2020.

[20] Yansong Gao, Change Xu, Derui Wang, Shiping Chen, Damith C Ranasinghe, and Surya Nepal. Strip: A defence against trojan attacks on deep neural networks. In *ACSAC*, pages 113–125, 2019.

[21] Yunjie Ge, Qian Wang, Baolin Zheng, Xinlu Zhuang, Qi Li, Chao Shen, and Cong Wang. Anti-distillation backdoor attacks: Backdoors can really survive in knowledge distillation. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 826–834, 2021.

[22] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *IJRR*, 2013.

[23] Tianyu Gu, Kang Liu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Evaluating backdooring attacks on deep neural networks. *IEEE Access*, 2019.

[24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.

[25] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, pages 4700–4708, 2017.

[26] Shanjiaoyang Huang, Weiqi Peng, Zhiwei Jia, and Zhuowen Tu. One-pixel signature: Characterizing cnn models for backdoor detection. In *ECCV*, 2020.

[27] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and< 0.5 mb model size. *arXiv preprint arXiv:1602.07360*, 2016.

[28] Jinyuan Jia, Xiaoyu Cao, and Neil Zhenqiang Gong. Certified robustness of nearest neighbors against data poisoning attacks. In *AAAI Conference on Artificial Intelligence*, 2020.

[29] Panagiota Kiourti, Kacper Wardega, Susmit Jha, and Wenchao Li. Trojdrl: evaluation of backdoor attacks on deep reinforcement learning. In *2020 57th ACM/IEEE Design Automation Conference (DAC)*, pages 1–6. IEEE, 2020.

[30] Soheil Kolouri, Aniruddha Saha, Hamed Pirsiavash, and Heiko Hoffmann. Universal litmus patterns: Revealing backdoor attacks in cnns. In *CVPR*, 2020.

[31] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.

[32] Keita Kurita, Paul Michel, and Graham Neubig. Weight poisoning attacks on pre-trained models. In *ACL*, 2020.

[33] Fredrik Larsson, Michael Felsberg, and P-E Forssen. Correlating fourier descriptors of local patches for road sign recognition. *IET Computer Vision*, 5(4):244–254, 2011.

[34] Yige Li, Nodens Koren, Lingjuan Lyu, Xixiang Lyu, Bo Li, and Xingjun Ma. Neural attention distillation: Erasing backdoor triggers from deep neural networks. In *ICLR*, 2021.

[35] Yiming Li, Tongqing Zhai, Baoyuan Wu, Yong Jiang, Zhifeng Li, and Shutao Xia. Rethinking the trigger of backdoor attack. *arXiv preprint arXiv:2004.04692*, 2020.

[36] Junyu Lin, Lei Xu, Yingqi Liu, and Xiangyu Zhang. Composite backdoor attack for deep neural network by mixing existing benign features. In *CCS*, 2020.

[37] Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. *arXiv preprint arXiv:1312.4400*, 2013.

[38] Kang Liu, Brendan Dolan-Gavitt, and Siddharth Garg. Fine-pruning: Defending against backdooring attacks on deep neural networks. In *RAID*, 2018.

[39] Yingqi Liu, Wen-Chuan Lee, Guanhong Tao, Shiqing Ma, Yousra Aafer, and Xiangyu Zhang. Abs: Scanning neural networks for back-doors by artificial brain stimulation. In *CCS*, pages 1265–1282, 2019.

[40] Yingqi Liu, Shiqing Ma, Yousra Aafer, Wen-Chuan Lee, Juan Zhai, Weihang Wang, and Xiangyu Zhang. Trojaning attack on neural networks. In *NDSS*, 2018.

[41] Yunfei Liu, Xingjun Ma, James Bailey, and Feng Lu. Reflection backdoor: A natural backdoor attack on deep neural networks. In *ECCV*, 2020.

[42] Yuntao Liu, Yang Xie, and Ankur Srivastava. Neural trojans. In *ICCD*, 2017.

[43] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*, pages 3730–3738, 2015.

[44] Shiqing Ma, Yingqi Liu, Guanhong Tao, Wen-Chuan Lee, and Xiangyu Zhang. Nic: Detecting adversarial samples with neural network invariant checking. In *NDSS*.

[45] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018.

[46] Michael McCoyd, Won Park, Steven Chen, Neil Shah, Ryan Roggenkemper, Minjune Hwang, Jason Xinyu Liu, and David Wagner. Minority reports defense: Defending against adversarial patches. In *In-*

ternational Conference on Applied Cryptography and Network Security, pages 564–582. Springer, 2020.

[47] Andreas Mogelmose, Mohan Manubhai Trivedi, and Thomas B Moeslund. Vision-based traffic sign detection and analysis for intelligent driver assistance systems: Perspectives and survey. *T-ITS*, 13(4):1484–1497, 2012.

[48] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. Universal adversarial perturbations. In *CVPR*, pages 1765–1773, 2017.

[49] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011.

[50] Anh Nguyen and Anh Tran. Wanet–imperceptible warping-based backdoor attack. In *ICLR*, 2021.

[51] Thien Duc Nguyen, Phillip Rieger, Huili Chen, Hossein Yalame, Helen Möllering, Hossein Fereidooni, Samuel Marchal, Markus Miettinen, Azalia Mirhoseini, Shaza Zeitouni, Farinaz Koushanfar, Ahmad-Reza Sadeghi, and Thomas Schneider. Flame: Taming backdoors in federated learning. In *31st USENIX Security Symposium (USENIX Security 22)*, 2022.

[52] Augustus Odena, Vincent Dumoulin, and Chris Olah. Deconvolution and checkerboard artifacts. *Distill*, 1(10):e3, 2016.

[53] Vardan Papyan, XY Han, and David L Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40):24652–24663, 2020.

[54] Phillip Rieger, Thien Duc Nguyen, Markus Miettinen, and Ahmad-Reza Sadeghi. Deepsight: Mitigating backdoor attacks in federated learning through deep model inspection. In *the Network and Distributed System Security Symposium (NDSS)*, 2022.

[55] Ahmed Salem, Rui Wen, Michael Backes, Shiqing Ma, and Yang Zhang. Dynamic backdoor attacks against machine learning models. *arXiv preprint arXiv:2003.03675*, 2020.

[56] Hadi Salman, Andrew Ilyas, Logan Engstrom, Ashish Kapoor, and Aleksander Madry. Do adversarially robust imagenet models transfer better? In *NeurIPS*, 2020.

[57] Ali Shafahi, Mahyar Najibi, Zheng Xu, John Dickerson, Larry S Davis, and Tom Goldstein. Universal adversarial training. In *AAAI*, 2020.

[58] Guangyu Shen, Yingqi Liu, Guanhong Tao, Shengwei An, Qiuling Xu, Siyuan Cheng, Shiqing Ma, and Xiangyu Zhang. Backdoor scanning for deep neural networks through k-arm optimization. In *ICML*, 2021.

[59] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[60] Johannes Stallkamp, Marc Schlipsing, Jan Salmen, and Christian Igel. Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition. *Neural networks*, 32:323–332, 2012.

[61] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016.

[62] Di Tang, XiaoFeng Wang, Haixu Tang, and Kehuan Zhang. Demon in the variant: Statistical analysis of dnns for robust backdoor contamination detection. In *30th USENIX Security Symposium (USENIX Security 21)*, 2021.

[63] Guanhong Tao, Yingqi Liu, Guangyu Shen, Qiuling Xu, Shengwei An, Zhuo Zhang, and Xiangyu Zhang. Model orthogonalization: Class distance hardening in neural networks for better security. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022.

[64] Vale Tolpegin, Stacey Truex, Mehmet Emre Gursoy, and Ling Liu. Data poisoning attacks against federated learning systems. In *ESORICS*, pages 480–501, 2020.

[65] Florian Tramer and Dan Boneh. Adversarial training and robustness for multiple perturbations. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

[66] Brandon Tran, Jerry Li, and Aleksander Madry. Spectral signatures in backdoor attacks. In *NeurIPS*, pages 8000–8010, 2018.

[67] Alexander Turner, Dimitris Tsipras, and Aleksander Madry. Label-Consistent Backdoor Attacks code. https://github.com/MadryLab/label-consistent-backdoor-code.

[68] Alexander Turner, Dimitris Tsipras, and Aleksander Madry. Clean-label backdoor attacks. 2018.

[69] Akshaj Kumar Veldanda, Kang Liu, Benjamin Tan, Prashanth Krishnamurthy, Farshad Khorrami, Ramesh Karri, Brendan Dolan-Gavitt, and Siddharth Garg. Nnoculation: broad spectrum and targeted treatment of backdoored dnns. *arXiv preprint arXiv:2002.08313*.

[70] Bolun Wang, Yuanshun Yao, Shawn Shan, Huiying Li, Bimal Viswanath, Haitao Zheng, and Ben Y Zhao. Neural cleanse: Identifying and mitigating backdoor attacks in neural networks. In *S&P*, pages 707–723, 2019.

[71] Hongyi Wang, Kartik Sreenivasan, Shashank Rajput, Harit Vishwakarma, Saurabh Agarwal, Jy-yong Sohn, Kangwook Lee, and Dimitris Papailiopoulos. Attack of the tails: Yes, you really can backdoor federated learning. *NeurIPS*, 2020.

[72] Lun Wang, Zaynah Javed, Xian Wu, Wenbo Guo, Xinyu Xing, and Dawn Song. Backdoorl: Backdoor attack against competitive reinforcement learning. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI-21)*, 2021.

[73] Ren Wang, Gaoyuan Zhang, Sijia Liu, Pin-Yu Chen, Jinjun Xiong, and Meng Wang. Practical detection of trojan neural networks: Data-limited and data-free cases. In *ECCV*, 2020.

[74] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

[75] Eric Wong, Leslie Rice, and J Zico Kolter. Fast is better than free: Revisiting adversarial training. In *ICLR*, 2020.

[76] Dongxian Wu and Yisen Wang. Adversarial neuron pruning purifies backdoored deep models. *Advances in Neural Information Processing Systems*, 34, 2021.

[77] Chong Xiang, Arjun Nitin Bhagoji, Vikash Sehwag, and Prateek Mittal. Patchguard: A provably robust defense against adversarial patches via small receptive fields and masking. In *30th USENIX Security Symposium (USENIX Security 21)*, 2021.

[78] Chong Xiang, Saeed Mahloujifar, and Prateek Mittal. Patchcleanser: Certifiably robust defense against adversarial patches for any image classifier. *arXiv preprint arXiv:2108.09135*, 2021.

[79] Chulin Xie, Keli Huang, Pin-Yu Chen, and Bo Li. Dba: Distributed backdoor attacks against federated learning. In *ICLR*, 2019.

[80] Xiaojun Xu, Qi Wang, Huichen Li, Nikita Borisov, Carl A Gunter, and Bo Li. Detecting ai trojans using meta neural analysis. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2021.

[81] Yi Zeng, Han Qiu, Shangwei Guo, Tianwei Zhang, Meikang Qiu, and Bhavani Thuraisingham. Deepsweep: An evaluation framework for mitigating dnn backdoor attacks using data augmentation. *arXiv preprint arXiv:2012.07006*.

[82] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.

[83] Xiaoyu Zhang, Ajmal Mian, Rohit Gupta, Nazanin Rahnavard, and Mubarak Shah. Cassandra: Detecting trojaned networks from adversarial perturbations. *arXiv preprint arXiv:2007.14433*, 2020.

[84] Xinyang Zhang, Zheng Zhang, and Ting Wang. Trojaning language models for fun and profit. In *European S&P*, 2021.

[85] Yuting Zhang, Kibok Lee, and Honglak Lee. Augmenting supervised neural networks with unsupervised objectives for large-scale image classification. In *International Conference on Machine Learning*, pages 612–621, 2016.

[86] Junbo Zhao, Michael Mathieu, Ross Goroshin, and Yann Lecun. Stacked what-where auto-encoders. *arXiv preprint arXiv:1506.02351*, 2015.

[87] Pu Zhao, Pin-Yu Chen, Payel Das, Karthikeyan Natesan Ramamurthy, and Xue Lin. Bridging mode connectivity in loss landscapes and adversarial robustness. In *International Conference on Learning Representations (ICLR 2020)*, 2020.

APPENDIX

I. ESSENCE OF THE VARIOUS UPSAMPLING METHODS

Assume a 2×2 input matrix $\mathbf{X}$ on the left of Figure 15 and a convolution kernel parameterized by $\mathbf{W} \in \mathbb{R}^{2\times2}$. Zero-padding is used (demonstrated by the dotted cells). Output values can be derived from the values in the matrix and the parameter values through the following equations on the top (lines 1-4). The inverse equations are shown on the bottom (lines 5-8).

$$
\begin{aligned}
y_0 &= w_0 \cdot x_0 + w_1 \cdot x_1 + w_2 \cdot x_2 + w_3 \cdot x_3 \\
y_1 &= w_0 \cdot x_1 + \cancel{w_1 \cdot 0} + w_2 \cdot x_3 + \cancel{w_3 \cdot 0} \\
y_2 &= w_0 \cdot x_2 + w_1 \cdot x_3 + \cancel{w_2 \cdot 0} + \cancel{w_3 \cdot 0} \\
y_3 &= w_0 \cdot x_3 + \cancel{w_1 \cdot 0} + \cancel{w_2 \cdot 0} + \cancel{w_3 \cdot 0} \\
x_0 &= \tfrac{1}{w_0} y_0 - \tfrac{w_1}{w_0^2} y_1 - \tfrac{w_2}{w_0^2} y_2 - \tfrac{w_0 w_3 - 2 w_1 w_2}{w_0^3} y_3 \\
x_1 &= \tfrac{1}{w_0} y_1 - \tfrac{w_2}{w_0^2} y_3 \\
x_2 &= \tfrac{1}{w_0} y_2 - \tfrac{w_1}{w_0^2} y_3 \\
x_3 &= \tfrac{1}{w_0} y_3
\end{aligned}
\tag{13}
$$

Here, we leave the activation functions out for discussion simplicity. The zero terms are crossed out. Downsampling is normally employed after convolutional operations to reduce data dimensions. We use max-pooling with a stride of 2 as an illustrative example. The resulted value (after downsampling) is hence $\hat{y} = \max(y_i), i \in \{0, 1, 2, 3\}$. The goal of upsampling is to recover values x_i from $\hat{y}$. Firstly, the upsampling methods (①-③ in Figure 15) expand $\hat{y}$ to the original matrix shape in different ways. A transposed convolutional operation then follows right after that to recover the input values, i.e., $\hat{x}_0, ..., \hat{x}_3$. Suppose the transposed convolution kernel is parameterized by $\mathbf{K} \in \mathbb{R}^{2\times2}$. The upsampled values $\hat{x}_0, ..., \hat{x}_3$ can be formalized as follows for each method.

$$
\begin{aligned}
\hat{x}_0 = {} & ① \; k_0 \cdot \hat{y} && ② \; k_0 \cdot \hat{y} \\
\hat{x}_1 = {} & \quad k_1 \cdot \hat{y} + \cancel{k_0 \cdot 0} && \quad (k_1 + k_0) \cdot \hat{y} \\
\hat{x}_2 = {} & \quad k_2 \cdot \hat{y} + \cancel{k_0 \cdot 0} && \quad (k_2 + k_0) \cdot \hat{y} \\
\hat{x}_3 = {} & \quad k_3 \cdot \hat{y} + \cancel{k_2 \cdot 0} + \cancel{\;\cdot\;} && \quad (k_3 + k_2 + k_1 + k_0) \cdot \hat{y}
\end{aligned}
$$

$$
\begin{aligned}
& ③ \; k_0 \cdot \hat{y} \\
& \quad k_1 \cdot \hat{y} + k_0 \cdot \epsilon_1 \\
& \quad k_2 \cdot \hat{y} + k_0 \cdot \epsilon_2 \\
& \quad k_3 \cdot \hat{y} + k_2 \cdot \epsilon_1 + k_1 \cdot \epsilon_2 + k_0 \cdot \epsilon_3
\end{aligned}
\tag{14}
$$

The constant ϵ_1, ϵ_2, and ϵ_3 in ③ are values randomly sampled from a Gaussian distribution $\mathcal{N}(0, \sigma^2)$, where σ is trained from the activation values from the corresponding layer in the encoder. We use a zero-mean Gaussian as it is simpler and does not require training for the mean. The results of using trained mean are similar to using a zero-mean Gaussian and hence omitted. Observe that both upsampling methods ① and ② induce a fixed linear correlation among the decoded values $\hat{x}_i$ (through $\hat{y}$). For instance in method ①, $\frac{\hat{x}_2}{\hat{x}_1} \equiv \frac{k_2}{k_1}$. However, as we can see from Equation 13, such correlations do not exist in downsampling.

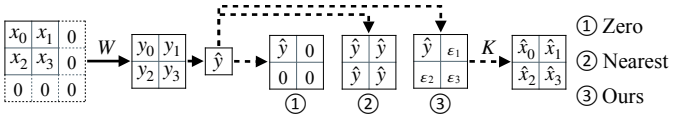

Fig. 15: Example for different upsampling methods

In our method ③, we mitigate such fixed correlations through randomization. Specifically, we leverage a Gaussian distribution to help approximating the possible disregarded values. In Equation 14, decoded values $\hat{x}_1, \hat{x}_2, \hat{x}_3$ have correlations with both the internal value $\hat{y}$ and the randomly sampled values $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$ (σ is collected from the activation values in the encoder). The magnitude of the randomization is controlled by the transposed convolution kernel k_i, allowing to express more complex relations. In addition, just like other randomization procedures (e.g., random dropout) used in deep neural network training, the randomization (in upsampling) also enhances the generality and robustness of the decoding kernels. During decoder training, we sample ϵ_i from Gaussian distributions for each upsampling layer.

II. REGIONAL TRANSFORMATION FOR SIG ATTACK

We use the SIG attack (the 4th column in Figure 1) that adds a strip-like pattern on images, as an example to study the capability of DECK in modeling various pervasive backdoors. As ABS [39] can invert feature map based triggers, we employ it as a baseline to compare with DECK. Specifically, we apply ABS and DECK to invert triggers for a poisoned ResNet20 model on CIFAR-10. The ASR of the inverted trigger is 39.00% by ABS and 97.00% by DECK. We also calculate the MSE between original poisoned images and inverted poisoned images. The distances are 1480.46 for ABS, and 1073.41 for DECK, delineating DECK can better simulate pervasive backdoors.

III. DETAILS OF EXPERIMENT SETUP

A. Datasets and Models

CIFAR-10 [31] is an object recognition dataset with 10 classification classes. It consists of 60,000 images and is divided into a training set (48,000 images), a validation set (2,000 images), and a test set (10,000 images).

SVHN [49] is a dataset contains house digital numbers extracted from Google Street View images. It has 73,257 training images and 26,032 test images. We divide the original training set into 67,257 images for training and 6,000 images for validation.

STL-10 [13] is an image recognition dataset with 10 classification classes. It consists of 5,000 training images and 8,000 test images.

LISA [47] is a U.S. traffic sign dataset with 47 different road signs. However, the number of images in different classes is not well balanced. Some classes have very few images, which is hard to train a reasonable model. We follow an existing work [18] by choosing 18 most common classes based on their numbers of training images. We use 5,635 images for training, 704 images for validation, and 704 images for testing.

GTSRB [60] is a German traffic sign recognition dataset with 43 classes. We split the dataset into a training set (35,289

images), a validation set (3,920 images), and a test set (12,630 images).

CelebA [43] is a face attributes dataset. It contains 10,177 identities with 202,599 face images. Each image has an annotation of 40 binary attributes. We follow [50] to select 3 out of 40 attributes, i.e., Heavy Makeup, Mouth Slightly Open, and Smiling, and create an 8-class classification task.

TrojAI [3] round 4 includes 16 types of model structures such as InceptionV3 [61], DenseNet121 [25], SqueezeNet [27], etc. The task of these models is to recognize synthetic street traffic signs with between 15 and 45 classes. Input images are constructed by compositing a foreground object, e.g., a synthetic traffic sign, with a random background images from five different dataset such as Cityscapes [14], KITTI [22], Swedish Roads [33], etc. A set of random transformations are applied during model training, such as blurring, lighting, shifting, titling, etc. Adversarial training such as PGD [45] and FBF [75] is also utilized to improve model quality. We use random seed 1030792629 to select 10 pre-trained models from the list of benign models from TrojAI round 4 [3]. For the filter attack, we use random seed 186270393 to select 34 poisoned models.

B. Baselines

NC [70] generates backdoors for each class and checks whether there exists an exceptionally small backdoor. If this is the case, it then injects the generated backdoor on 20% of the available training set (10% of the original training set) to retrain the subject model.

NAD [34] leverages the teacher-student structure to eliminate backdoors. It first finetunes the poisoned model on 5% of the training set. It uses this finetuned model as the teacher network, and the poisoned model as the student network. It then aims to reduce the internal feature differences between the teacher network and the student network by updating the student network. Finally, NAD outputs the student network as the cleaned model.

ANP [76] is based on the observation that backdoor related neurons are more sensitive to adversarial perturbations on their weights. It hence applies a mask on all the neurons in the model, adversarially perturbs neuron weights to increase the classification loss for a set of clean samples, and minimizes the size of mask. ANP then prunes neurons with small mask values, meaning that they have been compromised by backdoor attacks.

ABS [39] introduces a neuron stimulation analysis to expose abnormal behaviors of neurons in a deep neural network by increasing their activation values. Those neurons are regarded as compromised neurons and leveraged to reverse engineer backdoor triggers. ABS proposes a one-layer transformation to approximate/invert filter triggers. The inverted trigger is hence utilized to remove the injected backdoor in poisoned models following the unlearning procedure in NC [70] (see the above description).

Fine-pruning [38] prunes neurons that have low activation values for a set of clean samples. It then finetune the pruned model on a small set of clean samples.

TABLE VI: Comparison of different methods on hardening class distance for naturally trained models (cont.)

D	M	Method	Acc.	Rob.	Time	Dist.	Increase	ADeg.	RDeg.
CIFAR-10	ResNet50	Natural	92.71%	0.00%	74.11	23.05	-	-	-
		NC	91.54%	0.00%	183.00	24.40	5.78%	1.17%	0.00%
		NAD	92.65%	0.00%	6.01	23.10	0.48%	0.06%	0.00%
		ANP	92.27%	0.00%	4.66	23.35	1.35%	0.44%	0.00%
		UAP	91.83%	0.70%	345.61	24.10	5.08%	0.88%	0.00%
		MOTH	91.25%	0.00%	51.69	24.84	8.22%	1.46%	0.00%
		DECK	91.43%	0.00%	21.49	**42.42**	**83.96%**	1.28%	0.00%
LISA	CNN	Natural	97.30%	1.00%	0.15	42.51	-	-	-
		NC	82.67%	1.14%	8.27	46.86	15.46%	14.63%	0.00%
		NAD	95.45%	0.86%	0.15	44.79	8.58%	1.85%	0.14%
		ANP	96.45%	1.14%	0.07	45.68	11.30%	0.85%	0.00%
		UAP	95.60%	3.57%	1.79	39.35	-2.40%	1.70%	0.00%
		MOTH	96.31%	1.29%	13.09	50.16	18.69%	0.99%	0.00%
		DECK	96.73%	0.71%	3.82	**82.56**	**111.38%**	0.57%	0.29%
	ResNet20	Natural	98.86%	0.00%	1.70	29.79	-	-	
		NC	98.44%	0.00%	34.42	29.49	0.81%	0.42%	0.00%
		NAD	96.59%	0.00%	0.72	30.99	5.51%	2.27%	0.00%
		ANP	96.88%	0.00%	1.16	28.74	-1.47%	1.98%	0.00%
		UAP	96.16%	0.00%	6.33	29.58	2.67%	2.70%	0.00%
		MOTH	98.30%	0.14%	38.75	27.02	-5.55%	0.57%	0.00%
		DECK	98.44%	0.00%	4.42	**34.71**	**17.65%**	0.42%	0.00%
GTSRB	NiN	Natural	95.28%	13.90%	4.60	29.28	-	-	-
		NC	95.65%	13.30%	110.51	29.28	0.15%	0.00%	0.60%
		NAD	93.63%	12.20%	0.51	32.74	14.96%	1.65%	1.70%
		ANP	94.05%	13.30%	0.43	29.13	0.68%	1.23%	0.60%
		UAP	95.25%	15.80%	20.55	30.47	7.44%	0.03%	0.00%
		MOTH	95.55%	17.40%	74.01	30.75	8.00%	0.00%	0.00%
		DECK	96.33%	12.80%	45.71	**33.33**	**16.48%**	0.00%	1.10%
Average		Natural	96.04%	3.73%	20.14	31.16	-	-	-
		NC	92.08%	3.61%	84.05	32.51	5.55%	3.96%	0.12%
		NAD	94.58%	3.26%	1.85	32.91	7.38%	1.46%	0.46%
		ANP	94.91%	3.61%	1.58	31.73	2.97%	1.13%	0.11%
		UAP	94.71%	5.02%	93.57	30.88	3.20%	1.33%	0.00%
		MOTH	95.35%	4.71%	44.39	33.19	7.34%	0.69%	0.00%
		DECK	95.73%	3.38%	18.86	**48.26**	**57.37%**	0.31%	0.35%

MCR [87] linearly interpolates the weight parameters of two models. It also includes a set of trainable parameters during the interpolation. Specifically, the following equation is used to build a new model $\phi_\theta(t)$.

$$\phi_\theta(t) = (1-t)^2\omega_1 + 2t(1-t)\theta + t^2\omega_2, \;\; 0 \le t \le 1, \qquad (15)$$

where t is the interpolation hyper-parameter ranging from 0 to 1. ω_1 and ω_2 are the weight parameters of two pre-trained models, which are fixed. θ is a set of trainable parameters that have the same shape of ω_1 and ω_2. For eliminating backdoors in poisoned models, MCR uses the poisoned model and its finetuned version as the two endpoints (ω_1 and ω_2) and trains θ on a small set of clean samples. The best t is chosen for the interpolation based on the clean accuracy.

C. Adversarial Training and Robustness

PGD [45] is leveraged for creating adversarially trained models in Section V-B. We use L^∞ bound of $8/255$ for CIFAR-10, 0.03 for SVHN and GTSRB, and 0.1 for LISA. We follow the standard definition of L^∞ robustness by running PGD for multiple steps to attack the subject model [45] within the above given bound. The robustness measures how many generated adversarial examples are correctly classified as the ground truth by the subject model. For hardening with UAP, the L^∞ bound is determined based on the clean accuracy degradation. We use L^∞ bound of $4/255$ for CIFAR-10, 0.05 for SVHN, and 0.03 for LISA and GTSRB.

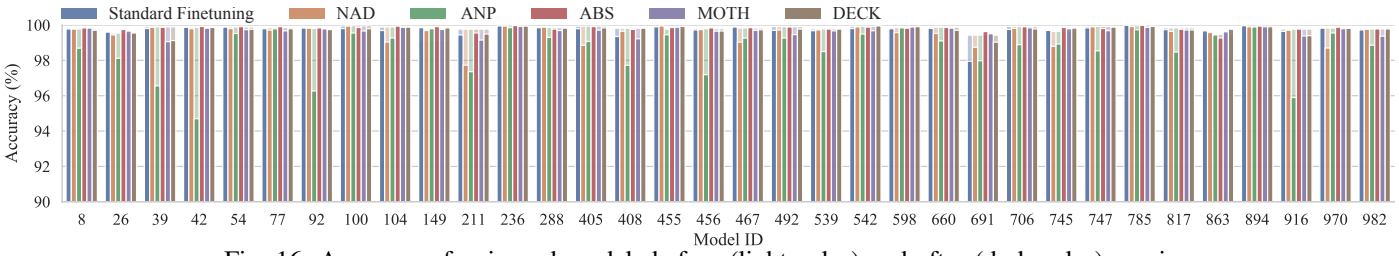

Fig. 16: Accuracy of poisoned models before (light color) and after (dark color) repair

D. Pervasive Backdoor Attacks

Deep Feature Space Trojan (DFST) [10] leverages a generative adversarial network (GAN) to inject a certain style (e.g., sunrise color style) to given training samples. It also introduces a detoxification procedure by iteratively training on ABS [39] reverse-engineered backdoors to reduce the number of compromised neurons that can be leveraged by existing scanners for successful detection. We follow the original paper and poison models with two settings: one-round detoxification and three-rounds detoxification.

Blend attack [9] injects a random perturbation pattern on the training samples of non-target classes and changes the ground truth labels of these samples to the target class (label 0). We use the random pattern reported in the original paper and poison 10% of the non-target-class training data with a blend ratio of $\alpha = 0.2$.

Sinusoidal Signal attck (SIG) [5] injects a strip-like pattern on the training samples of the target class and retains the original ground truth labels. We follow the setting in the original paper and generate the backdoor pattern using the horizontal sinusoidal function with $\Delta = 20$ and $f = 6$. We use label 0 as the target class and poison 8% of the training data in the target class.

WaNet [50] uses elastic image warping that deforms an image by applying the distortion transformation (e.g., distorting straight lines) as the backdoor. We download three backdoored models from the official repository [50], which are trained on CIFAR-10, GTSRB, and CelebA, respectively.

Clean Label attack [68] generates adversarial perturbations on the training samples in the target class using an adversarailly trained model. It then injects a 2×2 grid at the top left corner of the target-class inputs and retain their ground truth labels. We use L^∞ bound of $8/255$ for crafting adversarial perturbations, use label 3 as the target class, and poison 50% of the training data in the target class following the official repository [67].

Filter attack [39] applies Instagram filters on training samples and changes the ground truth labels of these samples to the target class. There are various filters can be used to poison data, such as Gotham filter, Nashville filter, Kelvin filter, Lomo filter, Toaster filter, etc.

E. Metrics

For measuring the distance $d_{i \to j}$, we select 100 random samples from the validation set of a victim class i. Our backdoor generation is then applied for 300 epochs to generate a backdoor that can induce misclassification to the target class j for 90% of those samples. We carry out the above evaluation process for three times using different random seeds to alleviate potential influence from random initialization, different sets of samples, etc. For TrojAI models, it is quite slow to measure the distance for every class pair (taking days to evaluate one single model). We hence randomly select 20 class pairs for evaluating each model. To avoid selecting the same set of class pairs for models with the same number of classes, we set the random seed to be the sum of 165838010 and the model id. The smallest distance obtained from the three trails is regarded as the final class distance.

TABLE VII: Comparison of different methods on hardening class distance for adversarially trained models (cont.)

D	M	Method	Acc.	Rob.	Time	Dist.	Increase	ADeg.	RDeg.
CIFAR-10	ResNet50	Adversarial	78.45%	42.30%	973.47	32.98	-	-	-
		NC	77.21%	42.70%	155.98	33.30	1.08%	1.24%	0.00%
		NAD	77.60%	44.50%	6.11	34.28	3.92%	0.85%	0.00%
		ANP	77.05%	33.70%	4.66	30.15	-8.49%	1.40%	8.60%
		MOTH	77.63%	41.70%	91.90	35.05	6.51%	0.82%	0.60%
		DECK	77.20%	42.70%	41.67	**39.81**	**21.70%**	1.25%	0.00%
LISA	CNN	Adversarial	75.43%	25.00%	5.00	97.82	-	-	-
		NC	73.44%	24.43%	8.54	87.82	-9.53%	1.99%	0.57%
		NAD	70.17%	23.57%	0.14	98.26	-2.02%	5.26%	1.43%
		ANP	73.86%	23.00%	0.07	86.14	-12.33%	1.57%	2.00%
		MOTH	74.72%	31.14%	17.47	95.72	-3.57%	0.71%	0.00%
		DECK	74.29%	30.86%	5.78	**105.47**	**15.30%**	1.14%	0.00%
	ResNet20	Adversarial	80.54%	36.57%	17.40	92.81	-	-	-
		NC	78.13%	41.86%	79.76	93.79	7.37%	2.41%	0.00%
		NAD	80.97%	40.00%	0.93	88.73	-1.08%	0.00%	0.00%
		ANP	80.26%	20.71%	1.18	81.04	-8.01%	0.28%	15.86%
		MOTH	81.25%	41.00%	43.34	77.27	-19.89%	0.00%	0.00%
		DECK	81.82%	36.71%	6.18	**101.55**	**17.80%**	0.00%	0.00%
GTSRB	NiN	Adversarial	90.96%	68.70%	54.98	50.02	-	-	-
		NC	91.55%	63.60%	111.72	36.11	-23.39%	0.00%	5.10%
		NAD	88.12%	54.30%	0.50	36.49	-20.58%	2.84%	14.40%
		ANP	89.67%	61.30%	0.54	37.11	-21.09%	1.29%	7.40%
		MOTH	90.32%	70.50%	71.29	48.54	2.00%	0.64%	0.00%
		DECK	90.37%	69.20%	27.83	**52.14**	**11.73%**	0.59%	0.00%
Average		Adversarial	81.35%	43.14%	262.71	68.41	-	-	-
		NC	80.08%	43.15%	89.00	62.75	-6.12%	1.27%	0.00%
		NAD	79.22%	40.59%	1.92	64.44	-4.94%	2.13%	2.55%
		ANP	80.21%	34.68%	1.61	58.61	-12.48%	1.13%	8.46%
		MOTH	80.98%	46.09%	56.00	64.15	-3.74%	0.37%	0.00%
		DECK	80.92%	44.87%	20.37	**74.74**	**16.63%**	0.42%	0.00%

IV. Additional Results on Enlarging Class Distance

From Table VI, observe that DECK can enlarge the class distance from 31.16 to 48.26 on average (57.37% improvement), with 0.31% and 0.35% degradation on accuracy and robustness, respectively. DECK has the largest distance improvement on CNN with LISA, with 111.38% enlargement. The three backdoor removal techniques, NC, NAD, and ANP,

TABLE VIII: Comparison of different methods on hardening class distance for TrojAI models

D	M	ID	Method	Accuracy	Time (m)	Distance	Increase	Degrad.
Swedish Roads	ResNet101	1	Natural	99.82%	1202.96	4.69	-	-
			NAD	99.67%	20.21	3.53	-4.97%	0.15%
			ANP	99.81%	10.24	4.04	-12.42%	0.01%
			MOTH	99.79%	186.09	4.61	27.76%	0.04%
			DECK	99.67%	143.52	13.10	259.66%	0.15%
KITTI City	VGG11bn	111	Natural	99.86%	582.55	3.42	-	-
			NAD	99.83%	21.26	3.80	11.53%	0.03%
			ANP	99.78%	5.24	2.78	-4.91%	0.08%
			MOTH	99.45%	97.16	3.58	14.88%	0.42%
			DECK	99.65%	77.38	3.82	31.14%	0.22%
	ResNet34	567	Natural	99.88%	339.35	4.37	-	-
			NAD	99.64%	9.18	2.80	-12.58%	0.25%
			ANP	99.77%	3.98	3.38	6.92%	0.11%
			MOTH	99.77%	62.86	3.96	17.10%	0.12%
			DECK	99.83%	87.60	8.58	80.78%	0.06%
	SqueezeNet1.0	730	Natural	99.94%	205.43	3.13	-	-
			NAD	99.83%	7.02	3.06	1.59%	0.11%
			ANP	99.87%	3.60	2.74	-1.70%	0.06%
			MOTH	99.76%	38.36	3.08	13.88%	0.17%
			DECK	99.87%	67.41	4.08	53.59%	0.06%
KITTI Residential	InceptionV3	692	Natural	99.84%	545.91	5.37	-	-
			NAD	99.66%	12.32	5.75	17.74%	0.18%
			ANP	99.30%	7.01	4.79	9.93%	0.54%
			MOTH	99.73%	282.45	4.50	16.20%	0.11%
			DECK	99.82%	78.87	5.77	48.05%	0.01%
	SqueezeNet1.1	810	Natural	99.88%	253.15	4.26	-	-
			NAD	99.72%	9.86	4.74	29.82%	0.15%
			ANP	99.88%	4.16	4.41	18.93%	0.00%
			MOTH	99.88%	33.99	4.91	35.78%	0.00%
			DECK	99.91%	177.18	6.34	57.15%	0.00%
KITTI Road	ResNet18	363	Natural	99.85%	561.59	4.10	-	-
			NAD	99.42%	14.10	5.01	21.25%	0.43%
			ANP	99.76%	5.36	5.03	13.88%	0.09%
			MOTH	99.57%	40.47	7.74	86.73%	0.28%
			DECK	99.78%	144.72	9.92	104.92%	0.07%
	DenseNet121	781	Natural	99.92%	654.97	2.93	-	-
			NAD	99.74%	17.30	2.44	-9.68%	0.18%
			ANP	99.94%	10.96	3.05	7.07%	0.00%
			MOTH	99.74%	121.62	3.31	22.98%	0.19%
			DECK	99.92%	98.31	4.10	40.47%	0.00%
Cityscapes	ShuffleNet1.0	210	Natural	99.96%	507.96	3.18	-	-
			NAD	99.95%	7.70	3.55	22.16%	0.01%
			ANP	99.94%	3.57	3.27	6.19%	0.02%
			MOTH	99.92%	50.48	4.65	70.05%	0.04%
			DECK	99.96%	170.38	5.73	88.17%	0.00%
	SqueezeNet1.0	621	Natural	99.93%	231.30	4.67	-	-
			NAD	99.92%	7.37	3.76	-2.35%	0.02%
			ANP	99.94%	3.14	3.31	-7.25%	0.00%
			MOTH	99.82%	83.70	6.58	38.83%	0.12%
			DECK	99.91%	78.57	7.83	124.15%	0.02%
Average			Natural	99.89%	508.52	4.01	-	-
			NAD	99.74%	12.63	3.84	7.45%	0.15%
			ANP	99.80%	5.73	3.68	3.66%	0.09%
			MOTH	99.74%	99.72	4.69	34.42%	0.15%
			DECK	99.83%	112.39	6.93	88.81%	0.06%

have limited performance on the class distance, with less than 8% improvement on average, which is consistent with the observation on CIFAR-10 and SVHN datasets and discussed in Section V-B. NC takes 84.05 minutes on average to harden a model, even longer than training the original model (20.14 minutes). NAD and ANP have low time costs, with 1.85 and 1.58 minutes overhead respectively. UAP can increase the class distance for 7.44% on NiN with GTSRB. But its performance on other datasets and models are still limited, with an average of 3.20% enlargement. In addition, UAP is extremely slow, taking 93.57 minutes on average, which is 4.6x longer than the natural training. MOTH can improve the class distance by more than 10% in one cases (CNN on LISA). However, MOTH cannot enlarge the class distance for ResNet20 on LISA. The time cost of MOTH is twice of training a natural model, with 44.39 minutes overhead. Overall, DECK has the largest class distance improvement for the evaluated models, substantially outperforming baselines. DECK has a reasonable hardening cost, with 18.86 minutes overhead on average.

Similar observations can also be made on adversarially trained models in Table VII. DECK is able to enlarge the class distance by 16.63% on average, with only 0.42% accuracy degradation and no robustness loss. NC, NAD, and ANP have worse performance on adversarially trained models than on naturally trained ones, with -6.12%, -4.94%, and -12.48% enlargement, respectively. In addition, NAD and ANP have more than 2.5% robustness degradation, which is critical to adversarially trained models. MOTH has a reasonable performance on CIFAR-10, improving the class distance by 6.51%. The results on LISA and GTSRB are inferior. The hardening time costs of all the methods are similar to those on naturally trained models.

V. COMPARISON WITH NEURAL COLLAPSE

Existing work [53] observes that by training the model to the terminal phase (training beyond zero misclassification error), the model can exhibit better generalization performance, better robustness, etc., which is called Neural Collapse. We follow the existing work [53] by training a ResNet20 model on CIFAR-10 to 250 epochs and 300 epochs. The class distance for the two models are 24.46 and 24.63, respectively, slightly larger than the original model (22.62). The model hardened by DECK has a distance of 34.60. This is because the normal training does not aim to enlarge pairwise class distances but only reduces prediction loss. More training epochs can hardly improve the distance.

VI. CORRELATION OF DISTANCE MEASURES BY MOTH AND DECK

We use the distance in MOTH to measure a hardened ResNet20 model on CIFAR-10 by DECK. The distance value increases from 53.49 to 64.29, whereas MOTH can increase the distance to 109.70. This shows that DECK can increase the MOTH distance to some extent. The distance measured in our paper increases from 22.62 to 25.92 for MOTH, to 34.60 for DECK. The two measures focus on different aspects of class distance, analogous to L^0 and L^2 norms in adversarial robustness. Existing works have shown that adversarial robustness on one norm does not promise robustness on another norm [65], which is consistent with the above observation in backdoor robustness.

VII. COMPARISON WITH MORE BASELINES ON ELIMINATING INJECTED PERVASIVE BACKDOORS

The results of two more baselines, namely, Fine-pruning and MCR on eliminating injected pervasive backdoors are shown in Table IX. Columns 6-7 present the accuracy and ASR of models repaired by Fine-pruning and columns 8-9 are for MCR. We also show the results of the original models and models repaired by DECK in columns 4-5 and columns 10-11 respectively for comparison. It is evident that

TABLE IX: Evaluation on backdoor attacks with more baselines

Backdoor Attack	Dataset	Model	Original		Fine-pruning		MCR		DECK	
			Accuracy	ASR	Accuracy	ASR	Accuracy	ASR	Accuracy	ASR
DFST	CIFAR-10	ResNet32 Detox1	89.95%	97.60%	88.20%	62.89%	87.95%	62.67%	88.22%	14.22%
		ResNet32 Detox3	90.93%	95.33%	88.09%	47.44%	82.84%	60.78%	88.11%	12.22%
		VGG13 Detox1	90.34%	95.89%	87.37%	51.11%	86.58%	90.33%	88.03%	2.00%
		VGG13 Detox3	91.29%	97.44%	88.84%	85.67%	88.81%	86.78%	89.08%	5.67%
	STL-10	ResNet32 Detox1	75.74%	97.67%	68.89%	96.11%	68.05%	84.67%	72.10%	2.67%
		ResNet32 Detox3	76.45%	99.00%	69.25%	88.22%	69.98%	81.89%	72.86%	4.78%
		VGG13 Detox1	72.18%	98.67%	67.12%	67.00%	66.06%	66.22%	68.61%	5.89%
		VGG13 Detox3	72.09%	98.89%	68.14%	49.44%	66.66%	79.67%	69.89%	12.33%
Blend	CIFAR-10	ResNet20	90.96%	99.96%	87.75%	3.63%	85.53%	63.58%	89.08%	0.00%
	SVHN	NiN	94.10%	92.37%	88.26%	23.75%	93.50%	0.59%	94.56%	0.85%
SIG	CIFAR-10	ResNet20	83.38%	93.30%	81.01%	76.29%	83.31%	16.63%	86.91%	3.97%
	SVHN	NiN	95.48%	92.46%	93.35%	23.49%	93.19%	45.16%	93.96%	0.46%
WaNet	CIFAR-10	ResNet18	94.15%	99.55%	89.14%	2.09%	93.29%	1.74%	91.12%	0.64%
	GTSRB	ResNet18	99.01%	98.94%	96.06%	63.40%	98.54%	10.47%	97.70%	0.30%
	CelebA	ResNet18	78.99%	99.08%	76.57%	18.07%	78.32%	16.21%	77.57%	8.12%
Clean Label	CIFAR-10	ResNet18	87.60%	98.36%	83.66%	26.77%	85.23%	13.91%	85.67%	4.22%
Average			**86.42%**	**97.16%**	**82.61%**	**49.09%**	**82.99%**	**48.83%**	**84.59%**	**4.90%**

TABLE X: Ablation study on effects of different design choices

Method	Accuracy	Time (min)	Distance	Increase	Difference
Natural	91.52%	56.77	22.62	-	-
DECK	90.31%	16.73	34.60	52.66%	-
w/ SWWAE [86], [85]	90.97%	16.48	23.75	4.63%	-48.03%
w/ Nearest [52]	90.32%	16.37	23.35	3.52%	-49.14%
w/ Zero [16]	89.83%	16.55	30.92	36.30%	-16.36%
w/ Pool2	90.81%	19.61	24.93	10.34%	-42.32%
w/o $\mathcal{L}_{content}$	90.13%	16.47	33.42	46.48%	-6.18%
w/o $\mathcal{L}_{SSIM}$	90.36%	17.25	29.15	27.91%	-24.75%
w/o $\mathcal{L}_{norm}$	90.27%	17.88	30.10	32.38%	-20.28%
w/o $\mathcal{L}_{smooth}$	90.13%	18.09	33.64	47.54%	-5.12%

both Fine-pruning and MCR have limited defense performance against DFST with only around 50% ASR reduction at best. As discussed in Section V-C, DFST leverages a detoxification procedure to make injected backdoors more robust against defenses. Fine-pruning is not able to identify compromised neurons and hence less effective. Since Standard Finetune fails to remove DFST-injected backdoors (see Table IV in Section V-C), there does not exist a backdoor-free model along the path explored by MCR from the poisoned model to the finetuned version. Fine-pruning can eliminate backdoors for a few models poisoned by Blend and WaNet (rows 11 and 15). However, it still fails to completely remove injected backdoors for the remaining cases. The observation is similar for MCR with two cases of the ASR lower than 5%. Overall, Fine-pruning/MCR can reduce the ASR to 49.09%/48.83% with 3.81%/3.43% accuracy degradation on average. In contrast, DECK has the ASR reduction from 97.16% to 4.90% and the accuracy degradation is only 1.83%, significantly outperforming the two baselines.

VIII. ABLATION STUDY

DECK features a few important design choices. In this section, we aim to study these choices individually to better understand their contributions to the performance. In particular, we first evaluate on different decoder designs as well as the position of the feature representation in the encoder. We then study the effects of four loss terms used in pervasive backdoor generation. The ablation study is conducted on a ResNet20 model on CIFAR-10 and the results are presented in Table X. Row 2 denotes the original model and row 3 the final result of DECK. Rows 4-6 show the results of using different decoders. Row 7 denotes using the Pool2 layer in the encoder as the feature representation. Rows 8-11 present the results of excluding each loss term individually during hardening. Observe that using SWWAE has only 4.63% class distance improvement. This is because the feature representation of backdoor samples can be very different from normal samples. The positional information of normal samples in the encoder leveraged by SWWAE does not help decoding the transformed representation that are similar to that of backdoor samples. Using Nearest has only 3.52% improvement and hardening using Zero leads to 16.36% improvement degradation. Nearest and Zero decoders impose a strong linear correlation among upsampled values, which reduces the quality of generated backdoor samples. The default setting of DECK employs the Pool1 layer of the encoder for extracting the feature representation, which is the first pooling layer in the encoder. Here, we study using a deeper layer for feature extraction. The result in row 7 shows that using a deeper layer reduces the quality of decoded samples as more errors are introduced during upsampling in the decoder and causes a smaller class distance improvement. For the four loss terms, we observe that $\mathcal{L}_{SSIM}$ has the largest impact on the final result (24.75% degradation without it). This is because the SSIM score directly constrains the quality of generated backdoor samples looking similar to original inputs. Without it, the backdoor samples can be too different from normal inputs and the model cannot learn the correct features. Loss $\mathcal{L}_{norm}$ on the normalization layer is also quite important (20.28% degradation without it) as it makes sure the normalized inputs are not far from the original distribution. Loss $\mathcal{L}_{content}$ constrains the difference of feature representations between backdoor samples and normal inputs. The purpose is similar to the SSIM loss and it can boost the performance by 6.18%. Loss $\mathcal{L}_{smooth}$ smooths the backdoor samples to improve the generated quality, which can provide 5.12% performance gain.