arXiv:2007.15147v3 [cs.LG] 17 Jun 2021

A General Framework For Detecting Anomalous Inputs to DNN Classifiers

Jayaram Raghuram [*1] **Varun Chandrasekaran** [*1] **Somesh Jha** [1 2] **Suman Banerjee** [1]

Abstract

Detecting anomalous inputs, such as adversarial and out-of-distribution (OOD) inputs, is critical for classifiers (including deep neural networks or DNNs) deployed in real-world applications. While prior works have proposed various methods to detect such anomalous samples using information from the internal layer representations of a DNN, there is a lack of consensus on a principled approach for the different components of such a detection method. As a result, often heuristic and one-off methods are applied for different aspects of this problem. We propose an unsupervised anomaly detection framework based on the internal DNN layer representations in the form of a meta-algorithm with configurable components. We proceed to propose specific instantiations for each component of the meta-algorithm based on ideas grounded in statistical testing and anomaly detection. We evaluate the proposed methods on well-known image classification datasets with strong adversarial attacks and OOD inputs, including an *adaptive attack* that uses the internal layer representations of the DNN (often not considered in prior work). Comparisons with five recently-proposed competing detection methods demonstrates the effectiveness of our method in detecting adversarial and OOD inputs.

1. Introduction

Deep neural networks (DNNs) have achieved impressive performance on a variety of challenging machine learning (ML) problems such as image classification, object detection, speech recognition, and natural language processing (He et al., 2015; Krizhevsky et al., 2017). However, it is well-known that DNN classifiers can be highly inaccurate (sometimes with high confidence) on test inputs from outside the training distribution (Nguyen et al., 2015; Szegedy et al., 2014; Hendrycks & Gimpel, 2017; Hein et al., 2019). Such anomalous inputs can arise in real-world settings either unintentionally due to external factors, or due to malicious adversaries that intend to cause prediction errors in the DNN and disrupt the system (Barreno et al., 2006; Biggio & Roli, 2018). Therefore, it is critical to have a defense mechanism that can detect such anomalous inputs, and take suitable corrective action (*e.g.*, abstain from predicting (Tax & Duin, 2008) or provide a more reliable class prediction).

In this work, we propose `JTLA` (Joint statistical Testing across DNN Layers for Anomalies), a general unsupervised framework for detecting anomalous inputs (including adversarial and OOD) to a DNN classifier using its layer representations. `JTLA` utilizes the rich information at the intermediate layer representations of a DNN to obtain a better understanding of the patterns produced by anomalous inputs for detection. Our method is *unsupervised*, *i.e.*, it does not utilize any specific class(es) of known anomalous samples for learning or tuning its parameters. While a number of prior works have addressed the problem of adversarial and OOD detection (Feinman et al., 2017; Xu et al., 2018; Li & Li, 2017; Lee et al., 2018; Ma et al., 2018; Roth et al., 2019), including ones that utilize intermediate layer representations of a DNN (Li & Li, 2017; Meng & Chen, 2017; Xu et al., 2018; Lee et al., 2018; Zheng & Hong, 2018; Ma et al., 2018; Papernot & McDaniel, 2018; Miller et al., 2019; Yang et al., 2020; Sastry & Oore, 2020), some key limitations persist, that we propose to address in this work.

Limitations of prior work. First, a number of existing detection methods (Feinman et al., 2017; Lee et al., 2018; Ma et al., 2018; Yang et al., 2020) being supervised, have to be presented with a broad sampling of known anomalous samples for training (*e.g.*, different adversarial attacks of varying strength). Such methods typically also need to configure hyper-parameters based on the known anomalous samples from the training set (*e.g.*, using cross-validation). As a result, they often do not generalize well to *unknown* anomalies (*e.g.*, novel or adaptive attacks). It has been shown that a majority of the current detection methods fail to handle unseen and adaptive adversaries that are aware of the defense mechanism (Carlini & Wagner, 2017a; Tramèr et al., 2020). **Second**, detection methods that use only the input, output (pre-softmax), or a specific DNN layer (Roth et al.,

[*]Equal contribution [1]Computer Sciences, University of Wisconsin, Madison, USA. [2]XaiPient Inc., Princeton, NJ, USA. Correspondence to: Jayaram Raghuram <jayaramr@cs.wisc.edu>.

Proceedings of the 38^{th} International Conference on Machine Learning, PMLR 139, 2021. Copyright 2021 by the author(s).

2019; Hendrycks & Gimpel, 2017; Feinman et al., 2017) do not jointly exploit the properties exhibited by anomalous inputs across the layers. **Third**, although methods that utilize information from multiple layers (listed earlier) propose specific test statistics or features calculated from the layer representations (*e.g.*, local intrinsic dimensionality (Ma et al., 2018)), there is a lack of a *general anomaly detection framework* where one can plug-in test statistics, aggregation, and scoring methods suitable for the detection task. **Fourth**, existing unsupervised detection methods that are based on density (generative) modeling of the DNN layer representations (Zheng & Hong, 2018; Miller et al., 2019; Feinman et al., 2017) are not well-suited to handle the (often very) high dimensional layer representations. **Finally**, we observe that existing detection methods often do not utilize the predicted class of the DNN to focus on class-conditional properties of the layer representations (Ma et al., 2018; Li & Li, 2017; Xu et al., 2018; Yang et al., 2020), which can lead to improved detection performance. While prior works such as (Roth et al., 2019; Miller et al., 2019; Zheng & Hong, 2018; Sastry & Oore, 2020) are exceptions to this, there is still need for a unified approach in this regard.

Our contributions can be summarized as follows:

- We propose a general unsupervised framework `JTLA` for detecting anomalous inputs to a DNN using its layer representations. We first present a meta-algorithm and describe its components in general terms (§ 3). We then propose specific methods for realizing the components in a principled way (§ 4). The proposed framework is modular, and a number of prior works for anomaly detection based on the layer representations can be cast into this meta-algorithm.
- The importance of designing an adaptive, defense-aware adversary has been underscored in the literature (Carlini & Wagner, 2017a; Tramèr et al., 2020). We propose and evaluate against an adversarial attack that focuses on defenses (such as ours) that use the k-nearest neighbors of the layer representations of the DNN (§ 5).
- We report extensive experimental evaluations comparing `JTLA` with five baseline methods on three image classification datasets trained with suitably-complex DNN architectures. For adversarial detection, we evaluate on three well-known whitebox attacks and our proposed defense-aware attack (§ 6 and Appendix E) [1].

2. Related Works

We provide a brief review of related works on adversarial and OOD detection, focusing on methods that use the internal layer representations of a DNN. A detailed discussion of closely-related prior works, and how they fit into the proposed anomaly detection framework is provided in Appendix A.4. Recent surveys on adversarial learning and anomaly detection for DNNs can be found in (Biggio & Roli, 2018; Miller et al., 2020; Bulusu et al., 2020).

Prior works on adversarial and OOD detection can be broadly categorized into unsupervised and supervised methods. Supervised methods such as (Lee et al., 2018), (Ma et al., 2018), and (Yang et al., 2020) use a training set of adversarial or OOD samples (*i.e.*, known anomalies) to train a binary classifier that discriminates natural inputs from anomalies. They extract specific informative test statistics from the layer representations as features for the classifier. On the other hand, unsupervised methods such as (Roth et al., 2019; Zheng & Hong, 2018; Miller et al., 2019; Sastry & Oore, 2020; Li & Li, 2017; Xu et al., 2018), rely on interesting statistical properties and generative modeling of natural inputs at specific (*e.g.*, logit) or multiple layer representations of the DNN for detection. Works such as the trust score (Jiang et al., 2018), deep kNN (Papernot & McDaniel, 2018), and by Jha et al. (2019) have explored the problem of developing a confidence metric that can independently validate the predictions of a classifier. Inputs with low confidence scores are likely to be misclassified and hence are detected as anomalies.

3. Anomaly Detection Meta-algorithm

We first introduce the notation and problem setup, followed by a description of the proposed meta-algorithm.

3.1. Notations and Setup

Consider the conventional classification problem where the goal is to accurately classify an input $\mathbf{x} \in \mathcal{X}$ into one of m classes $[m] := \{1, \cdots, m\}$. We focus on DNN classifiers that learn a function of the form $\mathbf{F}(\mathbf{x}) = [F_1(\mathbf{x}), \cdots, F_m(\mathbf{x})]$, $\mathbf{F} : \mathcal{X} \mapsto \Delta_m$, where $\mathcal{X}$ is the space of inputs to the DNN and $\Delta_m = \{(p_1, \cdots, p_m) \in [0,1]^m : \sum_i p_i = 1\}$ is the space of output class probabilities. The class prediction of the DNN based on its output class probabilities is defined as $\widehat{C}(\mathbf{x}) = \arg\max_{c \in [m]} F_c(\mathbf{x})$. The multi-layer architecture of a DNN allows the input-output mapping to be expressed as a composition of multiple functions, *i.e.*, $\mathbf{F}(\mathbf{x}) = (\mathbf{g}_L \circ \mathbf{g}_{L-1} \cdots \circ \mathbf{g}_1)(\mathbf{x})$, where L is the number of layers. The output from an intermediate layer $\ell \in \{1, \cdots, L\}$ of the DNN, $\mathbf{f}_\ell(\mathbf{x}) = (\mathbf{g}_\ell \circ \cdots \circ \mathbf{g}_1)(\mathbf{x}) \in \mathbb{R}^{d_\ell}$, is referred to as its layer representation [2]. We also use the shorthand notation $\mathbf{x}^{(\ell)} = \mathbf{f}_\ell(\mathbf{x})$, with $\mathbf{x}^{(0)} = \mathbf{f}_0(\mathbf{x})$ denoting the vectorized input. The set of layers and distinct layer pairs are denoted by $\mathcal{L} = \{0, \cdots, L\}$ and $\mathcal{L}^2 = \{(\ell_1, \ell_2) \in \mathcal{L} \times \mathcal{L} : \ell_2 > \ell_1\}$. Table 1 in the

[1] The code base associated with our work can be found at: https://github.com/jayaram-r/adversarial-detection.

[2] Layers with tensor-valued outputs (*e.g.*, convolution) are flattened into vectors. Boldface symbols are used for vectors and tensors.

Appendix provides a quick reference for the notations.

We assume access to the trained DNN classifier to defend, and a labeled data set $\mathcal{D} = \{(\mathbf{x}_n, c_n),\ n = 1, \cdots, N\}$ that is different from the one used to train the DNN and does not contain any anomalous samples. We define an augmented data set $\mathcal{D}_a = \{(\mathbf{x}_n^{(0)}, \cdots, \mathbf{x}_n^{(L)}, c_n, \hat{c}_n),\ n = 1, \cdots, N\}$ that is obtained by passing samples from $\mathcal{D}$ through the DNN and extracting their layer representations $\mathbf{x}_n^{(\ell)} = \mathbf{f}_\ell(\mathbf{x}_n),\ \ell \in \mathcal{L}$ and the class prediction $\hat{c}_n = \widehat{C}(\mathbf{x}_n)$. We also define subsets of $\mathcal{D}_a$ corresponding to each layer $\ell \in \mathcal{L}$, and each predicted class $\hat{c} \in [m]$ or true class $c \in [m]$ respectively as:

$$\widehat{\mathcal{D}}_a(\ell, \hat{c}) = \{(\mathbf{x}_n^{(\ell)}, c_n, \hat{c}_n),\ n = 1, \cdots, N : \hat{c}_n = \hat{c}\},$$
$$\mathcal{D}_a(\ell, c) = \{(\mathbf{x}_n^{(\ell)}, c_n, \hat{c}_n),\ n = 1, \cdots, N : c_n = c\}.$$

3.2. Components of the Meta-algorithm

Algorithm 1 Meta-algorithm for Anomaly Detection

Inputs: Trained DNN $\mathbf{F}(\cdot)$, Dataset $\mathcal{D}$, Test input $\mathbf{x}$, FPR α or detection threshold τ.
Output: Detector decision – normal 0 or anomaly 1.

Preprocessing:
Calculate the detection threshold τ (if not specified).
Calculate the class prediction and layer representations of $\mathbf{x}$.
Create the data subsets corresponding to each layer, predicted class, and m true classes.

I. Test statistics (TS):
  for each layer ℓ:
    Calculate the TS at layer ℓ conditioned on the predicted class and the m candidate true classes.
  Compile the $m+1$ TS vectors from the layers.

II. Normalizing transformations:
  if multivariate normalization:
    Normalize each of the $m+1$ TS vectors.
  else
    for each layer ℓ:
      Normalize the $m+1$ TS from layer ℓ.
    for each distinct layer pair (ℓ_1, ℓ_2): **[optional]**
      Normalize the $m+1$ TS pairs from layers ℓ_1, ℓ_2.

III. Layerwise aggregation and scoring:
  if multivariate normalization:
    No need to aggregate the normalized TS.
  else
    Aggregate the normalized TS from the layers and layer pairs for the predicted class and each candidate true class.
  Calculate the final score from the $m+1$ aggregated normalized TS.

IV. Detection decision:
  Return anomaly (1) if the final score exceeds threshold; Else return normal (0).

The proposed meta-algorithm for detecting anomalous inputs to a DNN classifier based on its layer representations is given in Algorithm 1. A more formal version of the same can be found in Algorithm 2 in the Appendix. Details of the individual components of the meta-algorithm are discussed next. For this discussion, consider a test sample $\mathbf{x}$ whose true class is unknown, predicted class is $\widehat{C}(\mathbf{x}) = \hat{c}$, and layer representations are $\mathbf{x}^{(\ell)},\ \ell \in \mathcal{L}$.

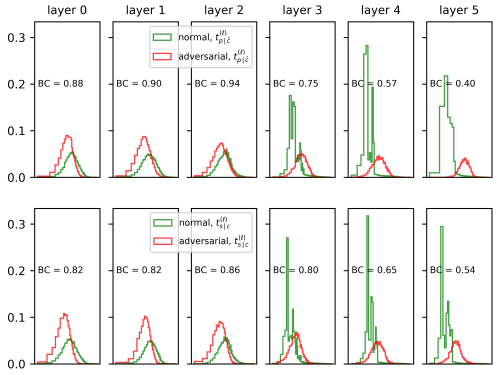

Figure 1: Distribution of test statistics corresponding to normal and adversarial samples (PGD, ℓ_∞ attack) from the layers of a DNN trained on the SVHN dataset. The top and bottom figures show the multinomial test statistic (§ 4.1) conditioned on the predicted and true class respectively. BC is the Bhattacharya coefficient, which is a measure of distribution overlap.

I. Test Statistics: In line with prior works that use the layer representations of a DNN for detection, we define test statistics at each layer that capture a statistical property of the layer representation useful for detection (*e.g.*, Mahalanobis distances (Lee et al., 2018)). The test statistics are defined to be conditioned on the predicted class and on each candidate true class (since the true class is unknown). The latter is particularly useful for adversarial samples, since they are known to have originated from one of the m classes. For a test input $\mathbf{x}$ predicted as class $\hat{c}$, the test statistic at any layer ℓ conditioned on the predicted class is defined as $T(\mathbf{x}^{(\ell)}, \hat{c}, \widehat{\mathcal{D}}_a(\ell, \hat{c}))$. It captures how anomalous the layer representation $\mathbf{x}^{(\ell)}$ is with respect to the distribution of natural inputs predicted into class $\hat{c}$ by the DNN. Similarly, a set of m test statistics at layer ℓ, conditioned on each candidate true class c, are defined as $T(\mathbf{x}^{(\ell)}, c, \mathcal{D}_a(\ell, c)),\ c \in [m]$. They capture how anomalous $\mathbf{x}^{(\ell)}$ is with respect to the distribution of natural inputs from a true class c. A mild requirement on the definition of the test statistic function is that larger values of $T(\cdot)$ correspond to larger deviations of the layer representation from the class-conditional distribution of natural inputs. When the input is clear from the context, we denote the test statistic random variables by $T^{(\ell)}_{p\,|\,\hat{c}} := T(\mathbf{x}^{(\ell)}, \hat{c}, \widehat{\mathcal{D}}_a(\ell, \hat{c}))$ and $T^{(\ell)}_{s\,|\,c} := T(\mathbf{x}^{(\ell)}, c, \mathcal{D}_a(\ell, c))$. Specific values of the test statistic are denoted by $t^{(\ell)}_{p\,|\,\hat{c}}$ and $t^{(\ell)}_{s\,|\,c}$. The vector of test statistics across the layers is defined as $\mathbf{t}_{p\,|\,\hat{c}} := [t^{(0)}_{p\,|\,\hat{c}}, \cdots, t^{(L)}_{p\,|\,\hat{c}}]$ given the predicted class $\hat{c}$, and as $\mathbf{t}_{s\,|\,c} := [t^{(0)}_{s\,|\,c}, \cdots, t^{(L)}_{s\,|\,c}],\ \forall c \in [m]$ given each candidate true class. In § 4.1, we propose a test statistic based on the multinomial likelihood ratio test (LRT) applied to

class counts from the k-nearest neighbors (kNN) of a layer representation. However, the above definitions are general and apply to test statistics proposed in prior works such as Gram matrix-based deviations (Sastry & Oore, 2020).

II. Distribution-Independent Normalization: In the absence of any prior assumptions, the class-conditional and marginal distributions of the test statistics are unknown and expected to change across the DNN layers (*e.g.*, see Fig. 1). Therefore, in order to effectively combine the test statistics from the layers for anomaly scoring, it is important to apply a normalizing transformation that (ideally) makes the transformed test statistics distribution independent. Some prior works partially address this using heuristic approaches such as z-score normalization (Roth et al., 2019) and scaling by the expected value (Sastry & Oore, 2020) in order to account for the distribution and range differences of the test statistics across the layers. We *propose two approaches for applying normalizing transformations* – the first one focuses on test statistics from the individual layers and layer pairs, and the second one focuses on the vector of test statistic across the layers. Considering test statistic pairs and the vector of test statistics allows our method to capture the joint effect of anomalous inputs on the layer representations.

Consider the first approach. In the meta-algorithm, such normalizing transformations are defined as $q(t^{(\ell)}_{p\,|\,\hat{c}})$ for the test statistic at layer ℓ conditioned on the predicted class $\hat{c}$, and $q(t^{(\ell)}_{s\,|\,c})$, $\forall c \in [m]$ for the test statistic at layer ℓ conditioned on each candidate true class. For each pair of layers $(\ell_1, \ell_2) \in \mathcal{L}^2$, $q(t^{(\ell_1)}_{p\,|\,\hat{c}}, t^{(\ell_2)}_{p\,|\,\hat{c}})$ and $q(t^{(\ell_1)}_{s\,|\,c}, t^{(\ell_2)}_{s\,|\,c})$, $\forall c \in [m]$ define the normalizing transformations for the corresponding test statistic pair conditioned on the predicted class $\hat{c}$ and on each candidate true class respectively. Since it is not efficient to include all the layer pairs beyond few tens of layers, this is specified as optional in Algorithm 1.

In the second approach, $q(\mathbf{t}_{p\,|\,\hat{c}})$ and $q(\mathbf{t}_{s\,|\,c})$, $\forall c \in [m]$ define the normalizing transformations for a vector of test statistics from the layers conditioned on the predicted class $\hat{c}$ and on each candidate true class respectively [3]. In § 4.2, we propose specific realizations for each case of the above normalizing transformations based on *class-conditional p-values*. They have the advantage of being nonparametric, and transform the test statistics into probabilities that, for natural inputs, will be approximately uniform on $[0, 1]$.

III. Layerwise Aggregation and Scoring: The normalized test statistics can be interpreted as anomaly scores that are each based on information from one or more layer representations and a specific (predicted or true) class. The goal of a scoring function is to aggregate the multiple anomaly scores in a principled way such that the *combined score is low* for inputs following the same distribution as normal inputs to the DNN, and *high for anomalies*. Prior works have taken approaches such as average or maximum of the normalized test statistics (Miller et al., 2019; Sastry & Oore, 2020), or a weighted sum of unnormalized test statistics, with the weights trained using a binary logistic classifier (Lee et al., 2018; Ma et al., 2018; Yang et al., 2020). In our meta-algorithm, we define an aggregation function $r(\cdot)$ that combines the set of all normalized test statistics from the individual layers and (optionally) layer pairs as follows: $q_{\text{agg}}(\mathbf{t}_{p\,|\,\hat{c}}) = r(Q_{p\,|\,\hat{c}})$ and $q_{\text{agg}}(\mathbf{t}_{s\,|\,c}) = r(Q_{s\,|\,c})$, $\forall c \in [m]$, where

$$
\begin{aligned}
Q_{p\,|\,\hat{c}} &= \{q(t^{(\ell)}_{p\,|\,\hat{c}}),\ \forall \ell \in \mathcal{L}\} \\
&\quad \cup \{q(t^{(\ell_1)}_{p\,|\,\hat{c}}, t^{(\ell_2)}_{p\,|\,\hat{c}}),\ \forall (\ell_1, \ell_2) \in \mathcal{L}^2\} \quad \text{and} \qquad (1) \\
Q_{s\,|\,c} &= \{q(t^{(\ell)}_{s\,|\,c}),\ \forall \ell \in \mathcal{L}\} \\
&\quad \cup \{q(t^{(\ell_1)}_{p\,|\,c}, t^{(\ell_2)}_{p\,|\,c}),\ \forall (\ell_1, \ell_2) \in \mathcal{L}^2\},\ \ \forall c \in [m]
\end{aligned}
$$

define the sets of normalized test statistics.

Motivated by ideas from multiple testing, *we propose specific aggregation functions* $r(\cdot)$ in § 4.3 for combining multiple p-values from the layers and layer pairs. For the normalization approach that directly transforms the test statistic vector from the layers, there is no need for an aggregation function; we simply set $q_{\text{agg}}(\mathbf{t}_{p\,|\,\hat{c}}) = q(\mathbf{t}_{p\,|\,\hat{c}})$ and $q_{\text{agg}}(\mathbf{t}_{s\,|\,c}) = q(\mathbf{t}_{s\,|\,c})$, $\forall c \in [m]$. The final anomaly score in the meta-algorithm is defined to be a simple function of the aggregate, normalized test statistics, *i.e.*, $S(q_{\text{agg}}(\mathbf{t}_{p\,|\,\hat{c}}), q_{\text{agg}}(\mathbf{t}_{s\,|\,1}), \cdots, q_{\text{agg}}(\mathbf{t}_{s\,|\,m}), \hat{c})$. We propose specific realizations of the score functions for adversarial and OOD detection in § 4.4.

IV. Detection Decision: The detection decision for a test input $\mathbf{x}$ predicted into class $\hat{c}$ is obtained by thresholding the final anomaly score as follows:

$$
\begin{aligned}
&\psi_\tau(\mathbf{x}^{(0)}, \cdots, \mathbf{x}^{(L)}, \hat{c}) = \\
&\mathbb{1}\left[S(q_{\text{agg}}(\mathbf{t}_{p\,|\,\hat{c}}), q_{\text{agg}}(\mathbf{t}_{s\,|\,1}), \cdots, q_{\text{agg}}(\mathbf{t}_{s\,|\,m}), \hat{c}) \geq \tau\right] \qquad (2)
\end{aligned}
$$

where decisions 0 and 1 correspond to natural and anomalous inputs respectively, and $\mathbb{1}[\cdot]$ is the indicator function. The threshold τ is usually set based on a false positive rate (FPR) that is suitable for the target application. In order to operate the detector at an FPR $\alpha \in (0, 1)$ (usually a small value *e.g.*, 0.01), the threshold can be set by estimating the FPR $\widehat{P}_{\text{F}}(\tau)$ from the set of natural inputs $\mathcal{D}_a$ as follows:

$$
\begin{aligned}
\tau_\alpha &= \sup\{\tau \in \mathbb{R} : \widehat{P}_{\text{F}}(\tau) \leq \alpha\}, \text{ where} \\
\widehat{P}_{\text{F}}(\tau) &= \frac{1}{N} \sum_{n=1}^{N} \psi_\tau(\mathbf{x}_n^{(0)}, \cdots, \mathbf{x}_n^{(L)}, \hat{c}_n). \qquad (3)
\end{aligned}
$$

This threshold choice ensures that natural inputs are accepted by the detector with a probability close to $1 - \alpha$.

[3] Although the function $q(\cdot)$ is different depending on the class, layer(s), and the number of test statistic inputs, we use the same overloaded notation for clarity.

4. A Realization of the Meta-algorithm

In this section, we propose concrete methods for realizing the components of the anomaly detection meta-algorithm.

4.1. Test Statistic Based on kNN Class Counts

Consider a set of natural inputs to the DNN that are predicted into a class $\hat{c} \in [m]$. The class counts from the kNN of its representations from a layer $\ell \in \mathcal{L}$ are expected to follow a certain distribution, wherein class $\hat{c}$ has a higher probability than the other classes. A similar observation can be made for natural inputs from a candidate true class $c \in [m]$. Let $(k_1^{(\ell)}, \cdots, k_m^{(\ell)})$ denote the tuple of class counts from the kNN $N_k^{(\ell)}(\mathbf{x}^{(\ell)})$ of a layer representation $\mathbf{x}^{(\ell)}$, such that $k_i^{(\ell)} \in \{0, 1, \cdots, k\}$ and $\sum_{i=1}^m k_i^{(\ell)} = k$. The null (natural) distribution of the kNN class counts at a layer ℓ conditioned on the predicted class $\hat{c}$ can be modeled using the following multinomial distribution

$$p(k_1^{(\ell)}, \cdots, k_m^{(\ell)} \,|\, \widehat{C} = \hat{c}) \;=\; k! \prod_{i=1}^{m} \frac{[\pi_{i\,|\,\hat{c}}^{(\ell)}]^{k_i^{(\ell)}}}{k_i^{(\ell)}!}, \quad (4)$$

where $(\pi_{1\,|\,\hat{c}}^{(\ell)}, \cdots, \pi_{m\,|\,\hat{c}}^{(\ell)})$ are the multinomial probability parameters specific to class $\hat{c}$ and layer ℓ (they are non-negative and sum to 1). We estimate these parameters from the labeled subset $\widehat{\mathcal{D}}_a(\ell, \hat{c})$ using maximum-a-posteriori (MAP) estimation with the Dirichlet conjugate prior distribution (Barber, 2012) [4]. For a test input $\mathbf{x}$ sampled from the natural data distribution that is predicted into class $\hat{c}$ by the DNN, we expect the multinomial distribution (4) to be a good fit for the class counts observed from its kNN at layer ℓ. In order to test whether the observed class counts $(k_1^{(\ell)}, \cdots, k_m^{(\ell)})$ from layer ℓ given predicted class $\hat{c}$ are consistent with distribution (4), we apply the well-known multinomial LRT (Read & Cressie, 2012), whose log-likelihood ratio statistic is given by

$$T(\mathbf{x}^{(\ell)}, \hat{c}, \widehat{\mathcal{D}}_a(\ell, \hat{c})) \;=\; \sum_{i=1}^{m} k_i^{(\ell)} \log \frac{k_i^{(\ell)}}{k\, \pi_{i\,|\,\hat{c}}^{(\ell)}}. \quad (5)$$

This test statistic is a class count deviation measure which is always non-negative, with larger values corresponding to a larger deviation from the null distribution (4). In a similar way, the test statistics conditioned on each candidate true class are defined as

$$T(\mathbf{x}^{(\ell)}, c, \mathcal{D}_a(\ell, c)) \;=\; \sum_{i=1}^{m} k_i^{(\ell)} \log \frac{k_i^{(\ell)}}{k\, \widetilde{\pi}_{i\,|\,c}^{(\ell)}}, \; \forall c \in [m]. \quad (6)$$

Here, $(\widetilde{\pi}_{1\,|\,c}^{(\ell)}, \cdots, \widetilde{\pi}_{m\,|\,c}^{(\ell)})$ are the multinomial parameters specific to class c and layer ℓ, which are estimated from the corresponding data subset $\mathcal{D}_a(\ell, c)$.

[4] We set the prior counts of the Dirichlet distribution to a small non-zero value to avoid 0 estimates for the multinomial parameters.

4.2. Normalizing Transformations Based on p-values

Recall that we are interested in designing a normalizing transformation that, for natural inputs, makes the transformed test statistics across the layers and classes follow the same distribution. One such approach is to use the p-value, that calculates the probability of a test statistic taking values (as or) more extreme than the observed value. More generally, a p-value is defined as any transformation of the test statistic (possibly a vector), following the null hypothesis distribution, into a uniformly distributed probability (Root et al., 2016). This provides a simple approach for normalizing the class-conditional test statistics in both the univariate and multivariate cases, as discussed next.

A. p-values at Individual Layers and Layer Pairs

For an input predicted into class $\hat{c}$ with class-conditional test statistics at a layer ℓ given by $t_{p\,|\,\hat{c}}^{(\ell)}, t_{s\,|\,1}^{(\ell)}, \cdots, t_{s\,|\,m}^{(\ell)}$, the normalizing p-value transformations are defined as [5]:

$$\begin{aligned} q(t_{p\,|\,\hat{c}}^{(\ell)}) \;&=\; \mathbb{P}(\, T_{p\,|\,\hat{c}}^{(\ell)} \geq t_{p\,|\,\hat{c}}^{(\ell)} \,|\, \widehat{C} = \hat{c}\,) \\ q(t_{s\,|\,c}^{(\ell)}) \;&=\; \mathbb{P}(\, T_{s\,|\,c}^{(\ell)} \geq t_{s\,|\,c}^{(\ell)} \,|\, C = c\,), \;\; \forall c \in [m]. \end{aligned} \quad (7)$$

Similarly the normalizing p-value transformations for test statistic pairs from layers (ℓ_1, ℓ_2) are defined as:

$$\begin{aligned} q(t_{p\,|\,\hat{c}}^{(\ell_1)}, t_{p\,|\,\hat{c}}^{(\ell_2)}) \;&=\; \mathbb{P}(\, T_{p\,|\,\hat{c}}^{(\ell_1)} \geq t_{p\,|\,\hat{c}}^{(\ell_1)}, T_{p\,|\,\hat{c}}^{(\ell_2)} \geq t_{p\,|\,\hat{c}}^{(\ell_2)} \,|\, \widehat{C} = \hat{c}\,) \\ q(t_{s\,|\,c}^{(\ell_1)}, t_{s\,|\,c}^{(\ell_2)}) \;&= \\ \mathbb{P}(\, T_{s\,|\,c}^{(\ell_1)} \geq t_{s\,|\,c}^{(\ell_1)}, T_{s\,|\,c}^{(\ell_2)} \geq t_{s\,|\,c}^{(\ell_2)} \,|\, C = c\,), \;\; \forall c \in [m]. \end{aligned} \quad (8)$$

Since the class-conditional distributions of the test statistics are unknown, we estimate the p-values using the empirical cumulative distribution function of the test statistics calculated from the corresponding data subsets of $\mathcal{D}_a$ [6].

B. Multivariate p-value Based Normalization

In this approach, we consider the class-conditional joint density of a test statistic vector from the layers, and propose a normalizing transformation $q : \mathbb{R}^{L+1} \mapsto [0, 1]$ based on the idea of multivariate p-values. Consider an input predicted into a class $\hat{c}$, that has a vector of test statistics $\mathbf{t}_{p\,|\,\hat{c}} = \mathbf{t}$ from the layers. Suppose $f_0(\mathbf{t}_{p\,|\,\hat{c}} \,|\, \hat{c})$ denotes the true null-hypothesis density of $\mathbf{t}_{p\,|\,\hat{c}}$ conditioned on the predicted class $\hat{c}$, then the region outside the level set of constant density equal to $f_0(\mathbf{t} \,|\, \hat{c})$ is given by $\{\mathbf{t}_{p\,|\,\hat{c}} \in \mathbb{R}^{L+1} \;:\; f_0(\mathbf{t}_{p\,|\,\hat{c}} \,|\, \hat{c}) \;<\; f_0(\mathbf{t} \,|\, \hat{c})\}$. The multivariate p-value for $\mathbf{t}$ is the probability of this region under the null hypothesis probability measure (Root et al., 2016).

We use the averaged localized p-value estimation method using kNN graphs (aK-LPE) proposed by (Qian & Saligrama,

[5] We use one-sided, right-tailed p-values since larger values of the test statistic correspond to a larger deviation.

[6] In our implementation, we averaged the p-value estimates from a hundred bootstrap samples in order to reduce the variance.

2012). The main idea is to define a score function based on nearest neighbor graphs $G(\mathbf{t})$ that captures the local relative density around $\mathbf{t}$. They show that a score function defined as the average distance from $\mathbf{t}$ to its $\frac{k}{2}$-th through $\frac{3k}{2}$-th nearest neighbors provides the following asymptotically-consistent p-value estimate:

$$q_{\text{lpe}}(\mathbf{t}) \;=\; \frac{1}{|\mathcal{D}_t|} \sum_{\mathbf{t}_n \in \mathcal{D}_t} \mathbb{1}[G(\mathbf{t}) \leq G(\mathbf{t}_n)], \qquad (9)$$

where $\mathcal{D}_t$ is a large sample of test statistic vectors corresponding to natural inputs. In our problem, we apply the above p-value transformation (using the appropriate data subsets) to normalize the $m+1$ test statistic vectors giving: $q_{\text{lpe}}(\mathbf{t}_{p\,|\,\hat{c}}), q_{\text{lpe}}(\mathbf{t}_{s\,|\,1}), \cdots, q_{\text{lpe}}(\mathbf{t}_{s\,|\,m})$.

4.3. Aggregation of p-values

The p-value based normalized test statistics capture the extent of deviation of the test statistics of an input relative to their distribution on natural inputs; smaller p-values correspond to a larger deviation. For approach A in § 4.2, we can consider each p-value to correspond to a hypothesis test involving a particular layer or layer pair. We are interested in combining the evidence from these multiple tests (Dudoit & Van Der Laan, 2007) into a single p-value for the overall problem of testing for natural versus anomalous inputs. We investigate two methods for combining p-values from multiple tests and define the corresponding aggregation functions. Note that there is no need to aggregate p-values for approach B in § 4.2, and we simply set $q_{\text{agg}}(\mathbf{t}_{p\,|\,\hat{c}}) = q_{\text{lpe}}(\mathbf{t}_{p\,|\,\hat{c}})$ and $q_{\text{agg}}(\mathbf{t}_{s\,|\,c}) = q_{\text{lpe}}(\mathbf{t}_{s\,|\,c}),\ \forall c \in [m]$.

Fisher's method (Fisher, 1992) provides a principled way of combining p-values from multiple independent tests based on the idea that, under the null hypothesis, the sum of the log of multiple p-values follows a χ^2-distribution. The aggregate p-value function based on this method is given by

$$\log\, q_{\text{fis}}(\mathbf{t}) \;=\; \log\, r(Q) \;=\; \sum_{q \in Q} \log q, \qquad (10)$$

where Q is one of the sets $Q_{p\,|\,\hat{c}}$ or $Q_{s\,|\,c}$ defined in Eq. (1), and $\mathbf{t}$ is the corresponding test statistic vector [7]. An apparent weakness of Fisher's method is its assumption of independent p-values. We briefly provide the aggregate p-value function for an alternate harmonic mean p-value (HMP) method for combining p-values from multiple dependent tests (Wilson, 2019), and discuss its details in Appendix B.

$$q_{\text{hmp}}(\mathbf{t})^{-1} \;=\; r(Q)^{-1} \;=\; \sum_{q \in Q} q^{-1}. \qquad (11)$$

4.4. Scoring for Adversarial and OOD Detection

We propose different score functions for detecting adversarial and general OOD inputs. An adversarial input predicted into class $\hat{c}$ by the DNN is expected to be anomalous at one or more of its layer representations relative to the distribution of natural inputs predicted into the same class. This implies that its aggregate p-value conditioned on the predicted class, $q_{\text{agg}}(\mathbf{t}_{p\,|\,\hat{c}})$, should have a small value. Moreover, since the adversarial input was created from a source class different from $\hat{c}$, it is expected to be a typical sample relative to the distribution of natural inputs from the unknown source class $c \neq \hat{c}$. This implies that its aggregate p-value conditioned on a candidate true class (different from $\hat{c}$) should have a relatively large value. Combining these ideas, we define the score function for adversarial inputs as

$$S(q_{\text{agg}}(\mathbf{t}_{p\,|\,\hat{c}}), q_{\text{agg}}(\mathbf{t}_{s\,|\,1}), \cdots, q_{\text{agg}}(\mathbf{t}_{s\,|\,m}), \hat{c}) \;=\; \log \left(\frac{\max\limits_{c \in [m] \setminus \{\hat{c}\}} q_{\text{agg}}(\mathbf{t}_{s\,|\,c})}{q_{\text{agg}}(\mathbf{t}_{p\,|\,\hat{c}})} \right). \qquad (12)$$

The table below provides additional insight on this score function by considering the numerator and denominator terms (inside the log) for different categories of input.

Input type & prediction	Numerator	Denominator	Score
$\mathbf{x}$ natural, $\widehat{C}(\mathbf{x}) = c$	Low	High	Low
$\mathbf{x}$ natural, $\widehat{C}(\mathbf{x}) \neq c$	High	High	Medium
$\mathbf{x}$ adversarial, $\widehat{C}(\mathbf{x}) \neq c$	High	Low	High

Similar to adversarial inputs, OOD inputs are also expected to exhibit anomalous patterns at the layers of the DNN relative to the distribution of natural inputs predicted into the same class. Since OOD inputs are not created by intentionally perturbing natural inputs from a true class different from the predicted class, we simplify score function (12) for OOD detection as follows

$$S(q_{\text{agg}}(\mathbf{t}_{p\,|\,\hat{c}}), q_{\text{agg}}(\mathbf{t}_{s\,|\,1}), \cdots, q_{\text{agg}}(\mathbf{t}_{s\,|\,m}), \hat{c}) \;=\; -\log\, q_{\text{agg}}(\mathbf{t}_{p\,|\,\hat{c}})$$

OOD inputs are expected to have a low aggregate p-value $q_{\text{agg}}(\mathbf{t}_{p\,|\,\hat{c}})$, and hence a high value for the above score.

4.5. Implementation and Computational Complexity

We briefly discuss some practical aspects of implementing JTLA efficiently. We apply the neighborhood preserving projection method (He et al., 2005) to perform dimensionality reduction on the DNN layer representations since they can be very high dimensional (details in Appendix D.3). In order to efficiently construct and query from kNN graphs at the layer representations, we use the fast approximate nearest neighbors method *NNDescent* (Dong et al., 2011) [8]. Together, these two techniques significantly reduce the memory utilization and running time of JTLA.

[7] $q_{\text{lpe}}(\cdot)$, $q_{\text{fis}}(\cdot)$, and $q_{\text{hmp}}(\cdot)$ are specific instances of $q_{\text{agg}}(\cdot)$.

[8] We use the following implementation of NNDescent: https://github.com/lmcinnes/pynndescent

The computational complexity of the proposed instantiation of JTLA at prediction (test) time can be expressed as $\mathrm{O}\left(L\left(d_{\max} N^{\rho} + m^2 + B\,N\right)\right)$ when layer pairs are not used, and $\mathrm{O}\left(L^2\left(d_{\max} N^{\rho} + m^2 + B\,N\right)\right)$ when layer pairs are used. Here $d_{\max}$ is the maximum dimension of the projected layer representations, m is the number of classes, N is the number of samples, B is the number of bootstrap replications used for estimating p-values, and $\rho \in (0, 1)$ is an unknown factor associated with the approximate nearest neighbor queries (that are sub-linear in N). The p-value calculation can be made faster and independent of N by pre-computing the empirical class-conditional CDFs. A comparison of the running time of JTLA with other detection methods can be found in Appendix E.6.

5. Defense-Aware Adaptive Attack

The importance of evaluating adversarial detection methods against an adaptive, defense-aware adversary has been highlighted in prior works (Carlini & Wagner, 2017a; Athalye et al., 2018; Tramèr et al., 2020). We consider a gray-box adversary that is assumed to have full knowledge of the DNN architecture and parameters, and partial knowledge of the detection method [9].

Consider a clean input sample $\mathbf{x}$ from class c that is correctly classified by the DNN. Let η_ℓ denote the distance between $\mathbf{x}^{(\ell)} = \mathbf{f}_\ell(\mathbf{x})$ and its k-th nearest neighbor from layer ℓ. The number of samples from any class i among the kNNs of $\mathbf{x}^{(\ell)}$, relative to the dataset $\mathcal{D}_a$, can be expressed as

$$k_i^{(\ell)} = \sum_{n=1\,:\,c_n=i}^{N} u(\eta_\ell - d(\mathbf{f}_\ell(\mathbf{x}), \mathbf{f}_\ell(\mathbf{x}_n))), \;\; i = 1, \cdots, m,$$

where $u(\cdot)$ is the unit step function. Consider the following probability mass function over the class labels: $p_i = k_i \,/\, \sum_{j=1}^{m} k_j$, $i \in [m]$, where $k_i = \sum_{\ell=0}^{L} k_i^{(\ell)}$ is the cumulative kNN count from class i across the layers. In order to fool a defense method relying on the kNN class counts from the layer representations, our attack finds an adversarial input $\mathbf{x}' = \mathbf{x} + \boldsymbol{\delta}$ with target class $c' \neq c$ that minimizes the following log-ratio of probabilities:

$$\log \frac{p_c}{p_{c'}} = \log k_c - \log k_{c'} = \log \sum_{\ell=0}^{L} k_c^{(\ell)} - \log \sum_{\ell=0}^{L} k_{c'}^{(\ell)}, \tag{13}$$

subject to a penalty on the norm of the perturbation $\boldsymbol{\delta}$ [10]. To address the non-smoothness arising from the step function in the class counts, we use the Gaussian (RBF) kernel $h_\sigma(\mathbf{x}, \mathbf{y}) = e^{-\frac{1}{\sigma^2} d(\mathbf{x},\mathbf{y})^2}$ to obtain a smooth approximation of the class counts $k_i^{(\ell)}$. The attack objective function to minimize is a weighted sum of the ℓ_2-perturbation norm and the kernel-smoothed log-ratio of probabilities, given by

$$J(\boldsymbol{\delta}) = \|\boldsymbol{\delta}\|_2^2 + \lambda \log \sum_{\ell=0}^{L} \sum_{\substack{n=1\,:\\ c_n=c}}^{N} h_{\sigma_\ell}(\mathbf{f}_\ell(\mathbf{x}+\boldsymbol{\delta}), \mathbf{f}_\ell(\mathbf{x}_n)) - \lambda \log \sum_{\ell=0}^{L} \sum_{\substack{n=1\,:\\ c_n=c'}}^{N} h_{\sigma_\ell}(\mathbf{f}_\ell(\mathbf{x}+\boldsymbol{\delta}), \mathbf{f}_\ell(\mathbf{x}_n)). \tag{14}$$

Here $\sigma_\ell > 0$ is the kernel scale for layer ℓ and $\lambda > 0$ is a constant that sets the relative importance of the terms in the objective function. The method used for setting the kernel scale per layer and minor extensions of the proposed attack are described in Appendix C. Details of the optimization method and the choice of λ are given in Appendix D.4. In our experiments, we chose the class with the second highest probability predicted by the DNN as the target attack class.

6. Experimental Results

We evaluated JTLA on the following well-known image classification datasets: CIFAR-10 (Krizhevsky et al., 2009), SVHN (Netzer et al., 2011), and MNIST (LeCun et al., 1998). We used the training partition provided by the datasets for training standard CNN architectures, including a Resnet for CIFAR-10. We performed class-stratified 5-folds cross-validation on the test partition provided by the datasets; the training folds are used for estimating the detector parameters, and the test folds are used solely for calculating performance metrics (which are then averaged across the test folds). We used the Foolbox library (Rauber et al., 2017) for generating adversarial samples from the following attack methods: (i) Projected Gradient Descent (PGD) with ℓ_∞ norm (Madry et al., 2018), (ii) Carlini-Wagner (CW) attack with ℓ_2 norm (Carlini & Wagner, 2017b), and (iii) Fast gradient sign method (FGSM) with ℓ_∞ norm (Goodfellow et al., 2015). We also implement and generate adversarial samples from the adaptive attack proposed in § 5. More details on the datasets, DNN architectures, and the attack parameters used are provided in Appendix D.

Methods Compared. We evaluated the following two variants of JTLA using the multinomial test statistic: 1) p-value normalization at the layers and layer pairs using Fisher's method for aggregation, 2) multivariate p-value normalization based on the aK-LPE method. The score functions from § 4.4 for adversarial and OOD detection are used for the respective tasks. The number of nearest neighbors is the **only hyperparameter** of the proposed instantiation of JTLA. This is set to be a function of the number of in-distribution training samples n using the heuristic $k = \lceil n^{0.4} \rceil$.

We compared against the following recently-proposed methods: (i) Deep Mahalanobis detector (Mahalanobis) (Lee et al., 2018), (ii) Local Intrinsic Dimensionality detector

[9] For example, the detection threshold and specific layers of the DNN used may be unknown.

[10] A similar type of attack on kNN-based models has been recently proposed in (Sitawarin & Wagner, 2020).

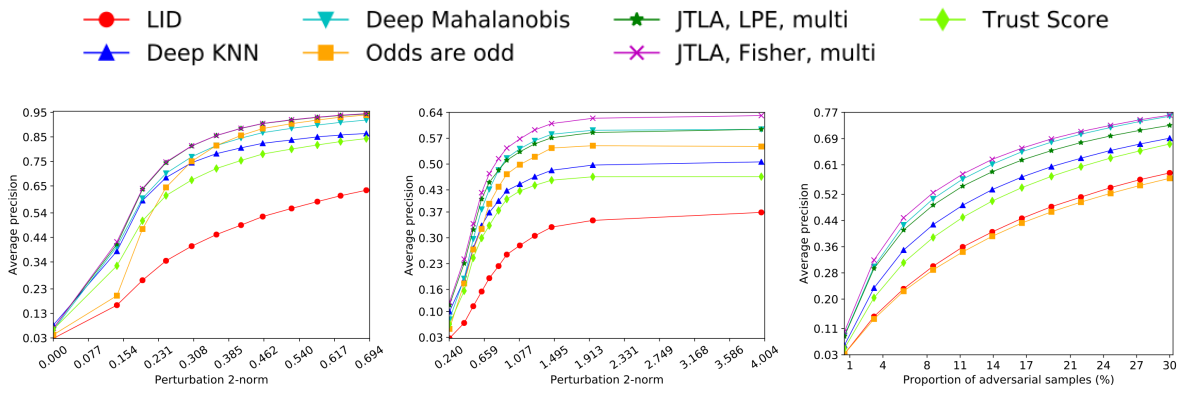

(a) CW, confidence $= 0$ (b) Adaptive attack (c) PGD, $\epsilon = 1/255$

Figure 2: Adversarial detection performance on CIFAR-10 under different attacks.

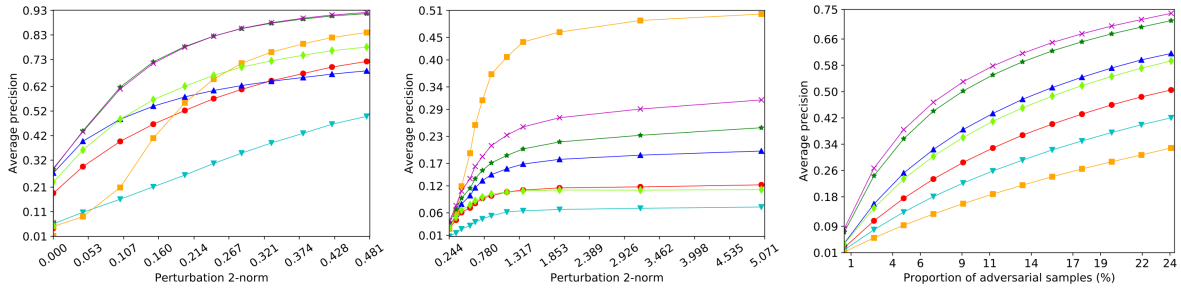

(a) CW, confidence $= 0$ (b) Adaptive attack (c) PGD, $\epsilon = 1/255$

Figure 3: Adversarial detection performance on SVHN under different attacks.

(`LID`) (Ma et al., 2018), (iii) The odds are odd detector (`Odds`) (Roth et al., 2019), (iv) Deep kNN (`DKNN`) (Papernot & McDaniel, 2018), and (v) Trust Score (`Trust`) (Jiang et al., 2018). `Mahalanobis` and `LID` are supervised (they utilize adversarial or outlier data from the training folds), while the remaining methods are unsupervised. `LID` and `Odds` are excluded from the OOD detection experiment because they specifically address adversarial samples. Details on the implementation, hyperparameters, and layer representations used by the methods can be found in Appendix D.3.

Performance metrics. We evaluate detection performance using the precision-recall (PR) curve (Davis & Goadrich, 2006; Flach & Kull, 2015) and the receiver operating characteristic (ROC) curve (Fawcett, 2006). We use *average precision* as a threshold-independent metric to summarize the PR curve, and partial area under the ROC curve below FPR α (*pAUC*-α) as the metric to summarize low-FPR region of the ROC curve. Note that both the metrics do not require the selection of a threshold. We do not report the area under the entire ROC curve because it is skew-insensitive and tends to have optimistic values when the fraction of anomalies is very small (Ahmed & Courville, 2020).

6.1. Detecting Adversarial Samples

Figures 2 and 3 show the average precision of the detection methods as a function of the perturbation ℓ_2 norm of the adversarial samples generated by the CW (confidence $= 0$) and adaptive attack methods. For the PGD attack ($\epsilon = 1 / 255$), the proportion of adversarial samples is shown on the x-axis instead of the perturbation norm because most of the samples from this attack have the same norm value. We observe that in almost all cases, `JTLA` outperforms the other baselines. Methods such as `Mahalanobis`, `Odds`, and `DKNN` perform well in some cases but fail on others, while `LID` performs poorly in nearly all scenarios. We observe an outlying trend in Figure 3b, where `Odds` outperforms `JTLA` on the adaptive attack applied to SVHN. However, a comparison of the pAUC-0.2 metric for this scenario (Figure 10 in Appendix E.4) reveals that `JTLA` has higher pAUC-0.2 for low perturbation norm (where adversarial samples are likely to be more realistic and harder to detect). We provide additional results in Appendix E that include: (i) attack transfer and attacks of varying strength, (ii) evaluation of the pAUC-0.2 metric, (iii) results on the MNIST dataset, and (iv) results on the FGSM attack.

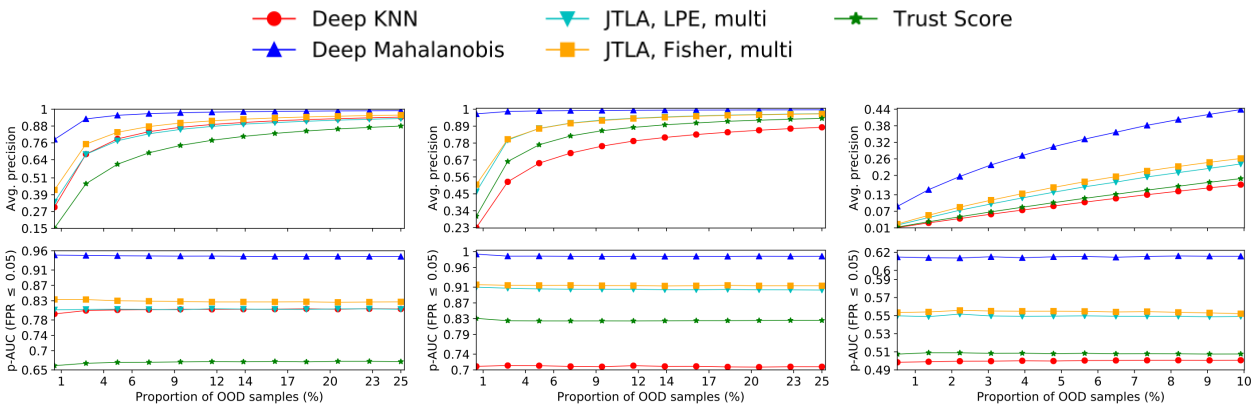

(a) CIFAR-10 vs. SVHN (b) MNIST vs. not-MNIST (c) CIFAR-10 vs. CIFAR-100

Figure 4: Comparison of OOD detection performance.

6.2. Detecting Out-Of-Distribution Samples

We evaluated OOD detection using the following image dataset pairs, with the first dataset used as in-distribution (inliers) and the second dataset used as out-distribution (outliers): 1) MNIST vs. Not-MNIST (Bulatov, 2011), 2) CIFAR-10 vs. SVHN. While a majority of papers on OOD detection evaluate using such dataset pairs, the importance of evaluating against outliers that are semantically meaningful has been emphasized by Ahmed & Courville (2020). Therefore, we performed an experiment where object classes from the CIFAR-100 dataset (not in CIFAR-10) are treated as outliers relative to CIFAR-10. This is a more realistic and challenging task since novel object categories can be considered semantically-meaningful anomalies. Using the same 5-fold cross-validation setup, we compared the performance of `JTLA` with `Mahalanobis`, `DKNN`, and `Trust` (`Odds` and `LID` are excluded because they focus on adversarial examples). Since `Mahalanobis` is a supervised method, it uses both inlier and outlier data from the training folds, while the remaining methods (all unsupervised) use only the inlier data from the training folds. To promote fairness in the comparison, we excluded outlier data corresponding to one half of the classes from the training folds, and included outlier data from *only* the excluded classes in the test folds. Additionally, in the test folds we included image samples with random pixel values uniformly selected from the same range as valid images. The number of random samples is set equal to the average number of test-fold samples from a single class. Figure 4 shows the average precision and pAUC-0.05 as a function of the proportion of OOD samples on the OOD detection tasks [11]. We observe that `JTLA` outperforms the unsupervised methods `DKNN` and `Trust` in all cases, but does not achieve the very good performance of `Mahalanobis`. This should be considered in light of the fact that `Mahalanobis` uses outlier samples from the training folds to train a classifier and tune a noise parameter (Lee et al., 2018). However, in real-world settings, one is unlikely to have the prior knowledge and sufficient number (and variety) of outlier samples for training.

[11] We report pAUC below 5% FPR because the methods achieve high detection rates at very low FPR.

6.3. Ablation Studies

We performed ablation studies to gain a better understanding of the different components of the proposed method. Specifically, we evaluated 1) the relative performance of the proposed p-value based normalization and aggregation methods, 2) the performance improvement from testing at layer pairs in addition to the individual layers, 3) the relative performance of using only the last few layers compared to using all the layer representations, and 4) the relative performance of the two scoring methods in § 4.4. These results are discussed in Appendix E.2.

7. Conclusions

We presented `JTLA`, a general framework for detecting anomalous inputs to a DNN classifier based on joint statistical testing of its layer representations. We presented a general meta-algorithm for this problem, and proposed specific methods for realizing the components of this algorithm in a principled way. The construction of `JTLA` is modular, allowing it to be used with a variety of test statistics proposed in prior works. Extensive experiments with strong adversarial attacks (including an adaptive defense-aware attack we proposed) and anomalous inputs to DNN image classifiers demonstrate the effectiveness of our method.

Acknowledgements

We thank the anonymous reviewers for their useful feedback that helped improve the paper. VC, JR, and SB were supported in part through the following US NSF grants: CNS-1838733, CNS-1719336, CNS-1647152, CNS-1629833, CNS-1942014, CNS-2003129, and an award from the US Department of Commerce with award number

70NANB21H043. SJ was partially supported by Air Force Grant FA9550-18-1-0166, the NSF Grants CCF-FMitF-1836978, SaTC-Frontiers-1804648 and CCF-1652140, and ARO grant number W911NF-17-1-0405.

References

Ahmed, F. and Courville, A. C. Detecting semantic anomalies. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, New York, NY, USA, February 7-12, 2020*, pp. 3154–3162. AAAI Press, 2020.

Akhtar, N. and Mian, A. S. Threat of adversarial attacks on deep learning in computer vision: A survey. *IEEE Access*, 6:14410–14430, 2018.

Amsaleg, L., Chelly, O., Furon, T., Girard, S., Houle, M. E., Kawarabayashi, K.-i., and Nett, M. Estimating local intrinsic dimensionality. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 29–38. ACM, 2015.

Athalye, A., Carlini, N., and Wagner, D. A. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *Proceedings of the 35th International Conference on Machine Learning, ICML*, volume 80 of *Proceedings of Machine Learning Research*, pp. 274–283. PMLR, 2018.

Barber, D. *Bayesian reasoning and machine learning*. Cambridge University Press, 2012. ISBN 0521518148.

Barreno, M., Nelson, B., Sears, R., Joseph, A. D., and Tygar, J. D. Can machine learning be secure? In *Proceedings of the 2006 ACM Symposium on Information, Computer and Communications Security (ASIACCS)*, pp. 16–25. ACM, 2006. doi: 10.1145/1128817.1128824.

Biggio, B. and Roli, F. Wild patterns: Ten years after the rise of adversarial machine learning. *Pattern Recognition*, 84:317–331, 2018. doi: 10.1016/j.patcog.2018.07.023.

Biggio, B., Nelson, B., and Laskov, P. Poisoning attacks against support vector machines. In *Proceedings of the 29th International Conference on Machine Learning, ICML*. icml.cc / Omnipress, 2012.

Bulatov, Y. NotMNIST dataset. http://yaroslavvb.com/upload/notMNIST/, 2011.

Bulusu, S., Kailkhura, B., Li, B., Varshney, P. K., and Song, D. Anomalous instance detection in deep learning: A survey. *CoRR*, abs/2003.06979, 2020. URL https://arxiv.org/abs/2003.06979.

Carlini, N. and Wagner, D. A. Adversarial examples are not easily detected: Bypassing ten detection methods. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security, AISec@CCS*, pp. 3–14. ACM, 2017a. doi: 10.1145/3128572.3140444.

Carlini, N. and Wagner, D. A. Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy (S&P)*, pp. 39–57. IEEE Computer Society, 2017b. doi: 10.1109/SP.2017.49.

Chandola, V., Banerjee, A., and Kumar, V. Anomaly detection: A survey. *ACM computing surveys (CSUR)*, 41(3): 1–58, 2009.

Davis, J. and Goadrich, M. The relationship between Precision-Recall and ROC curves. In *Proceedings of the 23rd international conference on Machine learning*, pp. 233–240, 2006.

Dong, W., Moses, C., and Li, K. Efficient k-nearest neighbor graph construction for generic similarity measures. In *Proceedings of the 20th international conference on World wide web*, pp. 577–586, 2011.

Dudoit, S. and Van Der Laan, M. J. *Multiple testing procedures with applications to genomics*. Springer Science & Business Media, 2007.

Fawcett, T. An introduction to ROC analysis. *Pattern recognition letters*, 27(8):861–874, 2006.

Fawzi, A., Moosavi-Dezfooli, S., and Frossard, P. Robustness of classifiers: from adversarial to random noise. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems*, pp. 1624–1632, 2016.

Fawzi, A., Fawzi, O., and Frossard, P. Analysis of classifiers' robustness to adversarial perturbations. *Machine Learning*, 107(3):481–508, 2018.

Feinman, R., Curtin, R. R., Shintre, S., and Gardner, A. B. Detecting adversarial samples from artifacts. *CoRR*, abs/1703.00410, 2017. URL http://arxiv.org/abs/1703.00410.

Fisher, R. A. Statistical methods for research workers. In *Breakthroughs in statistics*, pp. 66–70. Springer, 1992.

Flach, P. and Kull, M. Precision-recall-gain curves: PR analysis done right. In *Advances in neural information processing systems*, pp. 838–846, 2015.

Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. In *3rd International Conference on Learning Representations, Conference Track Proceedings*, 2015.

He, K., Zhang, X., Ren, S., and Sun, J. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *IEEE International Conference*

on Computer Vision (ICCV), pp. 1026–1034. IEEE Computer Society, 2015. doi: 10.1109/ICCV.2015.123.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

He, X., Cai, D., Yan, S., and Zhang, H.-J. Neighborhood preserving embedding. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 2, pp. 1208–1213. IEEE, 2005.

Hein, M., Andriushchenko, M., and Bitterwolf, J. Why relu networks yield high-confidence predictions far away from the training data and how to mitigate the problem. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 41–50. Computer Vision Foundation / IEEE, 2019. doi: 10.1109/CVPR.2019.00013.

Hendrycks, D. and Gimpel, K. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *5th International Conference on Learning Representations, Conference Track Proceedings*. OpenReview.net, 2017.

Jha, S., Raj, S., Fernandes, S. L., Jha, S. K., Jha, S., Jalaian, B., Verma, G., and Swami, A. Attribution-based confidence metric for deep neural networks. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems*, pp. 11826–11837, 2019.

Jiang, H., Kim, B., Guan, M. Y., and Gupta, M. R. To trust or not to trust a classifier. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems*, pp. 5546–5557, 2018.

Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.

Krizhevsky, A., Sutskever, I., and Hinton, G. E. ImageNet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.

Kurakin, A., Goodfellow, I. J., and Bengio, S. Adversarial machine learning at scale. In *5th International Conference on Learning Representations, Conference Track Proceedings*. OpenReview.net, 2017.

LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Lee, K., Lee, K., Lee, H., and Shin, J. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems*, pp. 7167–7177, 2018.

Li, X. and Li, F. Adversarial examples detection in deep networks with convolutional filter statistics. In *IEEE International Conference on Computer Vision (ICCV)*, pp. 5775–5783. IEEE Computer Society, 2017.

Ma, X., Li, B., Wang, Y., Erfani, S. M., Wijewickrema, S. N. R., Schoenebeck, G., Song, D., Houle, M. E., and Bailey, J. Characterizing adversarial subspaces using local intrinsic dimensionality. In *6th International Conference on Learning Representations, Conference Track Proceedings*. OpenReview.net, 2018.

Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. In *6th International Conference on Learning Representations, Conference Track Proceedings*. OpenReview.net, 2018.

Meng, D. and Chen, H. MagNet: A two-pronged defense against adversarial examples. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pp. 135–147. ACM, 2017.

Miller, D. J., Wang, Y., and Kesidis, G. When not to classify: Anomaly detection of attacks (ADA) on DNN classifiers at test time. *Neural Computation*, 31(8):1624–1670, 2019.

Miller, D. J., Xiang, Z., and Kesidis, G. Adversarial learning targeting deep neural network classification: A comprehensive review of defenses against attacks. *Proceedings of the IEEE*, 108(3):402–433, 2020.

Moosavi-Dezfooli, S., Fawzi, A., and Frossard, P. Deepfool: A simple and accurate method to fool deep neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2574–2582. IEEE Computer Society, 2016. doi: 10.1109/CVPR.2016.282.

Moosavi-Dezfooli, S., Fawzi, A., Fawzi, O., and Frossard, P. Universal adversarial perturbations. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 86–94. IEEE Computer Society, 2017.

Muñoz-González, L., Biggio, B., Demontis, A., Paudice, A., Wongrassamee, V., Lupu, E. C., and Roli, F. Towards poisoning of deep learning algorithms with back-gradient optimization. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security, AISec@CCS*, pp. 27–38. ACM, 2017. doi: 10.1145/3128572.3140451.

Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning. 2011.

Nguyen, A. M., Yosinski, J., and Clune, J. Deep neural networks are easily fooled: High confidence predictions

for unrecognizable images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 427–436. IEEE Computer Society, 2015.

Papernot, N. and McDaniel, P. D. Deep k-nearest neighbors: Towards confident, interpretable and robust deep learning. *CoRR*, abs/1803.04765, 2018. URL http://arxiv.org/abs/1803.04765.

Papernot, N., McDaniel, P. D., Jha, S., Fredrikson, M., Celik, Z. B., and Swami, A. The limitations of deep learning in adversarial settings. In *IEEE European Symposium on Security and Privacy, EuroS&P*, pp. 372–387. IEEE, 2016. doi: 10.1109/EuroSP.2016.36.

Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. Automatic differentiation in PyTorch. 2017.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. Scikit-learn: Machine learning in Python. *Journal of machine learning research*, 12(Oct):2825–2830, 2011.

Qian, J. and Saligrama, V. New statistic in p-value estimation for anomaly detection. In *IEEE Statistical Signal Processing Workshop (SSP)*, pp. 393–396. IEEE, 2012. doi: 10.1109/SSP.2012.6319713.

Rauber, J., Brendel, W., and Bethge, M. Foolbox: A python toolbox to benchmark the robustness of machine learning models. In *Reliable Machine Learning in the Wild Workshop, 34th International Conference on Machine Learning*, 2017.

Read, T. R. and Cressie, N. A. *Goodness-of-fit statistics for discrete multivariate data*. Springer Science & Business Media, 2012.

Root, J., Saligrama, V., and Qian, J. Learning minimum volume sets and anomaly detectors from KNN graphs. *CoRR*, abs/1601.06105, 2016. URL http://arxiv.org/abs/1601.06105.

Roth, K., Kilcher, Y., and Hofmann, T. The odds are odd: A statistical test for detecting adversarial examples. In *Proceedings of the 36th International Conference on Machine Learning, ICML*, volume 97 of *Proceedings of Machine Learning Research*, pp. 5498–5507. PMLR, 2019.

Ruder, S. An overview of gradient descent optimization algorithms. *CoRR*, abs/1609.04747, 2016. URL http://arxiv.org/abs/1609.04747.

Sastry, C. S. and Oore, S. Detecting out-of-distribution examples with gram matrices. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, volume 119 of *Proceedings of Machine Learning Research*, pp. 8491–8501. PMLR, 2020.

Sitawarin, C. and Wagner, D. A. Minimum-norm adversarial examples on KNN and KNN-based models. In *IEEE Security and Privacy Workshops, SP Workshops*, pp. 34–40. IEEE, 2020. doi: 10.1109/SPW50608.2020.00023.

Steinhardt, J., Koh, P. W., and Liang, P. Certified defenses for data poisoning attacks. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems*, pp. 3517–3529, 2017.

Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I. J., and Fergus, R. Intriguing properties of neural networks. In *2nd International Conference on Learning Representations, Conference Track Proceedings*, 2014.

Tax, D. M. J. and Duin, R. P. W. Growing a multi-class classifier with a reject option. *Pattern Recognition Letters*, 29(10):1565–1570, 2008.

Tramèr, F., Carlini, N., Brendel, W., and Madry, A. On adaptive attacks to adversarial example defenses. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems*, 2020.

Wilson, D. J. The harmonic mean p-value for combining dependent tests. *Proceedings of the National Academy of Sciences*, 116(4):1195–1200, 2019.

Xu, W., Evans, D., and Qi, Y. Feature squeezing: Detecting adversarial examples in deep neural networks. In *25th Annual Network and Distributed System Security Symposium (NDSS)*. The Internet Society, 2018.

Yang, P., Chen, J., Hsieh, C., Wang, J., and Jordan, M. I. ML-LOO: Detecting adversarial examples with feature attribution. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*, pp. 6639–6647. AAAI Press, 2020.

Yuan, X., He, P., Zhu, Q., and Li, X. Adversarial examples: Attacks and defenses for deep learning. *IEEE Transactions on Neural Networks and Learning Systems*, 30(9): 2805–2824, 2019. doi: 10.1109/TNNLS.2018.2886017.

Zhao, M. and Saligrama, V. Anomaly detection with score functions based on nearest neighbor graphs. In *Advances in Neural Information Processing Systems 22: 23rd Annual Conference on Neural Information Processing Systems*, pp. 2250–2258, 2009.

Zheng, Z. and Hong, P. Robust detection of adversarial attacks by modeling the intrinsic properties of deep neural networks. In *Advances in Neural Information Processing*

Systems 31: Annual Conference on Neural Information Processing Systems, pp. 7924–7933, 2018.

Appendix

Algorithm 2 Meta-algorithm for Anomaly Detection

Inputs: Trained DNN $\mathbf{F}(\cdot)$, Dataset $\mathcal{D}$, Test input $\mathbf{x}$, FPR α or detection threshold τ.
Output: Detector decision – normal 0 or anomaly 1.

Preprocessing:
Calculate the detection threshold τ (if not specified).
Class prediction of $\mathbf{x}$: $\hat{c} = \arg\max_{c\in[m]} F_c(\mathbf{x})$.
Layer representations of $\mathbf{x}$: $\mathbf{x}^{(\ell)} = \mathbf{f}_\ell(\mathbf{x})$, $\forall \ell \in \mathcal{L}$.
Data subsets: $\widehat{\mathcal{D}}_a(\ell, \hat{c})$ and $\mathcal{D}_a(\ell, c)$, $\forall\, \ell \in \mathcal{L},\ c \in [m]$.

I. Test statistics:
  for each layer $\ell \in \mathcal{L}$:
    Calculate $t^{(\ell)}_{p\,|\,\hat{c}} = T(\mathbf{x}^{(\ell)}, \hat{c}, \widehat{\mathcal{D}}_a(\ell, \hat{c}))$.
    Calculate $t^{(\ell)}_{s\,|\,c} = T(\mathbf{x}^{(\ell)}, c, \mathcal{D}_a(\ell, c)),\ \forall c \in [m]$.
  Compile the vectors: $\mathbf{t}_{p\,|\,\hat{c}}, \mathbf{t}_{s\,|\,1}, \cdots, \mathbf{t}_{s\,|\,m}$.

II. Normalizing transformations:
  if multivariate normalization:
    Calculate $q(\mathbf{t}_{p\,|\,\hat{c}}), q(\mathbf{t}_{s\,|\,1}), \cdots, q(\mathbf{t}_{s\,|\,m})$.
  else
    for each layer $\ell \in \mathcal{L}$:
      Calculate $q(t^{(\ell)}_{p\,|\,\hat{c}}), q(t^{(\ell)}_{s\,|\,1}), \cdots, q(t^{(\ell)}_{s\,|\,m})$.
    for each layer pair $(\ell_1, \ell_2) \in \mathcal{L}^2$: **[optional]**
      Calculate $q(t^{(\ell_1)}_{p\,|\,\hat{c}}, t^{(\ell_2)}_{p\,|\,\hat{c}})$.
      Calculate $q(t^{(\ell_1)}_{s\,|\,1}, t^{(\ell_2)}_{s\,|\,1}), \cdots, q(t^{(\ell_1)}_{s\,|\,m}, t^{(\ell_2)}_{s\,|\,m})$.

III. Layerwise aggregation and scoring:
  if multivariate normalization:
    Set $q_{\text{agg}}(\mathbf{t}_{p\,|\,\hat{c}}) = q(\mathbf{t}_{p\,|\,\hat{c}})$.
    Set $q_{\text{agg}}(\mathbf{t}_{s\,|\,c}) = q(\mathbf{t}_{s\,|\,c}),\ \forall c \in [m]$.
  else
    Create the sets $Q_{p\,|\,\hat{c}}$ and $Q_{s\,|\,c},\ \forall c \in [m]$.
    Calculate $q_{\text{agg}}(\mathbf{t}_{p\,|\,\hat{c}}) = r(Q_{p\,|\,\hat{c}})$.
    Calculate $q_{\text{agg}}(\mathbf{t}_{s\,|\,c}) = r(Q_{s\,|\,c}),\ \forall c \in [m]$.
  Calculate the final score:
  $S(q_{\text{agg}}(\mathbf{t}_{p\,|\,\hat{c}}), q_{\text{agg}}(\mathbf{t}_{s\,|\,1}), \cdots, q_{\text{agg}}(\mathbf{t}_{s\,|\,m}), \hat{c})$.

IV. Detection decision:
  if $S(q_{\text{agg}}(\mathbf{t}_{p\,|\,\hat{c}}), q_{\text{agg}}(\mathbf{t}_{s\,|\,1}), \cdots, q_{\text{agg}}(\mathbf{t}_{s\,|\,m}), \hat{c}) \geq \tau$:
    Return anomaly (1).
  else
    Return normal (0).

The appendices are organized as follows.

- Appendix A provides a background on adversarial attacks and defenses, and discusses prior works in more detail.
- Appendix B describes the harmonic mean p-value (HMP) method for aggregating p-values.
- Appendix C discusses details of the adaptive attack method that were left out of § 5.
- Appendix D discusses the implementation details such as computing platform, datasets, DNN architectures, experimental setup, and the method implementation details, hyperparameters etc.
- Appendix E discusses additional experiments that we performed including ablation studies, performance on attack transfer and attacks of varying strength, and a comparison of running times.

A. Background and Related Works

We provide a background on adversarial attacks and defense methods, followed by a detailed discussion of related works.

A.1. Adversarial Attacks

Adversarial attacks may be broadly classified into training-time or data poisoning attacks, and test-time or evasion attacks (Biggio & Roli, 2018). Data poisoning attacks focus on tampering with the training set of a learning algorithm by introducing malicious input patterns that steer the learning algorithm to a sub-optimal solution, causing degradation in performance (Biggio et al., 2012; Muñoz-González et al., 2017; Steinhardt et al., 2017). Evasion attacks focus on tampering with test inputs to an already-trained ML model such that the model predicts incorrectly on them. In both cases, an adversary aims to modify the inputs in such a way that the changes are not easily perceived by a human or flagged by a detector. For example, a test image of the digit 1 may be modified by introducing minimal perturbations to the pixels such that a classifier is fooled into classifying it as the digit 7, while the image still looks like the digit 1 to a human. In this work, we focus on detecting test-time evasion attacks.

Given a test input $\mathbf{x} \in \mathcal{X}$ from class c, adversarial attack methods aim to create a minimally-perturbed input $\mathbf{x}' = \mathbf{x} + \boldsymbol{\delta}$ that is mis-classified by the classifier either into a specific class (targeted attack) or into any class other than c (untargeted attack). This is formulated as an optimization problem, which in its most general form looks like

$$\begin{aligned} &\min_{\boldsymbol{\delta}} \ \|\boldsymbol{\delta}\|_p \quad \text{s.t.} \\ &\mathbf{x} + \boldsymbol{\delta} \ \in \ \mathcal{X} \\ &\widehat{C}(\mathbf{x} + \boldsymbol{\delta}) \ = \ c' \quad \text{(targeted)} \\ \text{or } &\widehat{C}(\mathbf{x} + \boldsymbol{\delta}) \ \neq \ c \quad \text{(untargeted)} \end{aligned}$$

A number of adversarial attack methods have been proposed based on this general formulation (Szegedy et al., 2014; Goodfellow et al., 2015; Madry et al., 2018; Carlini & Wagner, 2017b; Moosavi-Dezfooli et al., 2016; Papernot et al., 2016; Kurakin et al., 2017; Moosavi-Dezfooli et al., 2017). Some of the well-known methods for generating adversarial data include the fast gradient sign method (FGSM) (Goodfellow et al., 2015), projected gradient descent (PGD) attack (Madry et al., 2018), Carlini-Wagner attack (Carlini & Wagner, 2017b), and DeepFool attack (Moosavi-Dezfooli et al., 2016). These attack methods can be categorized into *white-box* or *black-box* depending on the extent of their knowledge about the classifier's structure, parameters, loss function, and learning algorithm. Most of the commonly

Table 1: Notations and Definitions

Term	Description
$[m] := \{1, \cdots, m\}$	Set of classes.
$\mathcal{L} = \{0, \cdots, L\}$, $\mathcal{L}^2 = \{(\ell_1, \ell_2) \in \mathcal{L} \times \mathcal{L} : \ell_2 > \ell_1\}$	Set of layers and distinct layer pairs.
$\|\cdot\|_p$	ℓ_p norm of a vector.
$d(\mathbf{x}, \mathbf{y})$	Distance between a pair of vectors; cosine distance unless specified otherwise.
$\mathbb{1}[\cdot]$	Indicator function mapping an input condition to 1 if true and 0 if false.
$h_\sigma(\mathbf{x}, \mathbf{y}) = e^{-\frac{1}{\sigma^2} d(\mathbf{x}, \mathbf{y})^2}$	Gaussian or RBF kernel.
$\mathbf{F}(\mathbf{x}) = [F_1(\mathbf{x}), \cdots, F_m(\mathbf{x})]$, $\mathbf{F} : \mathcal{X} \mapsto \Delta_m$	Input-output mapping learned by the DNN classifier.
$\Delta_m = \{(p_1, \cdots, p_m) \in [0,1]^m : \sum_i p_i = 1\}$	Space of output probabilities for the m classes.
$\widehat{C}(\mathbf{x}) = \arg\max_{c \in [m]} F_c(\mathbf{x})$	Class prediction of the DNN.
$\mathbf{x}^{(\ell)} := \mathbf{f}_\ell(\mathbf{x}) \in \mathbb{R}^{d_\ell}$, $\ell = 0, 1, \cdots, L$	Flattened layer representations of the DNN ($\ell = 0$ refers to the input).
C and $\widehat{C}$	Random variables corresponding to the true class and DNN-predicted class.
$\mathcal{D}_a = \{(\mathbf{x}_n^{(0)}, \cdots, \mathbf{x}_n^{(L)}, c_n, \hat{c}_n),\ n = 1, \cdots, N\}$	Labeled dataset with the layer representations, true class, and predicted class.
$\widehat{\mathcal{D}}_a(\ell, \hat{c})$	Subset of $\mathcal{D}_a$ corresponding to layer ℓ and predicted class $\hat{c}$.
$\mathcal{D}_a(\ell, c)$	Subset of $\mathcal{D}_a$ corresponding to layer ℓ and true class c.
$T(\mathbf{x}^{(\ell)}, \hat{c}, \widehat{\mathcal{D}}_a(\ell, \hat{c})) = T^{(\ell)}_{p \mid \hat{c}}$	Test statistic from layer ℓ conditioned on the predicted class $\hat{c}$.
$T(\mathbf{x}^{(\ell)}, c, \mathcal{D}_a(\ell, c)) = T^{(\ell)}_{s \mid c}$, $c \in [m]$	Test statistic from layer ℓ conditioned on a candidate true class c.
$\mathbf{t}_{p \mid \hat{c}} = [t^{(0)}_{p \mid \hat{c}}, \cdots, t^{(L)}_{p \mid \hat{c}}]$	Vector of test statistics from the layers conditioned on the predicted class $\hat{c}$.
$\mathbf{t}_{s \mid c} = [t^{(0)}_{s \mid c}, \cdots, t^{(L)}_{s \mid c}]$, $c \in [m]$	Vector of test statistics from the layers conditioned on a candidate true class c.
$q_{\text{agg}}(\mathbf{t}_{p \mid \hat{c}}), q_{\text{agg}}(\mathbf{t}_{s \mid 1}), \cdots, q_{\text{agg}}(\mathbf{t}_{s \mid m})$	Aggregate normalized test statistic from the predicted class and candidate true classes.
$(k_1^{(\ell)}, \cdots, k_m^{(\ell)}) \in \{0, 1, \cdots, k\}^m$ s.t. $\sum_i k_i^{(\ell)} = k$	Class counts from the k-nearest neighbors of a representation vector from layer ℓ.
$N_k^{(\ell)}(\mathbf{x}^{(\ell)}) \subset \{1, \cdots, N\}$	Index set of the k-nearest neighbors of a layer representation $\mathbf{x}^{(\ell)}$ relative to $\mathcal{D}_a$.

used test-time attacks on DNN classifiers are strongly white-box in that they assume a complete knowledge of the system. For a detailed discussion and taxonomy of adversarial attack methods, we refer readers to the recent surveys (Akhtar & Mian, 2018; Yuan et al., 2019).

A.2. Adversarial Defenses

On the defense side of adversarial learning, the focus can be broadly categorized into (1) *adversarial training* - where the objective is to train robust classifiers that generalize well in the face of adversarial attacks (Madry et al., 2018; Fawzi et al., 2016; 2018; Goodfellow et al., 2015), (2) *adversarial detection* - where the objective is to detect inputs that are adversarially manipulated by identifying signatures or patterns that make them different from natural inputs seen by the classifier at training time (Feinman et al., 2017; Xu et al., 2018; Lee et al., 2018; Ma et al., 2018). One approach to adversarial training involves augmenting the training set of the classifier with adversarial samples created from one or multiple attack methods along with their true labels, and retraining the classifier on the augmented training set (possibly initialized with parameters from a prior non-adversarial training) (Goodfellow et al., 2015). A limitation of this approach is that the resulting classifier may still fail on attacks that were not seen by the classifier during training. This has lead to research in the direction of robust optimization, where the learning objective of the classifier (usually an empirical risk function) is modified into a min-max optimization problem, with the inner maximization performed on a suitably chosen perturbation set (*e.g.*, an ℓ_∞ norm ball) (Madry et al., 2018). Adversarial detection, on the other hand, does not usually involve special training or modification of the underlying classifier to predict accurately on adversarial inputs. Instead, the focus is on methods that can detect adversarial inputs while operating at low false positive (natural inputs detected as adversarial) rates using ideas from the anomaly detection literature. The detector flags inputs that are suspicious and likely to be misclassified by the classifier so that they may be analyzed by an expert (possibly another ML system) for decision making down the line. There have been a plethora of works on the defense side adversarial learning. Recent surveys on this topic can be found in (Biggio & Roli, 2018; Miller et al., 2020).

A.3. Adversarial Samples as Anomalies

Adversarial detection is closely related to the problem of anomaly or outlier detection (Chandola et al., 2009; Zhao & Saligrama, 2009) with some important distinctions. The objective of anomaly detection is to determine if an input follows a pattern that is significantly different from that observed on a given data set. Stated differently, an input is said to be anomalous if it has a very low probability under the reference marginal distribution $p_0(\mathbf{x})$ underlying a given data set. On the other hand, adversarial in-

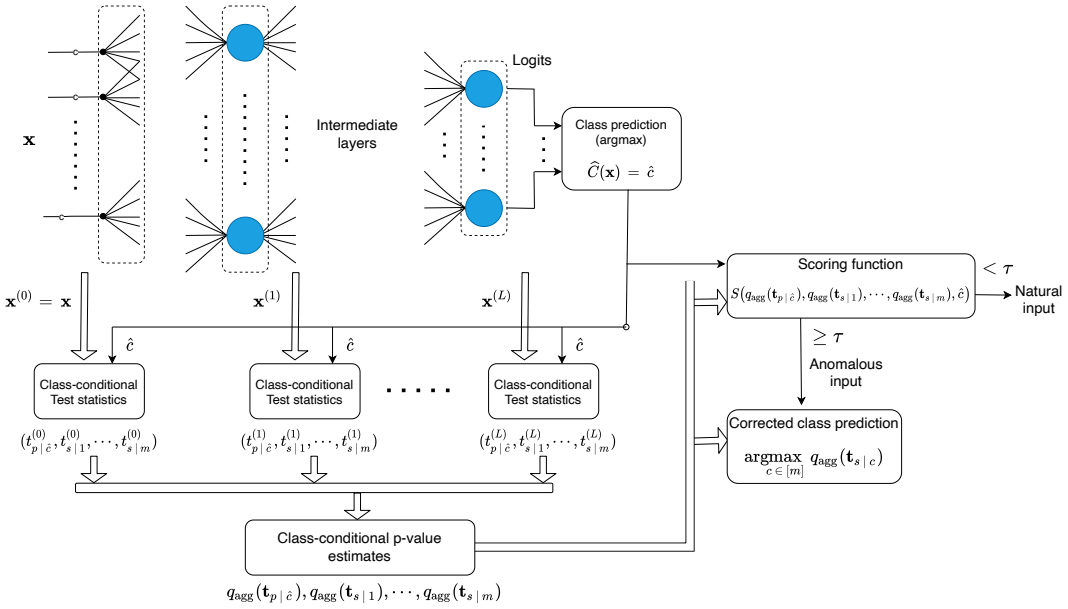

Figure 5: Overview of the proposed detection framework. The DNN classifies a test input $\mathbf{x}$ into one of the m classes. JTLA uses the predicted class and the layer representations to calculate $m+1$ class-conditional test statistics at each layer. The test statistics from the layers are aggregated into normalized p-value estimates, which are then used by the scoring function to determine whether the input is natural or anomalous.

puts from a test-time attack are not necessarily anomalous with respect to the marginal data distribution because of the way they are created by minimally perturbing a valid input $\mathbf{x} \sim p_0(\cdot \,|\, c)$ from a given class c. It is useful to consider the following notion of adversarial inputs. Suppose a clean input $\mathbf{x}$ from class c (*i.e.*, $\mathbf{x} \sim p_0(\cdot \,|\, c)$) is perturbed to create an adversarial input $\mathbf{x}'$ that appears to be from the same class c according to the true (Bayes) class posterior distribution, *i.e.*, $c = \arg\max_k p_0(k \,|\, \mathbf{x}')$. However, it is predicted into a different class c' by the classifier, *i.e.*, $\widehat{C}(\mathbf{x}') = \arg\max_k F_k(\mathbf{x}') = c'$. From this standpoint, we hypothesize that an adversarial input $\mathbf{x}'$ is likely to be a typical sample from the conditional distribution $p_0(\cdot \,|\, c)$ of the true class, while it is also likely to be an anomalous sample from the conditional distribution $p_0(\cdot \,|\, c')$ of the predicted class. We note that (Miller et al., 2019) use a similar hypothesis to motivate their detection method.

A.4. Related Works

We categorize and review some closely-related prior works on adversarial and OOD detection. For works that are based on multiple layer representations of a DNN, we discuss how the methods fit into the proposed meta-algorithm for anomaly detection.

Supervised Methods

In (Lee et al., 2018), the layer representations of a DNN are modeled class-conditionally using multivariate Gaussian densities, which leads to the Mahalanobis distance confidence score being used as a test statistic (feature) at the layers. The feature vector of Mahalanobis distances from the layers is used to train a binary logistic classifier for discriminating adversarial (or OOD samples) from natural samples. While this method uses the class-conditional densities of the layer representations, it uses a rather simple parametric model based on multivariate Gaussians, which may not be suitable for the intermediate layer representations (issues include high dimensionality, non-negative activations, and multimodality). Also, this method does not use the predicted class of the DNN; instead it finds the "closest" class corresponding to each layer representation. In the context of the meta-algorithm, this method does not explicitly apply a normalizing transformation to the test statistics. However, the Mahalanobis distance can be interpreted as the negative-log-density, which follows the Chi-squared distribution (at each layer) under the Gaussian density assumption. Finally, we note that using a binary logistic classifier is equivalent to using a weighted linear combination of the test statistics for scoring as shown in Appendix A.5.

(Ma et al., 2018) propose using the local intrinsic dimensionality (LID) of the layer representations as a test statistic for characterizing adversarial samples. Similar to the approach of (Lee et al., 2018), they calculate a test statistic vector of LID estimates from the layer representations, which is used for training a logistic classifier for discriminating adversarial samples from natural samples. In the context of the meta-algorithm, this method does not calculate class-

conditional test statistics. The LID is estimated based on the marginal distribution of the layer representation manifold. The LID test statistics are not normalized in any way, and use of the logistic classifier implies that this method also uses a weighted linear combination of the test statistics for scoring.

In (Yang et al., 2020), feature attribution methods are used to characterize the input and intermediate layer representations of a DNN. They find that adversarial inputs drastically alter the feature attribution map compared to natural inputs. Statistical dispersion measures such as the interquartile range (IQR) and median absolute deviation (MAD) are used to quantify the dispersion in the distribution of attribution values, which are then used as test statistics (features) to train a logistic classifier for discriminating adversarial samples from natural samples. In the context of the meta-algorithm, this method does not calculate class-conditional test statistics. The test statistics based on IQR or MAD are not normalized in any way. Similar to (Lee et al., 2018) and (Ma et al., 2018), this method uses a weighted linear combination of the test statistics for scoring.

Unsupervised Methods

(Roth et al., 2019) show that the log-odds ratio of inputs to a classifier (not necessarily a DNN) can reveal some interesting properties of adversarial inputs. They propose test statistics based on the expected log-odds of noise-perturbed inputs from different source and predicted class pairs. These test statistics are z-score normalized and thresholded to detect adversarial inputs. Their method does not use multiple layer representations for detection. They also propose a reclassification method for correcting the classifier's prediction on adversarial inputs. This is based on training logistic classifiers (one for each predicted class) that use the expected noise-perturbed log-odds ratio as features.

In (Zheng & Hong, 2018), the class-conditional distributions of fully-connected intermediate layer representations are modeled using Gaussian mixture models [12]. Inputs with a likelihood lower than a (class-specific) threshold under each class-conditional mixture model are rejected as adversarial. A key limitation of this method is that it is challenging to accurately model the high-dimensional layer representations of a DNN using density models such as Gaussian mixtures.

In (Miller et al., 2019), an anomaly detection method focusing on adversarial attacks is proposed, which in its basic form can be described as follows. The class-conditional density of the layer representations of a DNN are modeled using Gaussian mixture models, and they are used to estimate a Bayes class posterior at each layer. The Kullback-Leibler divergence between the class posterior of the DNN (based on the softmax function) and the estimated Bayes class posterior from each layer are used as test statistics for detecting adversarial samples. A number of methods are proposed for combining these class-conditional test statistics from the layers (*e.g.*, maximum and weighted sum across the layers). Their method does not apply any explicit normalization of the test statistics from the layers.

(Sastry & Oore, 2020) propose a method for detecting OOD samples based on analyzing the Gram matrices of the layer representations of a DNN. The Gram matrices of different orders (order 1 corresponds to the standard definition) capture the pairwise correlations between the elements of a layer representation. At training time, their detector records the element-wise minimum and maximum values of the Gram matrices (of different orders) from a training set of natural inputs, conditioned on each predicted class. The extent of deviation from the minimum and maximum values observed at training time is used to calculate a class-specific deviation test statistic at each layer. The final score for OOD detection is obtained by adding up the normalized deviations from the layers, where the normalization factor for a layer is the expected deviation observed on a held-out validation dataset. We note that the deviation statistic based on Gram matrices from the layer activations proposed in (Sastry & Oore, 2020) can be used as a test statistic for `JTLA` as well.

Confidence Metrics for Classifiers

Works such as the trust score (Jiang et al., 2018) and deep KNN (Papernot & McDaniel, 2018) have explored the problem of developing a confidence metric that can be used to independently validate the predictions of a classifier. Inputs with a low confidence score are expected to be misclassified and can be flagged as potentially adversarial or OOD. Deep KNN (Papernot & McDaniel, 2018) uses the class labels of the k-nearest neighbors of DNN layer representations to calculate a non-conformity score corresponding to each class. Large values of non-conformity corresponding to the predicted class indicate that an input may not have a reliable prediction. The method calculates empirical p-values of the non-conformity scores to provide a confidence score, credibility score, and an alternate (corrected) class prediction for test inputs. We note that the deep KNN test statistic based on the kNN class counts from the predicted class $k_{\hat{c}}^{(\ell)}$ and its complement $k - k_{\hat{c}}^{(\ell)}$ can be considered as a binomial specialization of the multinomial test statistic (§ 4.1). The trust score (Jiang et al., 2018) estimates the α-high density (level) set for each class, and calculates the distances from a test point to the α-high density sets from the classes to define a confidence metric. These methods can also be categorized as unsupervised.

[12] Convolutional layers are not modeled in their approach.

A.5. Scoring with a Logistic classifier

Consider a binary logistic classifier that takes a vector of test statistics $\mathbf{t}$ as input and produces an output probability for class 1 given by

$$P(Y = 1 \,|\, \mathbf{t}) \;=\; \frac{1}{1 \,+\, \exp(-\mathbf{w}^T\,\mathbf{t} \,-\, b)},$$

where $\mathbf{w}$ and b are weight vector and bias parameters. A detection decision of 1 (anomaly) is made when the output probability exceeds a threshold τ. It is easy to see that this results in the following decision rule:

$$\psi(\mathbf{x}^{(0)}, \cdots, \mathbf{x}^{(L)}, \hat{c}) \;=\; \begin{cases} 1 & \text{if } \mathbf{w}^T\,\mathbf{t} \,\geq\, \log\frac{\tau}{1-\tau} \,-\, b \\ 0 & \text{otherwise} \end{cases}$$

In other words, the score function is a weighted linear combination of the test statistics from the layers.

B. Harmonic Mean p-value Method

The harmonic mean p-value (HMP) (Wilson, 2019) is a recently proposed method for combining p-values from a large number of dependent tests. It is rooted in ideas from Bayesian model averaging and has some desirable properties such as robustness to positive dependency between the p-values, and an ability to detect small statistically-significant groups from a large number of p-values (tests). The main result of (Wilson, 2019) can be summarized as follows. Given a set of p-values $p_1, \cdots, p_m$ from m hypothesis tests, the weighted harmonic p-value of any subset $\mathcal{R} \subset \{1, \cdots, m\}$ of the p-values is given by

$$p_{\text{agg}}^{-1} \;=\; \frac{\sum\limits_{i \in \mathcal{R}} w_i p_i^{-1}}{\sum\limits_{i \in \mathcal{R}} w_i},$$

where the weights w_i are non-negative and satisfy $\sum_{i=1}^m w_i \,=\, 1$. In our problem, we apply the HMP method with the weights all set to the same value, resulting in the p-value aggregation function $r(\cdot)$ defined in Eq. (11).

In our experiments, we found the HMP method to have comparable or slightly worse performance than Fisher's method of combining p-values. Results from these ablation experiments can be found in Appendix E.2.

C. Details and Extensions of the Adaptive Attack

We first discuss the method we used for setting the scale of the Gaussian kernel per layer in the adaptive attack method of § 5. We then discuss an untargeted variant of the proposed adaptive attack, followed by an alternate formulation for the attack objective function.

C.1. Setting the Kernel Scale

For a given clean input $\mathbf{x}$, the scale of the Gaussian kernel for each layer σ_ℓ determines the effective number of samples that contribute to the soft count that approximates the kNN counts in Eq. (14). Intuitively, we would like the kernel to have a value close to 1 for points within a distance of η_ℓ (the kNN radius centered on $\mathbf{f}_\ell(\mathbf{x})$), and decay rapidly to 0 for points further away. Let $\mathcal{N}_k(\mathbf{f}_\ell(\mathbf{x}))$ denote the index set of the k-nearest neighbors of $\mathbf{f}_\ell(\mathbf{x})$ from the ℓ-th layer representations of the dataset $\mathcal{D}_a$. The probability of selecting the k-nearest neighbors from the set of N samples in $\mathcal{D}_a$ can be expressed as

$$s_1(\sigma_\ell) \;=\; \frac{\sum\limits_{n \,\in\, \mathcal{N}_k(\mathbf{f}_\ell(\mathbf{x}))} h_{\sigma_\ell}(\mathbf{f}_\ell(\mathbf{x}), \mathbf{f}_\ell(\mathbf{x}_n))}{\sum\limits_{n=1}^{N} h_{\sigma_\ell}(\mathbf{f}_\ell(\mathbf{x}), \mathbf{f}_\ell(\mathbf{x}_n))}.$$

We could choose σ_ℓ to maximize this probability and push it close to 1. However, this is likely to result in a very small value for σ_ℓ, which would concentrate all the probability mass on the nearest neighbor. To ensure that the probability mass is distributed sufficiently uniformly over the k-nearest neighbors we add the following normalized entropy [13] term as a regularizer

$$s_2(\sigma_\ell) \;=\; -\,\frac{1}{\log k} \sum_{i \,\in\, \mathcal{N}_k(\mathbf{f}_\ell(\mathbf{x}))} p_i(\sigma_\ell)\, \log p_i(\sigma_\ell),$$

where

$$p_i(\sigma_\ell) \;=\; \frac{h_{\sigma_\ell}(\mathbf{f}_\ell(\mathbf{x}), \mathbf{f}_\ell(\mathbf{x}_i))}{\sum\limits_{j \,\in\, \mathcal{N}_k(\mathbf{f}_\ell(\mathbf{x}))} h_{\sigma_\ell}(\mathbf{f}_\ell(\mathbf{x}), \mathbf{f}_\ell(\mathbf{x}_j))}.$$

Both terms $s_1(\sigma_\ell)$ and $s_2(\sigma_\ell)$ are bounded to the interval $[0, 1]$. We find a suitable σ_ℓ by maximizing a convex combination of the two terms, *i.e.*,

$$\max_{\sigma_\ell \,\in\, \mathbb{R}_+} \;(1 - \alpha)\, s_1(\sigma_\ell) \;+\; \alpha\, s_2(\sigma_\ell). \tag{15}$$

We set α to 0.5 in our experiments, and used a simple line search to find the approximate maximizer of Eq. (15).

C.2. Untargeted Attack Formulation

The formulation in § 5, where a specific class $c' \neq c$ is chosen, is used to create a targeted attack. Alternatively, a simple modification to Eq. (13) that considers the log-odds of class c, $\log\frac{p_c}{1-p_c}$, can be used to create an untargeted attack, resulting in the following objective function to be minimized:

[13] The entropy is divided by the maximum possible value of $\log k$ to scale it to the range $[0, 1]$.

$$J(\boldsymbol{\delta}) \;=\; \|\boldsymbol{\delta}\|_2^2 \;+\; \lambda \log \sum_{\ell=0}^{L} \sum_{\substack{n=1: \\ c_n=c}}^{N} h_{\sigma_\ell}(\mathbf{f}_\ell(\mathbf{x}+\boldsymbol{\delta}), \mathbf{f}_\ell(\mathbf{x}_n)) \\ -\; \lambda \log \sum_{\ell=0}^{L} \sum_{\substack{n=1: \\ c_n\neq c}}^{N} h_{\sigma_\ell}(\mathbf{f}_\ell(\mathbf{x}+\boldsymbol{\delta}), \mathbf{f}_\ell(\mathbf{x}_n)). \qquad (16)$$

C.3. Alternate Attack Loss Function

Starting from equation Eq. (13), but now assuming that the probability estimate for each class based on the kNN model factors into a product of probabilities across the layers of the DNN (*i.e.*, independence assumption), we get

$$\log \frac{p_c}{p_{c'}} \;=\; \log \frac{k_c}{k_{c'}} \;=\; \log \frac{\prod_\ell k_c^{(\ell)} \,/\, k}{\prod_\ell k_{c'}^{(\ell)} \,/\, k} \\ =\; \sum_{\ell=0}^{L} \log k_c^{(\ell)} \;-\; \sum_{\ell=0}^{L} \log k_{c'}^{(\ell)}.$$

Using the soft count approximation based on the Gaussian kernel (as before) leads to the following alternative loss function for the targeted adaptive attack

$$J(\boldsymbol{\delta}) \;=\; \|\boldsymbol{\delta}\|_2^2 \;+\; \lambda \sum_{\ell=0}^{L} \log \sum_{\substack{n=1: \\ c_n=c}}^{N} h_{\sigma_\ell}(\mathbf{f}_\ell(\mathbf{x}+\boldsymbol{\delta}), \mathbf{f}_\ell(\mathbf{x}_n)) \\ -\; \lambda \sum_{\ell=0}^{L} \log \sum_{\substack{n=1: \\ c_n=c'}}^{N} h_{\sigma_\ell}(\mathbf{f}_\ell(\mathbf{x}+\boldsymbol{\delta}), \mathbf{f}_\ell(\mathbf{x}_n)). \qquad (17)$$

In contrast to Eq. (14), this loss function considers the class probability estimates from each layer, instead of a single class probability estimate based on the cumulative kNN counts across the layers. A special case of the loss function Eq. (17) that includes only the final (logit) layer of the DNN $\mathbf{f}_L(\mathbf{x})$ can be directly compared to the Carlini-Wagner attack formulation (Carlini & Wagner, 2017b). With this formulation, the logit layer representations of an adversarial input will be guided closer to the neighboring representations from class c', and away from the neighboring representations from class c.

D. Additional Implementation Details

D.1. Computing Platform

Our experiments were performed on a single server running Ubuntu 18.04 with 128 GB memory, 4 NVIDIA GeForce RTX 2080 GPUs, and 32 CPU cores.

D.2. Datasets & DNN Architectures

A summary of the image classification datasets we used, the architecture and test set performance of the corresponding DNNs are given in Table 2. Each dataset has 10 image classes. We followed recommended best practices for training DNNs on image classification problems (using techniques like Dropout). We did not train a DNN on the NotMNIST dataset because this dataset is used only for evaluation in the OOD detection experiments.

D.3. Method Implementations

The code associated with our work can be accessed here [14]. We utilized the authors original implementation for the following methods: (i) deep mahalanobis detector [15], (ii) the odds are odd detector [16]. We implemented the remaining detection methods, viz. `LID`, `DKNN`, and `Trust`, and this is available as part of our released code. All implementations are in `Python3` and are based on standard scientific computing libraries such as `Numpy`, `Scipy`, and `Scikit-learn` (Pedregosa et al., 2011). We used `PyTorch` as the deep learning and automatic differentiation backend (Paszke et al., 2017). We perform approximate nearest neighbor search using the `NNDescent` method (Dong et al., 2011) to efficiently construct and query from kNN graphs at the DNN layer representations. We used the implementation of NNDescent provided by the library `PyNNDescent` [17]. Our implementation of `JTLA` is highly modular, allowing for new test statistics to be easily plugged in to the existing implementation. We provide implementations of the following test statistics: (i) multinomial class counts, (ii) binomial class counts, (iii) trust score, (iv) local intrinsic dimensionality, and (v) average kNN distance.

In our experiments with `JTLA`, where Fisher's method or HMP method are used for combining the p-values, we included p-values estimated from the individual layers (listed in Tables 3, 4, and 5) and from all distinct layer pairs. The number of nearest neighbors k is the **only hyperparameter** of the proposed method. This is set to be a function of the number of in-distribution training samples n using the heuristic $k = \lceil n^{0.4} \rceil$. For methods like `DKNN`, `LID`, and `Trust` that also depend on k, we found that setting k in this way produces comparable results to that obtained over a range of k values. Therefore, to maintain consistency, we set k for all the (applicable) methods using the rule $k = \lceil n^{0.4} \rceil$.

For the method `Trust`, we applied the trust score to the input, logit (pre-softmax) layer, and the fully-connected layer prior to the logit (pre-logit) layer. Since it was reported in (Jiang et al., 2018) that the trust score works well mainly in low-dimensional settings, we applied the same dimension-

[14] https://github.com/jayaram-r/adversarial-detection

[15] https://github.com/pokaxpoka/deep_Mahalanobis_detector

[16] https://github.com/yk/icml19_public

[17] https://github.com/lmcinnes/pynndescent

Table 2: Datasets & DNN Architectures.

Dataset	Input dimension	Train size	Test size	Test accuracy (%)	DNN architecture
MNIST (LeCun et al., 1998)	$28 \times 28 \times 1$	50,000	10,000	99.12	2 Conv. + 2 FC layers (LeCun et al., 1998)
SVHN (Netzer et al., 2011)	$32 \times 32 \times 3$	73,257	26,032	89.42	2 Conv. + 3 FC layers
CIFAR-10 (Krizhevsky et al., 2009)	$32 \times 32 \times 3$	50,000	10,000	95.45	ResNet-34 (He et al., 2016)
Not-MNIST (Bulatov, 2011)	$28 \times 28 \times 1$	500,000	18,724	N/A	N/A

ality reduction that was applied to JTLA (see Tables 3, 4, and 5) on the input and pre-logit layer representations. We found the pre-logit layer to produce the best detection performance in our experiments. Hence, we report results for Trust with the pre-logit layer in all our experiments. The constant α, which determines the fraction of samples with lowest empirical density to be excluded from the density level sets, is set to 0 in our experiments. We explored a few other values of α, but did not find a significant difference in performance. This is consistent with the observation in Section 5.3 of (Jiang et al., 2018).

The methods Mahalanobis and LID train a logistic classifier to discriminate adversarial samples from natural samples. The regularization constant of the logistic classifier is found by searching (over a set of 10 logarithmically spaced values between 0.0001 to 10000) for the value that leads to the largest average test-fold area under the ROC curve, in a 5-fold stratified cross-validation setup.

For the method Odds, the implementation of the authors returns a binary (0 / 1) detection decision instead of a continuous score that can be used to rank adversarial samples. The binary decision is based on applying z-score normalization to the original score, and comparing it to the 99.9-th percentile of the standard Gaussian density. Instead of using the thresholded decision, we used the z-score-normalized score of Odds in order to get a continuous score that is required by the metrics average precision and pAUC-α.

Details on the DNN Layer Representations

The DNN layers used in our experiments and their raw dimensionality are listed for the three datasets in Tables 3, 4, and 5. An exception to this is the LID method, for which we follow the implementation of (Ma et al., 2018) and calculate the LID features from all the intermediate layers. For DKNN and LID, we did not apply any dimensionality reduction or pre-processing of the layer representations to be consistent with the respective papers. For Mahalanobis, the implementation of the authors performs global average pooling at each convolutional layer to transform a $C \times W \times H$ tensor (with C channels) to a C-dimensional vector.

For JTLA, we use the neighborhood preserving projection (NPP) method (He et al., 2005) to perform dimensionality reduction on the layer representations. We chose NPP because it performs an efficient linear projection that attempts to preserve the neighborhood structure in the original space as closely as possible in the projected (lower dimensional) space. Working with the lower dimensional layer representations mitigates problems associated with the curse of dimensionality, and significantly reduces the memory utilization and the running time of JTLA. The original dimension, intrinsic dimension estimate, and the projected dimension of the layer representations for the three datasets are listed in Tables 3, 4, and 5. We did not apply dimension reduction to the logit layer because it has only 10 dimensions.

Table 3: Layer representations of the DNN trained on MNIST. The output of the block listed in the first column is used as the layer representation.

Layer block	Layer index	Original dimension	Intrinsic dimension	Projected dimension
Input	0	784	13	31
Conv-1 + ReLu	1	21632	22	53
Conv-2 + Maxpool + Dropout	2	9216	18	77
FC-1 + ReLu + Dropout	3	128	9	90
FC-2 (Logit)	4	10	6	10

Table 4: Layer representations of the DNN trained on SVHN. The output of the block listed in the first column is used as the layer representation.

Layer block	Layer index	Original dimension	Intrinsic dimension	Projected dimension
Input	0	3072	18	43
Conv-1 + ReLu	1	57600	38	380
Conv-2 + ReLu + Maxpool + Dropout	2	12544	42	400
FC-1 + ReLu + Dropout	3	512	12	120
FC-2 + ReLu + Dropout	4	128	7	10
FC-3 (Logit)	5	10	4	10

The procedure we used for determining the projected dimension is summarized as follows. We used the training partition of each dataset (that was used to train the DNN) for this task in order to avoid introducing any bias on the test partition (which is used for the detection and corrected classification experiments). At each layer, we first estimate the intrinsic dimension (ID) as the median of the LID estimates of the samples, found using the method of (Amsaleg et al., 2015). The ID estimate d_{int} is used as a lower bound

for the projected dimension. Using a 5-fold stratified cross-validation setup, we search over 20 linearly spaced values in the interval $[d_{\text{int}}, 10\, d_{\text{int}}]$ for the projected dimension (found using NPP) that results in the smallest average test-fold error rate for a standard k-nearest neighbors classifier. The resulting projected dimensions for each dataset (DNN architecture) are given in Tables 3, 4, and 5.

Table 5: Layer representations of the ResNet-34 trained on CIFAR-10. The output of the block listed in the first column is used as the layer representation. Legend: RB - Residual block, BN - BatchNorm

Layer block	**Layer index**	**Original dimension**	**Intrinsic dimension**	**Projected dimension**
Input	0	3072	25	48
Conv-1 + BN + ReLU	1	65536	33	330
RB-1	2	65536	58	580
RB-2	3	32768	59	590
RB-3	4	16384	28	28
RB-4	5	8192	15	15
2D Avg. Pooling	6	512	9	9
FC (Logit)	7	10	8	10

Note on Performance Calculation

The following procedure is used to calculate performance metrics as a function of the perturbation norm of adversarial samples. Suppose there are N_a adversarial samples and N_n natural samples in a test set, with $N = N_a + N_n$. Define the maximum proportion of adversarial samples $p_a = N_a \,/\, N$. The adversarial samples are first sorted in increasing order of their perturbation norm. The proportion of adversarial samples is varied over 12 equally-spaced values between 0.005 and $\min(0.3, p_a)$. For a given proportion p_i, the top $N_i = \lceil p_i\, N \rceil$ adversarial samples (sorted by norm) are selected. The perturbation norm of all these adversarial samples will be below a certain value; this value is shown on the x-axis of the performance plots. The performance metrics (average precision, pAUC-α etc.) are then calculated from the N_i adversarial samples and the N_n natural samples.

In order to calculate the performance metrics as a function of the proportion of adversarial or OOD (anomalous) samples, the only difference is that we do not sort the anomalous samples in a deterministic way. For a given proportion p_i, we select $N_i = \lceil p_i\, N \rceil$ anomalous samples at random uniformly, and calculate the performance metrics based on the N_i anomalous and N_n natural samples. To account for variability, this is repeated over 100 random subsets of N_i anomalous samples each, and the median value of the performance metrics is reported. Note that the detection methods need to score the samples only once, and the above calculations can be done based on the saved scores.

D.4. Details on the Adversarial Attacks

We list below the parameters of the adversarial attack methods we implemented using Foolbox (Rauber et al., 2017). We utilize the same variable names used by Foolbox in order to enable easy replication.

- PGD attack (Madry et al., 2018): ℓ_∞ norm with the norm-ball radius ϵ linearly spaced in the interval $[\frac{1}{255}, \frac{21}{255}]$. Using the notation of the Foolbox library: `epsilon = [`$\frac{1}{255}, \frac{3}{255}, \cdots, \frac{21}{255}$`], stepsize = 0.05, binary_search = False, iterations = 40, p_norm = inf.`
- Carlini-Wagner (CW) attack (Carlini & Wagner, 2017b): ℓ_2 norm with the confidence parameter of the attack varied over the set $\{0, 6, 14, 22\}$. Using the notation of the Foolbox library: `confidence = [0,6,14,22], max_iterations = 750, p_norm = 2.`
- FGSM attack (Goodfellow et al., 2015): ℓ_∞ norm with maximum norm-ball radius $\epsilon_{\max} = 1$. Using the notation of the Foolbox library: `max_epsilon = 1, p_norm = inf.`

For the adversarial detection experiments in § 6.1, we reported results for the attack parameter setting that would produce adversarial samples with the lowest perturbation norm (least perceptible), in order to make the detection task challenging. This corresponds to $\epsilon = \frac{1}{255}$ for the PGD attack, confidence $= 0$ for the CW attack, and $\epsilon_{\max} = 1$ for the FGSM attack. These same parameters are also used in the experiments in Appendix F.2, F.3, F.4, F.6, and F.7. In Appendix F.1, the strength of the attack is varied and the attack parameters used are described there. The adversarial samples are generated *once* from the clean samples corresponding to each train fold and test fold (from cross-validation), and saved to files for reuse in all the experiments. This way, we ensure that there is no randomness in the experiments arising due to the adversarial sample generation.

Adaptive Attack:

Here we provide additional details on the adaptive (defense-aware) attack method proposed in § 5. The constant λ in the objective function controls the perturbation norm of the adversarial sample, and smaller values of λ lead to solutions with a smaller perturbation norm. We follow the approach of (Carlini & Wagner, 2017b) to set λ to the smallest possible value that leads to a successful adversarial perturbation. This is found using a bisection line search over ten steps. An adversarial input is considered successful if it modifies the initially-correct class prediction of the defense method. We also follow the approach of (Carlini & Wagner, 2017b) to implicitly constraint the adversarial inputs to lie

in the same range as the original inputs using the hyperbolic tangent and its inverse function. Suppose the inputs lie in the range $[a, b]^d$, the following sequence of transformations is applied to each component of the vectors

$$\begin{aligned} z_i &= \tanh^{-1}(2\,\frac{x_i - a}{b - a} - 1) \\ z'_i &= z_i + w_i \\ x'_i &= (b - a)\,\frac{1}{2}\,(1 + \tanh(z_i + w_i)) + a \end{aligned}$$

effectively allowing the perturbation $\mathbf{w}$ to be optimized over $\mathbb{R}^d$. The resulting unconstrained optimization is solved using the RMSProp variant of stochastic gradient descent with adaptive learning rate (Ruder, 2016). The maximum number of gradient descent steps is set to 1000. For all experiments based on the adaptive attack method, the attack targets the variant of `JTLA` based on p-value normalization at the individual layers and layer pairs, Fisher's method for p-value aggregation, and the multinomial test statistic.

E. Additional Experiments

In this section, we supplement the results in § 6 with more extensive experiments.

E.1. Attack Transfer and Attacks of Varying Strength

We evaluate the performance of the detection methods on a task with different adversarial attacks used in the train and test sets. The strength of the attacks are also varied in both the train and test sets. Recall that we use a 5-fold cross-validation framework for evaluation. Hence, for a train/test split corresponding to fold $i \in \{1, 2, 3, 4, 5\}$, adversarial samples from Attack A are generated in the train split, while adversarial samples from Attack B are generated in the test split. Supervised methods, `Mahalanobis` and `LID`, learn using both the adversarial and clean samples from the train split, while the unsupervised methods, `JTLA`, `Odds`, `DKNN`, and `Trust`, learn using only the clean samples from the train split.

ℓ_2-CW attack to ℓ_∞-PGD attack

The ℓ_∞-PGD (Madry et al., 2018) attack is applied on the test split with perturbation strength parameter ϵ varied over the values $\{1/255, 3/255, 5/255, 7/255, 9/255\}$. For each clean sample, one of these ϵ values is randomly selected and used to create an attack sample. The ℓ_2-CW attack (Carlini & Wagner, 2017b) is applied to the train split with the confidence parameter value randomly selected from the set $\{0, 6, 14, 22\}$. The result of this experiment is given in Figure 6, with the average precision of the compared methods plotted as a function of the perturbation norm on the x-axis. The ℓ_∞ norm is used on the x-axis to match the norm type used by the PGD attack (that is applied to the test split). We observe that `Mahalanobis` and both variants of `JTLA` have the best performance on CIFAR-10, while both variants of `JTLA` outperform the other methods on SVHN. On MNIST, `JTLA` based on the aK-LPE method (§ 4.2.B) and `Trust` have the best performance. The methods `Mahalanobis`, `Trust`, and `Odds` do not consistently have good performance, while the `LID` method has poor performance on all datasets.

ℓ_∞-PGD attack to ℓ_2-CW attack

In this experiment, the ℓ_2-CW attack is applied to the test split and the ℓ_∞-PGD attack is applied to the train split. The attack parameters (strength) are varied as described earlier. The result of this experiment is given in Figure 7, with the average precision of the methods plotted as a function of the ℓ_2 norm of the attack (since the ℓ_2-CW attack is applied to the test split). From the figure, we make the following observations. On CIFAR-10, `Odds` outperforms the other methods, followed by the two variants of `JTLA`. On SVHN, the two variants of `JTLA` have the best performance, followed by `Odds`. On MNIST, the methods `JTLA`, `Odds`, and `Trust` have very similar average precision that is higher than the other methods. The methods `Mahalanobis`, `DKNN`, and `LID` do not have good performance in this experiment. We hypothesize that the `Odds` method (Roth et al., 2019) works well in detecting CW attacks because it uses a test statistic based on the noise-perturbed log-odds ratio of all pairs of classes, which is well-matched to the CW attack that is based on skewing the log-odds ratio of the class pair involved in the attack.

E.2. Ablation Experiments

We discuss the ablation experiments that we performed to get a better understanding of the components of the proposed method. Specifically, we are interested in understanding

1. The relative performance of the proposed p-value based normalization methods (§ 4.2) and the p-value aggregation methods (§ 4.3).
2. The value of including p-values from test statistics at all layer pairs (in addition to the individual layers).
3. The relative performance of using only the last few layer representations for detection.
4. The relative performance of the two scoring methods in § 4.4 on the task of adversarial detection.

Tables 6, 7, and 8 summarize the results of these experiments on the task of detecting adversarial samples from the ℓ_∞-PGD attack with $\epsilon = 1/255$, and the ℓ_2-CW attack with confidence set to 0. DNNs trained on the SVHN and CIFAR-10 datasets (described earlier) are used, and average precision and pAUC-0.2 (partial AUROC below FPR $= 0.2$) are reported as detection metrics. Unlike the results in Fig. 2 and Fig. 3, we do not vary the proportion of adversarial samples, and report performance with all the adversarial

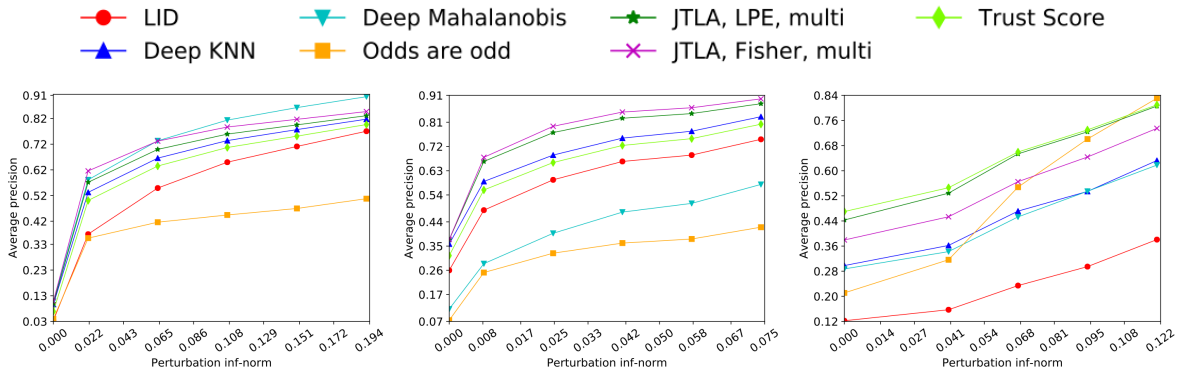

(a) CIFAR-10 (b) SVHN (c) MNIST

Figure 6: Average precision on the attack transfer experiment: CW to PGD attack.

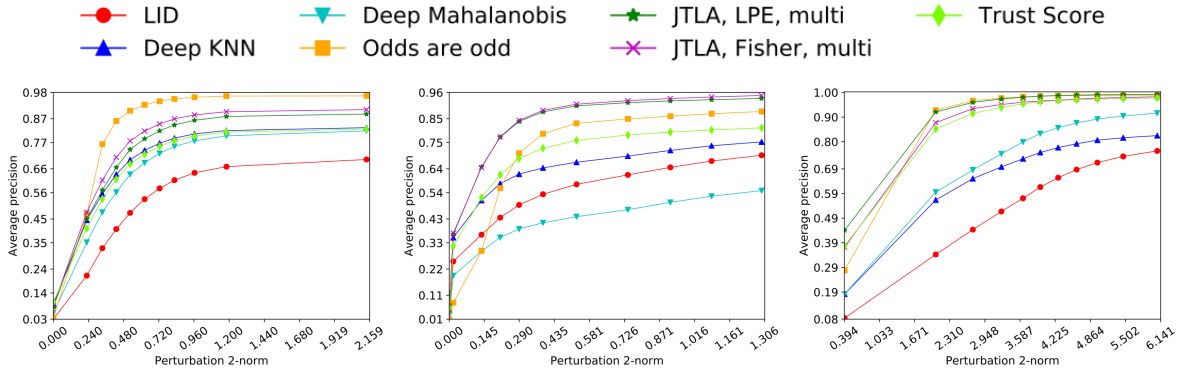

(a) CIFAR-10 (b) SVHN (c) MNIST

Figure 7: Average precision on the attack transfer experiment: PGD to CW attack.

samples included.

Table 6 focuses on points 1 and 2, and compares the proposed p-value normalization and aggregation methods. P-values from the layer pairs are included in the first case and not included in the second case. The best performing configuration (across the rows) is highlighted in bold. We observe that p-value normalization at the layers with aggregation using Fisher's method has the best performance in most cases. Including the layer pairs did not result in a significant improvement in these experiments. It is surprising that Fisher's method has better (in some cases comparable) performance compared to the HMP and the aK-LPE methods despite its simplistic assumption of independent tests (p-values). We believe this can be attributed to the fact that the multivariate p-values estimated by the aK-LPE method (Eq. (9)) require a large sample size to converge to their true values. Since we apply this estimator class conditionally, the moderate number of samples per class (ranging from 500 to 5000 in our experiments) may result in estimation errors. Also, we conjecture that Fisher's method achieves a higher detection rate (TPR) at the expense of a higher FPR, while the HMP method has a more conservative TPR with a lower FPR.

Table 7 focuses on point 3 and compares the performance of using all the layer representations (listed in tables 4 and 5) with the performance from using only the final one, two, or three layer representations [18]. We observe that including more layers generally increases the detection performance, confirming the intuition behind using multiple layers. For the CW attack on CIFAR-10, using only the final (logit) layer has comparable performance to using all the layers. This is consistent with the design of the CW attack based on only the logit layer representation.

Table 8 focuses on point 4 to understand if the score function (12) is better suited for adversarial samples since it considers the aggregate p-values from the candidate true classes. From the table, we observe that the adversarial score function clearly outperforms the OOD score function for each combination of normalization and aggregation method.

[18] We focus on the deeper layers since their representations are most useful for classification.

Table 6: Ablation experiment - performance of different p-value normalization and aggregation methods, and the effect of including/excluding layer pairs. The multinomial test statistic is used in all cases.

Normalization method	Aggregation method	SVHN, PGD ($\epsilon = 1/255$)		SVHN, CW (confidence = 0)		CIFAR-10, PGD ($\epsilon = 1/255$)		CIFAR-10, CW (confidence = 0)	
		average precision	pAUC-0.2	average precision	pAUC-0.2	average precision	pAUC-0.2	average precision	pAUC-0.2
p-values from layers & layer pairs (§4.2.A)	Fisher	0.7382	**0.8025**	0.9631	0.9213	0.7710	**0.7790**	0.9664	0.9377
	HMP	0.7296	0.7966	0.9611	0.9171	0.7614	0.7705	0.9653	0.9344
p-values from layers	Fisher	**0.7393**	0.8005	**0.9634**	**0.9214**	**0.7734**	0.7781	**0.9667**	**0.9380**
	HMP	0.7247	0.7925	0.9591	0.9140	0.7538	0.7617	0.9616	0.9315
Multivariate p-value (aK-LPE, §4.2.B)	None	0.7161	0.7840	0.9559	0.9042	0.7437	0.7518	0.9650	0.9296

Table 7: Ablation experiment: effect of including only the last few layers for detection.

Normalization method	Aggregation method	Layers included	SVHN, PGD ($\epsilon = 1/255$)		SVHN, CW (confidence = 0)		CIFAR-10, PGD ($\epsilon = 1/255$)		CIFAR-10, CW (confidence = 0)	
			average precision	pAUC-0.2	average precision	pAUC-0.2	average precision	pAUC-0.2	average precision	pAUC-0.2
p-values from layers & layer pairs (§4.2.A)	Fisher	All	**0.7382**	**0.8025**	**0.9631**	**0.9213**	**0.7710**	**0.7790**	0.9664	0.9377
		Final (logits)	0.6347	0.7396	0.9141	0.8538	0.7399	0.7636	0.9664	0.9371
		Last 2	0.6440	0.7431	0.9213	0.8599	0.7410	0.7650	**0.9688**	**0.9401**
		Last 3	0.6847	0.7674	0.9388	0.8857	0.7473	0.7657	0.9660	0.9358

E.3. Results on the FGSM Attack

Figure 8 presents the average precision of the different detection methods as a function of the ℓ_2 norm of perturbation for the FGSM attack method. It is clear that both variants of JTLA outperform the other methods, consistent with the trend observed on other attacks in § 6.

E.4. Evaluation of Partial Area Under the ROC Curve

Here we compare the performance of methods using pAUC-0.2, a metric calculating the partial area under the ROC curve for FPR below 0.2. Comparing the area under the entire ROC curve can lead to misleading interpretations because it includes FPR values that one would rarely choose to operate in. Therefore, to reflect realistic operating conditions, we chose a maximum FPR of 0.2. Recall that for the PGD attack we vary the proportion of adversarial samples along the x-axis because most of the samples from this attack have the same perturbation norm.

On the CIFAR-10 dataset (Figure 9), we observe that JTLA has better performance than the other methods in most cases. On the adaptive attack, Mahalanobis has slightly better performance than JTLA with the multivariate p-value estimation method (aK-LPE). We make similar observations on the SVHN dataset (Figure 10), with a minor exception on the adaptive attack where the Odds method performs better than JTLA as the perturbation norm increases.

On the MNIST dataset (Figure 11), we observe some different trends in the performance compared to the other datasets. On the CW and FGSM attacks, the methods Odds and Trust perform comparably or slightly better than JTLA (particularly the variant based on Fisher's method). On the adaptive attack, the performance of JTLA based on Fisher's method decreases significantly as the perturbation norm increases. On the other hand, the variant of JTLA based on the aK-LPE method outperforms the other methods on this attack. We think that this contrast in performance is due to the fact that the adaptive attack samples were optimized to fool the variant of JTLA based on Fisher's method. Also, attack samples with higher perturbation norm are more likely to be successful. On the PGD attack, Odds outperforms the other methods, but the pAUC-0.2 of all methods, except LID and DKNN, are higher than 0.95 in this case. We conjecture that the good performance of most methods on MNIST could be due to the simplicity of the input space and the classification problem.

E.5. Results on the MNIST Dataset

In Figure 12, we compare the average precision of different methods on the MNIST dataset for the CW, PGD, and adaptive attacks (results for the FGSM attack were presented in Appendix E.3). We observe that Odds has good performance on this dataset, outperforming JTLA in some cases. The method Trust (which uses the pre-logit, fully connected layer) also performs well on this dataset. This could be due to the fact that on the MNIST dataset, the attack samples exhibit very distinctive patterns at the logit and pre-logit DNN layers, which are the focus of the methods Odds and Trust. We note that Odds and Trust do not carry over this good performance to all datasets and attacks. Also, both variants of JTLA perform well in the low perturbation norm regime.

Table 8: Ablation experiment: comparison of the adversarial and OOD score functions (§4.4) on different adversarial attacks. The adversarial score function outperforms the OOD score function in all cases.

Normalization method	Aggregation method	Scoring method	SVHN, PGD ($\epsilon = 1/255$)		SVHN, CW (confidence = 0)		CIFAR-10, PGD ($\epsilon = 1/255$)		CIFAR-10, CW (confidence = 0)	
			average precision	pAUC-0.2	average precision	pAUC-0.2	average precision	pAUC-0.2	average precision	pAUC-0.2
p-values from layers & layer pairs (§4.2.A)	Fisher	Adversarial	**0.7382**	**0.8025**	**0.9631**	**0.9213**	**0.7710**	**0.7790**	**0.9664**	**0.9377**
		OOD	0.7078	0.7736	0.9524	0.8973	0.7492	0.7612	0.9620	0.9253
	HMP	Adversarial	**0.7296**	**0.7966**	**0.9611**	**0.9171**	**0.7614**	**0.7705**	**0.9653**	**0.9344**
		OOD	0.6946	0.7617	0.9482	0.8882	0.7396	0.7501	0.9599	0.9205
Multivariate p-value (aK-LPE, §4.2.B)	None	Adversarial	**0.7161**	**0.7840**	**0.9559**	**0.9042**	**0.7437**	**0.7518**	**0.9650**	**0.9296**
		OOD	0.6986	0.7664	0.9511	0.8941	0.7065	0.7247	0.9575	0.9149

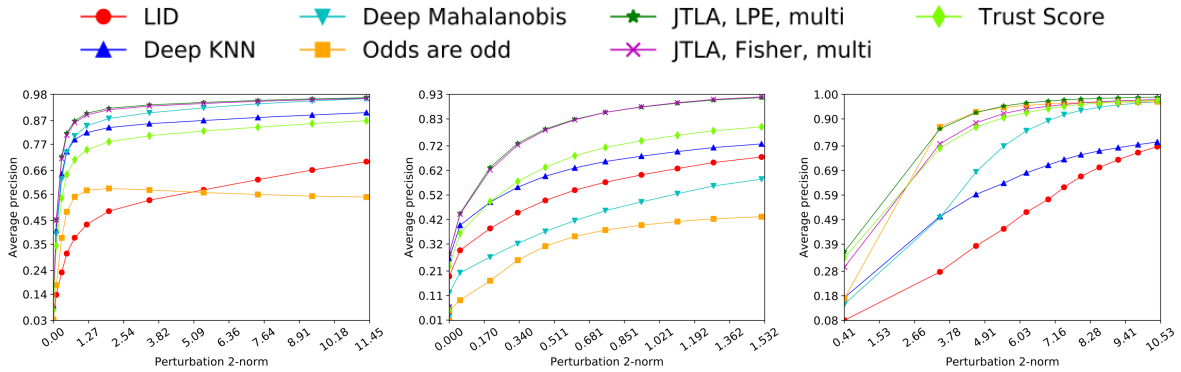

(a) CIFAR-10 (b) SVHN (c) MNIST

Figure 8: Average precision on the FGSM attack ($\epsilon_{\max} = 1$) for all datasets.

E.6. Running Time Comparison

Table 9 reports the wall-clock running time (in minutes) of the different detection methods per-fold, averaged across all attack methods. `Trust` consistently has the least running time, while both variants of `JTLA` have low running time as well. The running time of `Mahalanobis` is comparable to `JTLA` on MNIST and SVHN, but is higher on CIFAR-10. This is because `Mahalanobis` performs a search for the best noise parameter at each layer using 5-fold cross-validation, which takes a longer time on the Resnet-34 DNN for CIFAR-10. `Odds` and `LID` have much higher running time compared to the other methods. For each test sample, `Odds` computes an expectation over noisy inputs (from a few different noise parameters), which increases its running time as the size of the DNN increases. The computation involved in estimating the LID features at the DNN layers increases with the sample size used for estimation and the number of layers.

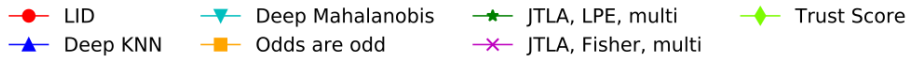

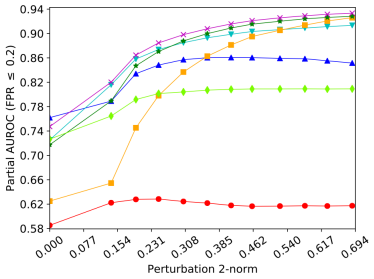

(a) CW, confidence $= 0$

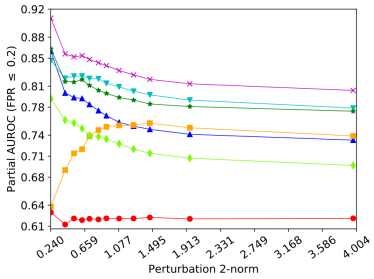

(b) Adaptive attack

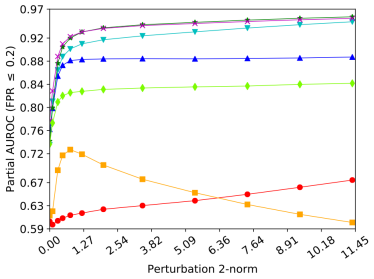

(c) FGSM, $\epsilon_{\max} = 1$

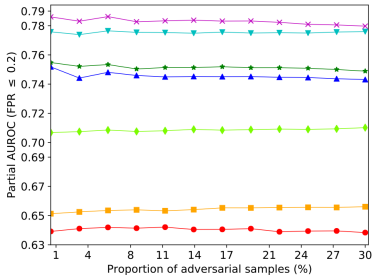

(d) PGD, $\epsilon = 1/255$

Figure 9: CIFAR-10 experiments: pAUC-0.2 for different attacks.

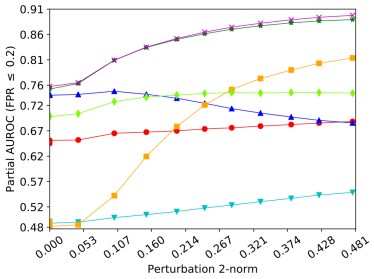

(a) CW, confidence $= 0$

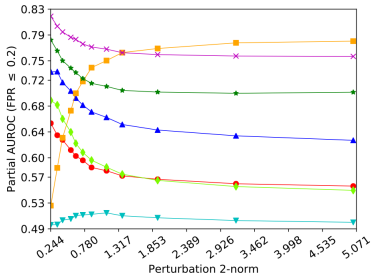

(b) Adaptive attack

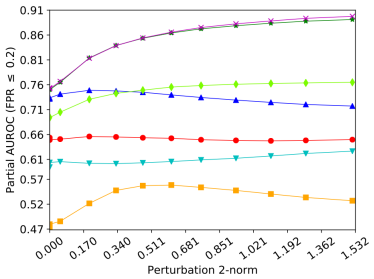

(c) FGSM, $\epsilon_{\max} = 1$

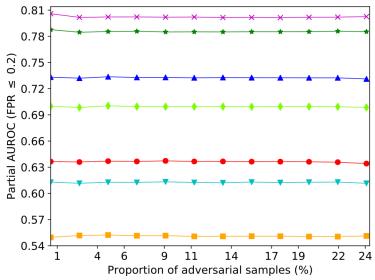

(d) PGD, $\epsilon = 1/255$

Figure 10: SVHN experiments: pAUC-0.2 for different attacks.

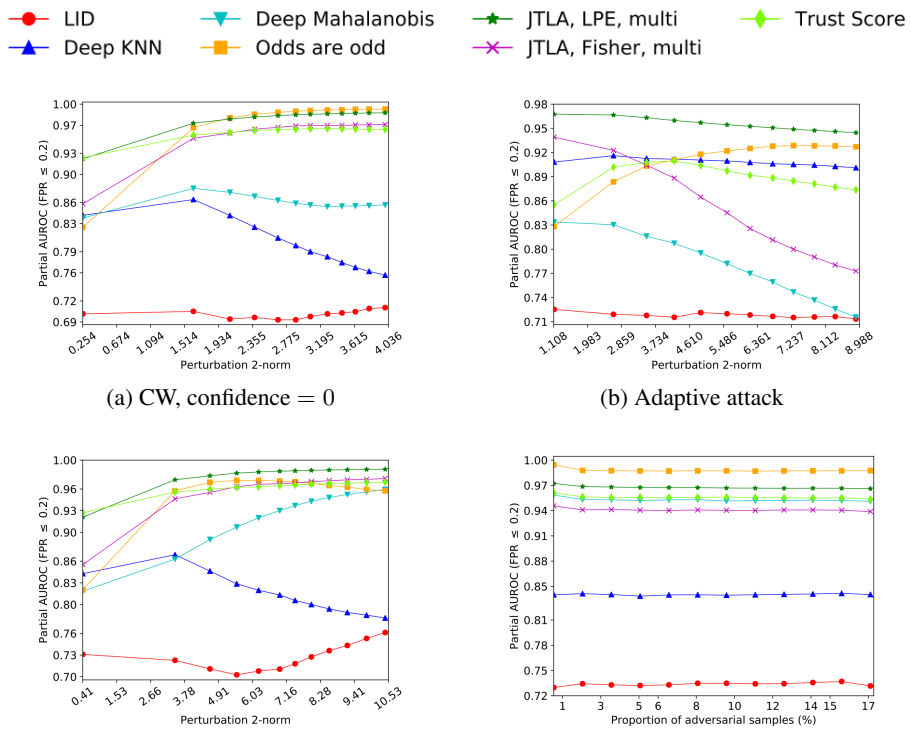

(a) CW, confidence $= 0$

(b) Adaptive attack

(c) FGSM, $\epsilon_{\max} = 1$

(d) PGD, $\epsilon = 1/255$

Figure 11: MNIST experiments: pAUC-0.2 for different attacks.

Table 9: Average wall-clock running time per-fold (in minutes) for the different detection methods.

Dataset	JTLA, Fisher	JTLA, LPE	Mahalanobis	Odds	LID	DKNN	Trust
CIFAR-10	2.73	2.18	15.08	142.94	49.74	11.99	0.53
SVHN	6.37	5.01	4.19	33.60	100.80	23.54	0.60
MNIST	0.92	0.85	1.18	6.79	6.96	1.73	0.24

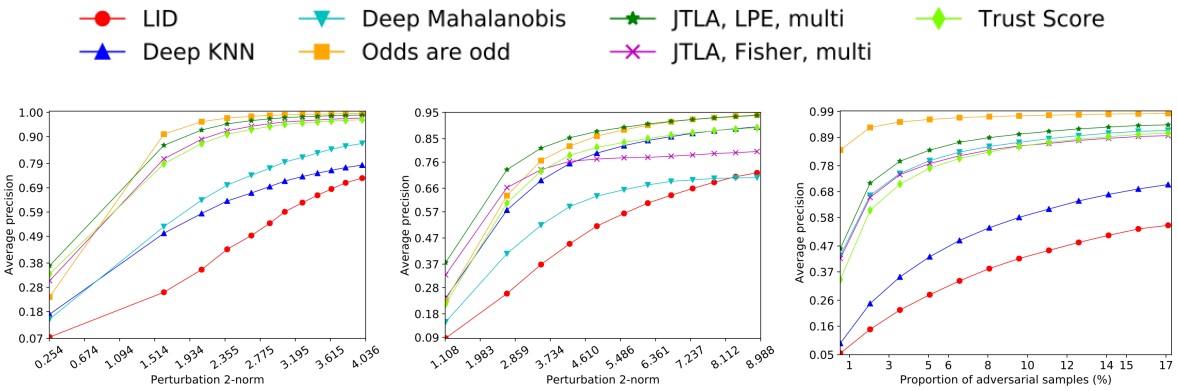

(a) CW, confidence $= 0$

(b) Adaptive attack

(c) PGD, $\epsilon = 1/255$

Figure 12: MNIST experiments: average precision for different attacks.