

Rethinking Reconstruction Autoencoder-Based Out-of-Distribution Detection

Yibo Zhou
Beihang University
ybzhou@impcas.ac.cn

Abstract

In some scenarios, classifier requires detecting out-of-distribution samples far from its training data. With desirable characteristics, reconstruction autoencoder-based methods deal with this problem by using input reconstruction error as a metric of novelty vs. normality. We formulate the essence of such approach as a quadruplet domain translation with an intrinsic bias to only query for a proxy of conditional data uncertainty. Accordingly, an improvement direction is formalized as maximumly compressing the autoencoder's latent space while ensuring its reconstructive power for acting as a described domain translator. From it, strategies are introduced including semantic reconstruction, data certainty decomposition and normalized L2 distance to substantially improve original methods, which together establish state-of-the-art performance on various benchmarks, e.g., the FPR@95%TPR of CIFAR-100 vs. TinyImagenet-crop on Wide-ResNet is 0.2%. Importantly, our method works without any additional data, hard-to-implement structure, time-consuming pipeline, and even harming the classification accuracy of known classes.

1. Introduction

Supervised discriminative deep classifiers are de facto designed under a static closed-world assumption where the data, which model face in the deployment environment, is supposed to be sampled from the same distribution as training set [12]. However, for applications in the wild, such as the safety-critical autonomous vehicles, test data is often hard to be known a priori. Worse, given that neural networks especially those built on relu or softmax can easily produce not only erroneous, but arbitrarily confident so, class predictions even for totally unrecognizable or irrelevant samples [9–11, 21], it is consequential to empower system to flag or abstain predictions from unknowns.

Out-of-distribution (OoD) detection is a binary classification of detecting inputs sampled from distribution different from training data [11]. Many existing methods rely on training or tuning with data labelled as OoD from other cat-

egories [28, 31], adversaries [15, 19] or the leave-out subset of training samples [30]. However, it is intractable to cover the full space of OoD particularly for data with large dimensions (e.g., an image) [20], resulting in that a supervised method capturing limited facets of OoD distribution hardly generalizes without a data selection bias [24]. Meanwhile, the classifier's accuracy on in-distribution (ID) images would be fluctuated by introducing OoD data with additional training objectives [28, 30]. Such factors expose the unsupervised nature of OoD detection.

In a class of unsupervised approaches, input reconstruction residual is considered a novelty metric to avoid the above problems [22]. The essential assumption is that an autoencoder learned to rebuild ID samples could not reconstruct OoD comparatively during testing [5]. However, it has been extensively reported that autoencoders can effectively rebuild kinds of OoD samples even better [7, 22], causing an inferior performance of such methods when applied to challenging multi-class OoD detection tasks. In this paper, we investigate in a precise manner the problematic reconstruction of outliers from the perspective of quadruplet domain translation by formalizing two concrete preconditions under which the reconstruction error of an input is a valid data uncertainty measure. First, its latent feature lies within the domain of encoded ID samples. Second, the decoder has adequate reconstructive power to bridge the domains of ID images and their latent representations.

Precondition 1 requires system to capture the outliers of latent representations. Without relying on kernel density estimators [1, 7, 23] performing weakly in high dimensions [17], similar to the scheme considering for compact latent representations [2, 3], we minimize a regularization loss to restrict ID latent features distributed within a certain space. In conjunction with the training of input reconstruction, it is explicated that, when this space is compressed sufficiently, any latent feature outlier will lie outside it.

There is a coupled problem that an over-restricted latent space might not provide sufficient reconstructive power for large-scale ID images, which violates precondition 2. To mitigate it, we change the reconstruction target from the image to its extracted activation vector (AV) feature to reduce the unnecessary requirement of the expressiveness of latent

space. As a further step, from the perspective of domain translation, we deduce a base equation to model the inherent connection between input reconstruction error and data certainty. By the probability chain rule, we further factorize it to express data certainty as a product of conditional densities defined by the layer-by-layer encoder feature reconstruction error. It is proved that, albeit the considerable information loss accumulated from the whole encoding process making it hard to recover input directly from a compact latent space to satisfy precondition 2, there exists an equivalent precondition feasibly to be met as it only requires to recover the information lost after each irreversible encoding layer respectively. Consequently, inspired by above concepts, we present a theoretically well-defined framework for OoD detection, namely layerwise semantic reconstruction.

In this framework, we employ only one fully connected (FC) layer and softmax function as the encoder architecture and leverage simple cross entropy loss to restrict the latent space. We provide both experimental evidences and mathematical insights, indicating that under such a setting the maximum value of a latent feature can be utilized as a domain affinity for filtering out OoD data potentially to be reconstructed. Since classifier is inclined to produce smaller neural activations on OoD data [25, 27], to render our method robust against it, reconstruction accuracy is evaluated with the proposed normalized L2 distance.

Our contribution is three-fold: First, we establish a novel perspective for understanding autoencoder-based OoD detections, figuring out one direction to improve them by maximumly condensing autoencoder’s latent space while reserving sufficient reconstructive power on ID data. Rooting from it, second, a framework of layerwise semantic reconstruction is developed. Third, along with ablation and robustness studies, we provide comprehensive analysis of our proposal using various benchmark datasets to demonstrate its efficacy, indicating that the potentiality of autoencoder-based methods is not as bleak as previously displayed.

- Our method delivers comparable performance to SOTA methods on various challenging benchmarks.
- As an auxiliary module, our OoD detector is trained in a way orthogonal to the classifier.
- Operated in an unsupervised mode of efficiency and applicability, our method requires no additional data.

2. Background and Related Work

An autoencoder consists of an encoder to project an input into a latent space with fewer dimensions (also known as bottleneck feature) and a decoder to recover the input from its latent representation. During training, input is processed sequentially through the encoder and decoder to minimize the reconstruction error which is a discrepancy

like L2 distance between the input and the reconstruction output from the decoder. The basic assumption in the reconstruction autoencoder-based OoD detector is that an autoencoder trained exclusively to recover ID samples could not succeed in reconstructing OoD samples. Hence, reconstruction error becomes a potentially effective decision function of OoD detection in testing. However, contrary to this expectation, autoencoders are reported to accurately reconstruct different types of OoD samples [18, 29].

It was pointed out in [6] that the gross noise within training samples makes it hard to learn a robust latent representation of the majority of ID samples. Hence, a neural network extension of robust principal component analysis (RPCA) [4] was proposed. In [7] the OoD detection methods of Mahalanobis distance and autoencoder were merged into a unified framework, supposing that the latter could be thus enhanced. Recent work on latent space autoregression [2] proposed to constrain the autoencoder from an identity function by applying an autoregressive density estimator to minimize the differential entropy of the distribution of ID latent features. Our proposal also takes advantage of restricting the distribution of ID latent features. Differently, we seek a maximumly compressed latent space that tightly covers the the domain of ID latent codes to get an approximation of this domain without parametric density estimator [26]. The comparison of both methods is detailed in the experiment section, showing that our method outperforms the existing method by considerable margins.

3. Hypothesis

In this section, we introduce the concepts of our framework. Formally, it is supposed that there is a domain $\mathcal{X} \subset \mathbb{R}^N$ characterized by all of the ID images. We consider a sample z as an OoD sample when it is not in \mathcal{X} . In essence, any $x \in \mathcal{X}$ is an ID sample and any $z \in \mathcal{Z} = \mathbb{C}_{\mathbb{R}^N} \mathcal{X}$ is an OoD. For a trained autoencoder consisting of an encoder $E(\cdot)$ and a decoder $D(\cdot)$ on a dataset of ID samples, let the domain spanned by $E(x), \forall x \in \mathcal{X}$ be $\mathcal{S}_{ID} \subset \mathbb{R}^M$. Under the hypothesis that *this decoder has sufficient reconstructive power to act as a domain translator of $\mathcal{S}_{ID} \mapsto \mathcal{X}$* , $E(\cdot)$ and $D(\cdot)$ are functions learning the forward and backward mapping relationships between \mathcal{X} and \mathcal{S}_{ID} :

$$E(x) \in \mathcal{S}_{ID}, \forall x \in \mathcal{X}, \tag{1}$$

$$D(f) \in \mathcal{X}, \forall f \in \mathcal{S}_{ID}. \tag{2}$$

Since \mathcal{X} and \mathcal{Z} might overlap when encoded, if we represent the domain of encoded OoD samples as \mathcal{S}_{OoD} , it could be decomposed as $\mathcal{S}_{OoD} = \mathcal{S}_{OoD} \cap \mathcal{S}_{ID} + \mathcal{S}_{OoD} \cap \mathbb{C}_{\mathbb{R}^M} \mathcal{S}_{ID}$. Using Eq.2, any OoD sample encoded into $\mathcal{S}_{ID} \cap \mathcal{S}_{OoD}$ can be decoded back to \mathcal{X} . For this part of OoD samples, we expect a larger reconstruction error since they are decoded into a different domain \mathcal{X} . Also, for these OoD samples, the

further they lie from \mathcal{X} in \mathbb{R}^N , the larger the reconstruction error might be. This is consistent with the characteristic of a measure of data certainty.

However, for an OoD sample z with latent feature $E(z) \in \mathcal{S}_{OoD} \cap \mathbb{C}_{\mathbb{R}^M} \mathcal{S}_{ID}$, it is not guaranteed to prevent the reconstruction of it. Although $E(z)$ can be far from the encoded ID data, z might lie on a manifold determined by the parameters and architecture of the autoencoder, on which data could be precisely reconstructed. An example is an OoD image of all zeros. After being processed by an autoencoder without shift operations, its reconstruction would be naturally all zeros, resulting in a perfect reconstruction of OoD. Naturally encapsulated with concept graphed in Fig.1, for an agnostic input r , given $E(r) \in \mathcal{S}_{ID}$, its reconstruction error is a valid measurement of data uncertainty $P(r \notin \mathcal{X})$. This concept is modeled as

$$P(r \in \mathcal{X} | E(r) \in \mathcal{S}_{ID}) = F(\text{Dist}(r, D(E(r)))) \tag{3}$$

for a decoder $D(\cdot)$ satisfying Eq.2, an appropriate distance metric $\text{Dist}(\cdot, \cdot)$ applied to compute reconstruction error and a certain monotonically decreasing function $F(\cdot) : \mathbb{R} \mapsto [0, 1]$ to map reconstruction error into probability. Noticing $P(r \in \mathcal{X}) = P(r \in \mathcal{X}, E(r) \in \mathcal{S}_{ID}) + P(r \in \mathcal{X}, E(r) \notin \mathcal{S}_{ID})$, by the definition of \mathcal{S}_{ID} we have $P(r \in \mathcal{X}, E(r) \notin \mathcal{S}_{ID}) = 0$, and therefore

$$\begin{aligned} P(r \in \mathcal{X}) &= P(r \in \mathcal{X}, E(r) \in \mathcal{S}_{ID}) + 0 \\ &= P(r \in \mathcal{X} | E(r) \in \mathcal{S}_{ID}) \cdot P(E(r) \in \mathcal{S}_{ID}) \tag{4} \\ &= F(\text{Dist}(r, D(E(r)))) \cdot P(E(r) \in \mathcal{S}_{ID}). \end{aligned}$$

To estimate $P(E(r) \in \mathcal{S}_{ID})$, instead of fitting $E(x) \in \mathcal{S}_{ID}$ extracted in a trained autoencoder with specific family of distributions like Gaussian to get a biased density estimator under restrictive assumption, we could apply an additional regularization loss on ID latent features in the training stage to force them to reside within a known compact space for a better approximation of \mathcal{S}_{ID} . This idea is represented in Fig.2. Since \mathcal{S}_{ID} is covered by this restricted latent space, with the increasing restrictive power imposed over ID latent features, to retain the power of the jointly learned reconstruction task, the model implicitly seek for a larger use of this latent space. It means that a restricted latent space could be more filled by \mathcal{S}_{ID} . Therefore, estimating $P(E(r) \in \mathcal{S}_{ID})$ is much equal to estimating the probability that $E(r)$ is in this restricted latent space.

However, this regularizer in the OoD detection setting is coupled simultaneously with the problem that an over-limited latent space can result in reconstruction difficulty in ID images. It potentially violates the fundamental assumption that $D(\cdot)$ acts as a domain translator of $\mathcal{S}_{ID} \mapsto \mathcal{X}$, pushing $F(\text{Dist}(r, D(E(r))))$ away from an accurate estimate of $P(r \in \mathcal{X} | E(r) \in \mathcal{S}_{ID})$. Thus, the main interest here is to maximumly condense the autoencoder’s latent space whilst retaining its power to reconstruct the ID data.

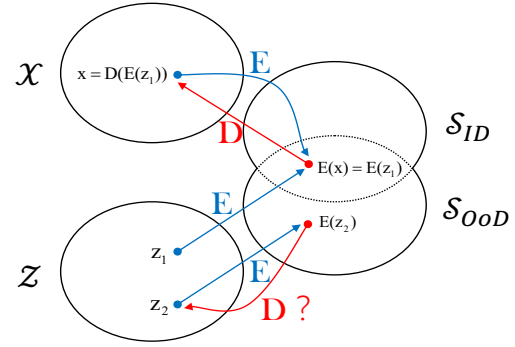


Figure 1. Illustration of the described quadruplet domain translation. For an OoD sample z_1 encoded into $\mathcal{S}_{ID} \cap \mathcal{S}_{OoD}$, its latent representation $E(z_1)$ is equal potentially to that of an ID sample x . Therefore, $E(z_1)$ can be decoded to a different sample x within \mathcal{X} , resulting in a large reconstruction error. However, for an OoD sample z_2 with latent representation $E(z_2)$ lying outside \mathcal{S}_{ID} , it offers no guarantee that it could not be reconstructed well.

Equipped with this perspective, images especially those with a large spatial resolution are not good selections for reconstruction target. Satisfying Eq.2 for image requires precisely recovering its myriad details irrelevant to OoD detection. Therefore, the unnecessary requirement rises for data collection, computation, and even worse, the size of latent space which is not desirable to estimate $P(E(r) \in \mathcal{S}_{ID})$.

Since the higher-level feature extracted in a classifier can be considered as a learned representation acquiring almost all the information about the original image except the less class-representative details [8] and the existence of a trained classifier is often implied in this topic, to avoid the complexity of reconstruction on pixel space and relief the susceptibility of original methods to the scale of image, we turn to reconstruct AV feature which is the output at penultimate layer of a classifier as the low dimensional semantic representation of image to quantify data normality. In the following, we will denote the AV feature of an agnostic image as v and the domain of AV features extracted from ID images as $\mathcal{V} \subset \mathbb{R}^H$, where H represents the dimensionality of v .

As the latent space is compressed constantly, the part of input information lost in it increases. Noticing that this unrecoverable part of information is accumulated from every single irreversible layer of the encoder, we factorize $P(v \in \mathcal{V})$ as follows, to sidestep the difficulty of recovering input directly from a latent space lacking of expressiveness: for an encoder $E(\cdot)$ with n layers, we represent $f_i(v), i = 1, \dots, n$ as the output feature at i^{th} layer given an input AV v . If we denote the domain of $f_i(v), v \in \mathcal{V}$ as \mathcal{F}_i , naturally $P(f_1(v) \in \mathcal{F}_1, f_2(v) \in \mathcal{F}_2, \dots, f_n(v) \in \mathcal{F}_n | v \in \mathcal{V}) = 1$. Therefore $P(v \in \mathcal{V}) = P(v \in \mathcal{V}) \cdot 1 = P(v \in \mathcal{V}) \cdot P(f_1(v) \in \mathcal{F}_1, f_2(v) \in \mathcal{F}_2, \dots, f_n(v) \in \mathcal{F}_n | v \in \mathcal{V}) = P(v \in \mathcal{V}, f_1(v) \in \mathcal{F}_1, f_2(v) \in \mathcal{F}_2, \dots, f_n(v) \in \mathcal{F}_n)$. Assuming that the forward propagation of this encoder forms




Figure 2. Illustration of the transition of \mathcal{S}_{ID} when the restrictive power imposed over latent codes increases. During training, any deviation of the latent codes from the latent space (red-dot circle) would be penalized greatly. By tightening this space sufficiently, in principle it would be mostly utilized to satisfy the jointly learned reconstruction task. Thus, detecting the outlier of \mathcal{S}_{ID} is approximately equal to identifying the feature outside this latent space.

a Markov process of $\mathbf{v} \xrightarrow{\text{layer1}} f_1(\mathbf{v}) \xrightarrow{\text{layer2}} f_2(\mathbf{v}) \xrightarrow{\text{layer3}} \dots \xrightarrow{\text{layer}n} f_n(\mathbf{v})$, any feature can be determined by any of its preceding feature. Hence, $P(f_i(\mathbf{v}) \in \mathcal{F}_i, f_{i+1}(\mathbf{v}) \in \mathcal{F}_{i+1}, \dots, f_n(\mathbf{v}) \in \mathcal{F}_n) = P(f_i(\mathbf{v}) \in \mathcal{F}_i)$, then we have

$$\begin{aligned}
 P(\mathbf{v} \in \mathcal{V}) &= P(\mathbf{v} \in \mathcal{V}, f_1(\mathbf{v}) \in \mathcal{F}_1, f_2(\mathbf{v}) \in \mathcal{F}_2, \dots, f_n(\mathbf{v}) \in \mathcal{F}_n) \\
 &= P(f_n(\mathbf{v}) \in \mathcal{F}_n) \cdot P(\mathbf{v} \in \mathcal{V} | f_1(\mathbf{v}) \in \mathcal{F}_1, \dots, f_n(\mathbf{v}) \in \mathcal{F}_n) \\
 &\cdot \prod_{i=1}^{n-1} P(f_i(\mathbf{v}) \in \mathcal{F}_i | f_{i+1}(\mathbf{v}) \in \mathcal{F}_{i+1}, \dots, f_n(\mathbf{v}) \in \mathcal{F}_n) \\
 &= P(f_n(\mathbf{v}) \in \mathcal{F}_n) \cdot P(\mathbf{v} \in \mathcal{V} | f_1(\mathbf{v}) \in \mathcal{F}_1) \\
 &\cdot \prod_{i=1}^{n-1} P(f_i(\mathbf{v}) \in \mathcal{F}_i | f_{i+1}(\mathbf{v}) \in \mathcal{F}_{i+1}).
 \end{aligned}$$

Extended from Eq.3, $F(\text{Dist}(f_i(\mathbf{v}), D_{i+1}(f_{i+1}(\mathbf{v}))))$ can be utilized as an estimation of $P(f_i(\mathbf{v}) \in \mathcal{F}_i | f_{i+1}(\mathbf{v}) \in \mathcal{F}_{i+1})$ for $i = 1, \dots, n-1$, where D_{i+1} is a decoder capable of recovering $f_i(\mathbf{v})$ from $f_{i+1}(\mathbf{v})$, $\forall \mathbf{v} \in \mathcal{V}$. After substitution, we get formally

$$\begin{aligned}
 P(\mathbf{v} \in \mathcal{V}) &= P(f_n(\mathbf{v}) \in \mathcal{F}_n) \cdot F(\text{Dist}(\mathbf{v}, D_1(f_1(\mathbf{v})))) \\
 &\cdot \prod_{i=1}^{n-1} F(\text{Dist}(f_i(\mathbf{v}), D_{i+1}(f_{i+1}(\mathbf{v}))))). \quad (5)
 \end{aligned}$$

In contrast to Eq.4, this reformulation in Eq.5 allows us to approximate $P(\mathbf{v} \in \mathcal{V})$ without direct input reconstruction from a restricted latent space, which requires recovering the considerable accumulated information loss at one round. Instead, it could be accomplished through a series of decoders D_i focusing on recovering the information lost after each single encoding layer, which is only a smaller divided part of the total information loss and therefore relatively easy to be recovered. Since the reconstructive ability of a decoder is closely related to its performance as an estimator of the defined conditional probability, Eq.5 allows a

better estimation of input certainty in practice, particularly when the latent space is over restricted.

4. Practical Approach

After establishing the framework, we must specify the form of $\text{Dist}(\cdot, \cdot)$, $F(\cdot)$ and the regularizer. Here we give our simple but effective designs in order.

By minimizing a reconstruction discrepancy like L_2 distance, we can train the decoders in Eq.5 (assuming that the reconstruction error follows a Gaussian distribution). However, such a norm-correlated distance metric is not appropriate for use as $\text{Dist}(\cdot, \cdot)$ in test time. Obviously, feature with smaller L_2 -norm tends to have smaller L_2 reconstruction loss, and corresponding proof is presented in the supplementary material. Unfortunately, the neural network is inclined to produce smaller neural activations on inputs that model is not familiar with [25,27], i.e., OoD. As a result, the normality measure in Eq.5 with L_1 or L_2 distance applied to evaluate reconstruction accuracy can produce a relatively larger value for OoD (see sec.5.3). This is opposite to what we expected. Thus, we propose the normalized L_2 distance (NL2) as the specific form of $\text{Dist}(\cdot, \cdot)$ to compute residuals

$$\text{Dist}(\mathbf{f}, \tilde{\mathbf{f}}) = \text{NL2}(\mathbf{f}, \tilde{\mathbf{f}}) = \left\| \frac{\mathbf{f}}{\|\mathbf{f}\|} - \frac{\tilde{\mathbf{f}}}{\|\tilde{\mathbf{f}}\|} \right\|, \quad (6)$$

where $\tilde{\mathbf{f}}$ represents the reconstruction of \mathbf{f} . It should be noted that $\tilde{\mathbf{f}}$ is normalized with $\|\tilde{\mathbf{f}}\|$ instead of the norm of itself. With NL2 distance, the reconstruction is assessed between the projection of \mathbf{f} on the surface of a unit hypersphere and its equally scaled reconstruction. Therefore, the negative influence of the feature norm is eliminated.

$F(\cdot)$ translates the reconstruction error into a conditional probability that can be measured as the inlier reconstruction error belongingness from previous discussion. Here, we interpret this belongingness as, for a given feature, the probability that its reconstruction error is smaller than that of an ID image's feature, which is naturally the complementary cumulative density function (CCDF) of the reconstruction error of ID features. Assuming that the reconstruction error of ID data distributes as Gaussian, we have

$$F(\text{NL2}(\mathbf{f}, \tilde{\mathbf{f}})) = \Psi(\text{NL2}(\mathbf{f}, \tilde{\mathbf{f}}) | \mu, \sigma + \epsilon), \quad (7)$$

in which $\Psi(\cdot)$ is the CCDF of a Gaussian and μ and σ are the parameters of Gaussian derived from validation data. Since a minor decrease leads to a substantial change of a value's numerical magnitude when it is close to 0, and our normality measure in Eq.5 is a product of multiple terms of Eq.7 within the range of 0 to 1, with increasing the number of encoder layers, the final estimation can collapse to zero rapidly even for training data. To avoid this situation, we introduce additive term ϵ in Eq.7 to prevent it from producing values smaller than 0.1 on features of ID validation data.

Algorithm 1 Training pipeline

Require: ID training set: $\{(\mathbf{x}_i, y_i)\}_{i=1}^k$, and ID validation set: $\{(\mathbf{x}_i, y_i)\}_{i=k+1}^n$

Require: Network $M(\cdot)$ fully trained on ID training set for classification of ID classes

- 1: Freeze all the parameters of network $M(\cdot)$, and jointly train $W \in \mathbb{R}^{H \times C}$ (C is the number of ID classes) and two decoders D_1 & D_2 to minimize the loss \mathcal{L}

$$\begin{aligned} \mathcal{L} &= \mathcal{L}_1 + \mathcal{L}_2 + \lambda \cdot \mathcal{L}_{regularizer} \\ \mathcal{L}_1 &= \sum_{i=1}^k \|\mathbf{v}_i - D_1(W\mathbf{v}_i)\|, \mathcal{L}_2 = \sum_{i=1}^k \left\| \frac{W\mathbf{v}_i}{T} - D_2\left(S\left(\frac{W\mathbf{v}_i}{T}\right)\right) \right\| \\ \mathcal{L}_{regularizer} &= - \sum_{i=1}^k \sum_{j=1}^C \mathbb{1}(j = y_i) \log S(W\mathbf{v}_i)_j, \end{aligned}$$

where \mathbf{v}_i is \mathbf{x}_i 's AV feature extracted in $M(\cdot)$ and λ is the weight of regularization loss

- 2: After training, compute:

$$\begin{aligned} (\mu_0, \sigma_0) &= \text{norm.fit}(\{S(\frac{W\mathbf{v}_i}{T})_{\bar{y}_i}\}_{i=k+1}^n) \\ (\mu_1, \sigma_1) &= \text{norm.fit}(\{\|\frac{\mathbf{v}_i}{\|\mathbf{v}_i\|} - \frac{D_1(W\mathbf{v}_i)}{\|\mathbf{v}_i\|}\|_{i=k+1}^n\}) \\ (\mu_2, \sigma_2) &= \text{norm.fit}(\{\|\frac{W\mathbf{v}_i}{\|W\mathbf{v}_i\|} - \frac{D_2(S(\frac{W\mathbf{v}_i}{T}))}{\|\frac{W\mathbf{v}_i}{T}\|}\|_{i=k+1}^n\}) \end{aligned}$$

- 3: **return** $D_1, D_2, W, (\mu_0, \sigma_0), (\mu_1, \sigma_1)$ and (μ_2, σ_2)
-

For an accurate estimate of $P(E(\mathbf{v}) \in \mathcal{S}_{ID})$, *i.e.*, $P(f_n(\mathbf{v}) \in \mathcal{F}_n)$ in Eq.5, it is essential to encode ID AVs into a depictable compact space closely covering \mathcal{S}_{ID} . Since a classifier $M(\cdot)$ was optimized to translate ID images into AV features capable of being converted to probability vectors close to their corresponding one-hot labels through a linear transformation (the last FC layer) followed by a softmax function, implicitly with an encoding process of $\mathbf{v} \rightarrow W\mathbf{v} \rightarrow \text{SoftMax}(W\mathbf{v})$ regularized under the same classification loss, for any $\mathbf{v} \in \mathcal{V}$ there exists $\delta \rightarrow 0$ satisfying

$$\begin{aligned} E(\mathbf{v}) = S(W\mathbf{v}) \in \mathcal{Q} &:= \{\mathbf{p} : 1 - p_y < \delta, \\ \sum_{i=1}^c p_i &= 1, p_i \geq 0 \text{ for } i = 1, 2, \dots, c\}, \end{aligned} \quad (8)$$

where $y \in Y := \{1, \dots, c\}$ represents the ground truth assigning membership to one of c ID classes and S denotes the softmax function. \mathcal{Q} is defined in a normalized space of probability, if the δ is sufficiently small, each point within \mathcal{Q} can potentially correspond to the predicted posterior probability of an ID data. Approximately, for $\forall \mathbf{p} \in \mathcal{Q}, \exists \mathbf{v} \in \mathcal{V}$ satisfying $E(\mathbf{v}) = S(W\mathbf{v}) = \mathbf{p}$, *i.e.*, for $\forall \mathbf{p} \in \mathcal{Q}, \mathbf{p} \in \mathcal{S}_{ID}$. Also, since $\mathcal{S}_{ID} \subset \mathcal{Q}$ from Eq.8, we emphasize that \mathcal{Q} is such a tractable compact space closely covering \mathcal{S}_{ID} conditioned on an adequately small δ .

Fundamentally, to ensure the δ as small as possible, we suggest not to involve additional loss term like weight decay

into the training of $M(\cdot)$ to distract it from generating ID AVs able to utterly minimize the classification loss, *i.e.*, the δ . Thus, the proposed encoder includes only one FC layer and the following softmax, and the regularization term is simply the same classification loss applied for the training of $M(\cdot)$ (here, we assumed cross entropy).

Since $\mathcal{S}_{ID} \approx \mathcal{Q}$ for a small δ , modeling $P(S(W\mathbf{v}) \in \mathcal{S}_{ID})$ is feasible by estimating the probability that $S(W\mathbf{v})_{\bar{y}}$ is larger than that of an ID data, where \bar{y} indexes the maximum value. Hence, similar to Eq.7, we express $P(S(W\mathbf{v}) \in \mathcal{S}_{ID})$ as the CDF $\Phi(S(W\mathbf{v})_{\bar{y}})$ of a (μ_0, σ_0) parameterized Gaussian. After employing the proposed two-layer encoder and substituting Eq.6 and 7 into Eq.5, we get the ultimate normality measure

$$\begin{aligned} P(\mathbf{v} \in \mathcal{V}) &= \Phi\left(S\left(\frac{W\mathbf{v}}{T}\right)_{\bar{y}} \middle| \mu_0, \sigma_0 + \epsilon_0\right) \\ &\cdot \Psi\left(\left\| \frac{\mathbf{v}}{\|\mathbf{v}\|} - \frac{D_1(W\mathbf{v})}{\|\mathbf{v}\|} \right\| \middle| \mu_1, \sigma_1 + \epsilon_1\right) \\ &\cdot \Psi\left(\left\| \frac{W\mathbf{v}}{\|W\mathbf{v}\|} - \frac{D_2\left(S\left(\frac{W\mathbf{v}}{T}\right)\right)}{\|\frac{W\mathbf{v}}{T}\|} \right\| \middle| \mu_2, \sigma_2 + \epsilon_2\right). \end{aligned} \quad (9)$$

Considering that softmax score in a calibrated probability space can better differentiate OoD from ID [16], we apply the temperature scaled logits $\frac{W\mathbf{v}}{T}$ and probability $S(\frac{W\mathbf{v}}{T})$ in Eq.9. As $W\mathbf{v} \rightarrow \frac{W\mathbf{v}}{T}$ is a bijective mapping, the term to recover information lost in it is ignored. By obtaining W, D_1, D_2 and parameters of distributions from the training pipeline summarized in Alg.1, we could compute the normality score by Eq.9 for each test image and determine the data abnormality relying on a threshold value that can be set without OoD data (see sec.5.3). If an input is recognized as ID, prediction output from the classifier is suggested for its label assignment among known classes.

5. Experiments

5.1. Experiment Settings

Network and Training Details: We adopt Dense-BC [13] and Wide-ResNet-28-10 [32] as the classifier $M(\cdot)$ for AV feature extraction. We do not apply weight decay for both models and use dropout with drop probability of 0.2 for Dense-BC. Other training details are identical to those given in [16]. Once $M(\cdot)$ has been trained on ID training set, we use the modules before the last FC layer as the feature extractor and further train an encoder along with two decoders for the proposed layerwise semantic reconstruction described in Alg.1. The FC layer of the proposed encoder is initialized with the parameters of the last FC layer in $M(\cdot)$. Moreover, each decoder is a three-layer FC network as a simple form among the possible variants, for which detailed architecture is listed in the supplementary material. For the training details, Adam solver is applied without Nesterov momentum.

Table 1. OoD detection results in CIFAR-10 and CIFAR-100. Our method is compared with three SOTA methods of DAC, ELOC and GODIN. For fair comparison, we use the results of DAC, ELOC and GODIN reported in original papers. If no result is reported as a certain setting, it is marked as $-$. Also, DAC and GODIN did not report experimental results in detection error and AUPR-in. For each evaluation metric, \uparrow means that larger value is better and \downarrow indicates that lower value is better. All values are percentages.

OoD Dataset			FPR@95%TPR \downarrow				Detection Error \downarrow		AUROC \uparrow				AUPR-In \uparrow	
			ELOC	DAC	GODIN	ours	ELOC	ours	ELOC	DAC	GODIN	ours	ELOC	ours
WRN-28-10	CIFAR-10	TINc	0.8	-	-	0.5	2.2	1.9	99.8	-	-	99.8	99.8	99.8
		TINr	2.9	1.9	-	1.5	3.8	3.1	99.4	99.5	-	99.5	99.4	99.6
		LSUNc	1.9	-	-	0.8	3.2	2.0	99.6	-	-	99.8	99.6	99.8
		LSUNr	0.9	1.5	-	0.5	2.5	2.2	99.7	99.6	-	99.7	99.7	99.7
		iSUN	-	-	-	2.9	-	4.0	-	-	-	-	99.2	-
WRN-28-100	CIFAR-100	TINc	9.2	-	-	1.5	6.7	3.4	98.2	-	-	99.4	98.4	99.5
		TINr	24.5	18.7	-	6.6	11.6	6.4	95.2	94.9	-	98.4	95.5	98.0
		LSUNc	14.2	-	-	3.7	8.2	4.8	97.4	-	-	99.0	97.6	99.1
		LSUNr	16.5	9.2	-	5.5	9.1	5.8	96.8	97.9	-	98.5	97.0	98.6
		iSUN	-	-	-	9.0	-	7.4	-	-	-	-	97.9	-
Dense-BC	CIFAR-10	TINc	1.2	-	6.6	3.7	2.6	4.6	99.7	-	98.7	98.9	99.7	99.1
		TINr	2.9	-	4.2	10.2	3.8	7.1	99.3	-	99.1	97.7	99.3	98.1
		LSUNc	3.4	-	8.5	1.7	4.1	3.5	99.3	-	98.3	99.4	99.3	99.5
		LSUNr	0.8	-	2.4	14.6	2.2	7.8	99.8	-	99.4	97.2	99.8	97.7
		iSUN	-	-	2.5	17.3	-	8.8	-	-	99.4	96.7	-	97.5
Dense-BC	CIFAR-100	TINc	8.3	-	12.2	14.8	6.3	6.6	98.4	-	97.6	97.1	98.6	98.0
		TINr	20.5	-	6.7	14.8	10.0	8.8	96.3	-	98.6	96.7	96.7	97.2
		LSUNc	14.7	-	25.0	7.8	8.5	5.5	97.4	-	95.3	98.0	97.6	98.7
		LSUNr	16.2	-	6.2	14.8	8.8	8.0	97.0	-	98.7	96.8	97.4	97.5
		iSUN	-	-	18.6	18.0	-	9.9	-	-	98.4	96.1	-	96.9

The learning rate starts at $1e-4$ and decays by a factor of 10 at 50% and 75% of total updates. Batch size and epoch number are 128 and 300, respectively. Horizontal flip with probability of 0.5 is the only data augmentation, and we do not apply weight decay, dropout, or other refinement tricks. For the hyper parameters, empirically loss weight λ is 1, temperature T is set as 100, and ϵ_i , which helps to avoid normality score from collapsing, is applied as $10 \times \sigma_i$.

Data and Evaluation Metric: 2000 images with size of 32×32 are randomly held out from the training set of CIFAR-10/100 [14] for validation use especially of the derivation of distribution parameters in Eq.9. Other 48000 images are utilized as ID dataset to train the networks. ID test set is the test split of CIFAR-10/100. OoD test sets include TinyImagenet-resize, TinyImagenet-crop, LSUN-resize, LSUN-crop and iSUN stated and released in [16]. Each test set contains 10000 images except iSUN (8925).

As for the evaluation metrics, since our method does not affect the class assignment of image as long as it is not detected as OoD, we only evaluate model’s performance of OoD/ID binary classification based on proposed normality score. Here we simply restate our evaluation metrics:

1. **FPR@95%TPR:** The value of False Positive Rate (FPR) when True Positive Rate (TPR) is 95%.
2. **AUROC:** Area Under the Receiver Operating Characteristic curve is the area under the TPR vs. FPR curve. It is a threshold independent metric and a higher value means a better trade off between TPR and FPR.
3. **AUPR:** Area Under the Precision-Recall curve shows the variation relation between precision and recall,

OoD Dataset	AUROC	
CIFAR-10	w/o. weight decay	
	TINc	98.2 / 98.9
	TINr	96.8 / 97.7
	LSUNc	99.2 / 99.4
	LSUNr	97.0 / 97.2
iSUN	96.7 / 96.7	
CIFAR-100	TINc	96.1 / 97.1
	TINr	94.5 / 96.7
	LSUNc	98.1 / 98.0
	LSUNr	93.8 / 96.8
	iSUN	93.1 / 96.1

Table 2. AUROC of OoD detection for Dense-BC trained with/without weight decay.




Figure 3. The distribution of softmax scores for models in Table.1 and 2. For better distinction we present the temperature-scaled probabilities.

which is also threshold independent. AUPR-in denotes the AUPR where IDs are specified as positive.

4. **Detection error:** Detection error is computed as the minimal value of $0.5(1-TPR) + 0.5FPR$ over all possible score thresholds, and could be interpreted as the minimum misclassification probability.

5.2. Benchmark Results

Following the benchmarks given in [16], we compare our method against three recent notable SOTA approaches: Ensemble of self supervised Leave-Out Classifiers (ELOC) [30], Generalized ODIN (GODIN) [12], and Deep Abstaining Classifier (DAC) [28]. It should be noted that both DAC and ELOC involve data treated as OoD during training. ELOC is an ensemble of multiple classifiers (in Table.1, five classifiers), and both ELOC and GODIN apply an input processing strategy requiring an additional inference and time-consuming gradient computation for each image.

The overall results are presented in Table.1. It highlights that for Wide-ResNet-28-10, across all settings, especially for CIFAR-100 and TinyImagenet-resized/LSUN-resized dataset pairs which are commonly considered as challenging cases, our prescription outperforms others with considerable margins. However, for Dense-BC our method delivers less comparable scores in a fraction of the cases.

The performance divergence between Wide-ResNet and Dense-BC is not surprising since the effectiveness of our approach leans on a sufficiently compact latent space to estimate $P(E(v) \in S_{ID})$. Specifically, our method benefits from a smaller δ in Eq.8, *i.e.*, larger prediction confidence

Table 3. The breakdown effect for each effective component of the proposed method on Wide-ResNet-28-10. The ID dataset is chosen as the much challenging CIFAR-100 for better comparison. The leftmost column is the experimental result of original LSA characterized with four latterly alternated components: 1. image: Image is selected as the reconstruction target from which reconstruction error to detect OoD is computed. 2. L2: The distance metric applied to evaluate reconstruction accuracy is L2 distance. 3. AutoReg: An autoregressive regressor is added to compress the latent space. 4.basic: Novelty score is computed by input reconstruction error instead of feature reconstruction errors. It is the framework of vanilla reconstruction autoencoder-based OoD detector introduced in sec.2. Each following column to the right represents a model modified from the one next to its left, with main modifications represented on the top, and the component after – is replaced with the component after +. The detailed configuration of each is placed in the supplementary material.

Methods		1 st LSA	→ 2 nd	→ 3 rd	→ 4 th	→ 5 th	→ 6 th
		image,L2,AutoReg,basic	-image+feature	-L2+NL2	-AutoReg+CE	-basic+layerwise	+epsilon
TINc	FPR@95%TPR ↓	42.0	99.5	5.9	3.0	0.2	1.5
	AUROC ↑	89.2	36.7	98.6	99.1	91.3	99.4
TINr	FPR@95%TPR ↓	51.2	78.8	20.3	17.7	2.5	6.6
	AUROC ↑	89.4	80.1	95.2	96.4	90.8	98.4
LSUNc	FPR@95%TPR ↓	55.8	100.0	4.7	4.1	0.8	3.7
	AUROC ↑	70.0	7.5	98.1	98.7	91.3	99.0
LSUNr	FPR@95%TPR ↓	28.2	65.6	20.0	17.7	1.7	5.5
	AUROC ↑	93.3	80.5	95.7	96.0	91.1	98.5
iSUN	FPR@95%TPR ↓	52.5	66.4	26.0	22.4	2.8	9.0
	AUROC ↑	89.4	82.3	94.8	95.8	90.8	97.9

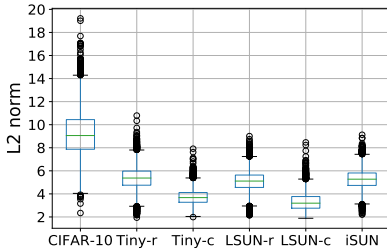


Figure 4. The distribution of AV features’ L2-norm for CIFAR-10 (ID) test set and various OoD datasets. Features are extracted in the Wide-ResNet.

of ID samples. For this purpose, as discussed previously, we remove the weight decay of 0.0001 when training the classifiers serving as AV feature extractor to gain performance boost (see Table.2). However, after we remove the weight decay for Dense-BC, the classification accuracy drops drastically from 77% to 72% for CIFAR-100 and 95% to 91% for CIFAR-10 (those for Wide-ResNet are 78% to 77% and 95% to 94%, respectively), making it harder for Dense-BC to generate ID AV features able to produce comparably high prediction confidence [11], as shown in Fig.3. Therefore, our method is much effective when employing classifier less dependent on weight decay as AV feature extractor.

5.3. Ablation Studies

To study the effectiveness and characteristics of each technical contribution, we employ the aforementioned latent space autoregression (LSA) [2], which is a leading improvement of autoencoder-based OoD detector as a basic method to be gradually modified towards our proposal, and compare the differences produced by each modification. To better discern between OoD detectors, we use CIFAR-100 as the ID dataset in the following ablation studies. The overall comparison is presented in Table.3.

AV feature as reconstruction target. We study the ef-

fect of changing the reconstruction target from the image to its AV feature. From column 1st to 2nd in Table.3, by simply use of the L2 reconstruction error computed from AV feature as decision function, it becomes much difficult to differentiate OoD from ID, even catastrophically. This is expectable since a classifier is inclined to generate AV features with smaller norm on OoD images (Fig.4), and such a phenomenon is negative for producing a larger reconstruction error for OoD samples. From column 2nd to 3rd, after altering to using the proposed normalized L2 distance as the evaluation metric of feature reconstruction, the performance in all settings is fostered utterly. That not only proves the effectiveness of the proposed NL2 distance metric, but also validates our point of view that semantic feature could enable autoencoder-based OoD detection better than image.

Cross entropy as latent space regularizer. Next, we substitute the original encoder and regularizer of LSA with the proposed ones, and apply a normality score based on Eq.4 (one decoder). The results in column 4th indicate that the detector with our encoder scores better in all the cases. Also, compared to the original autoregressor consisting of five masked FC layers with considerable parameters, our regularizer of simple cross entropy, utilizing the nature of ID AV features, is much effective in term of computation.

Layerwise reconstruction for certainty decomposition. With adjusting the basic framework to the one of layerwise reconstruction to detect OoD with the proposed certainty measure, column 5th and 6th both represent the complete versions of our method. The only difference is that in column 5th ϵ terms in Eq.9 are not applied. Comparing column 4th and 6th, our framework rooting from the proposed data certainty decomposition performs consistently better with restricted latent space. To further confirm this, we train two sequences of models with the same configura-

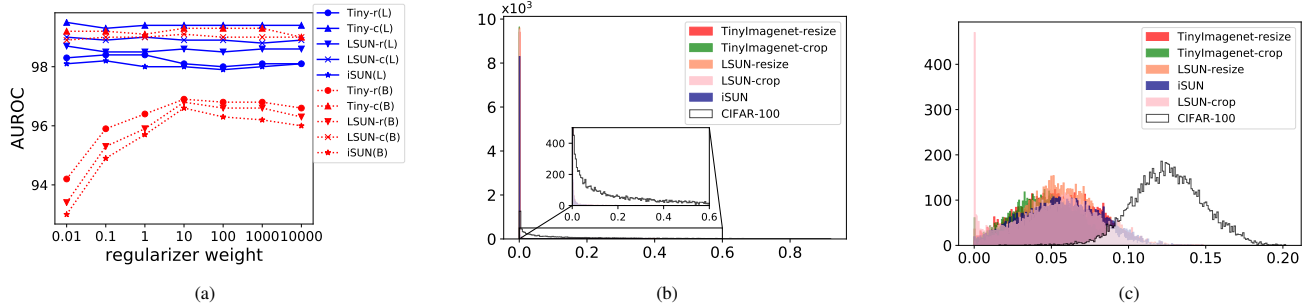


Figure 5. (a): AUROC as a function of the weight of regularizer λ for our framework of layerwise reconstruction (blue) vs. the basic framework (red). (b): The distributions of the normality scores computed from Eq.9 with ϵ_i applied as 0 and (c) with ϵ_i applied as $10 \times \sigma_i$.

tion as that of column 4th and 6th, respectively. They only differ in the loss weight λ of regularizer and their performance is compared in Fig.5a. It can be observed that for models built on the basic framework, by decreasing λ , the AUROC is improved at first and then degrades rapidly. This reveals that their performance is quite sensitive to the degree of restrictive power imposed over latent space. In contrast, our framework performs both superiority and robustness for a wide range of λ . Since the optimal weight is hard to set subtly without OoD validation data, our framework excels in both the detection performance and stability of OoD.

ϵ term. Additionally, column 5th and 6th indicate that ϵ terms help to enhance AUROC while increasing the FPR@95%TPR both to a significant extent. As shown in Fig.5b, without ϵ terms, the normality scores of OoD samples mostly condense to 0 while those of ID samples span over a wide space. It is demonstrated that under our framework the OoD sample yielding lower residuals can not have a high prediction confidence and vice versa. Especially for TinyImagenet-crop and LSUN-crop, the model in column 5th scores almost touch-the-ceiling in FPR@95%TPR. Hence, a high-recall threshold set over the output normality score is sufficient to ensure a considerably accurate OoD detection. Compared to the column 4th, column 6th has better results both for AUROC and FPR@95%TPR. Thus, if a much balanced performance between evaluation metrics is preferred, ϵ terms can be added to keep ID samples from producing low certainty (Fig.5c).

5.4. Robustness Exploration

Number of ID classes. Since the latent space employed for reconstruction has the same dimensions as the output of classifier, it may be argued that our approach can fail in cases that the number of ID classes is small. For this, we randomly select 2, 4, 6, 8 classes from CIFAR-10 and use them respectively as the ID dataset to assess the robustness of our method to the number of ID classes. As seen in Fig.6a, our method can yield convincing performance on these datasets until the one of 2 classes. One explanation

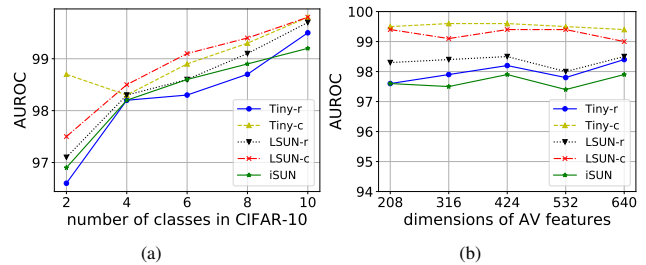


Figure 6. AUROC for OoD detection as function of (a): number of ID classes in CIFAR-10 and (b): dimensions of AV features serving as input information of autoencoder.

for this is that a latent space with smaller size is conducive to estimating $P(E(\mathbf{v}) \in \mathcal{S}_{ID})$. Also, the variation of ID AV features decreases while decreasing the number of classes, and thus less expressive power of latent space is required when the number of classes is small.

Dimensions of AV feature. Fig.6b reports a study questioning the robustness of our method to the dimensionality of AV features. For this, we train multiple Wide-ResNet-28-10 differing in the number of channels of the penultimate layer on CIFAR-100 and use them respectively to extract the AV features serving as the input of the autoencoder. The results prove the robustness of our proposal.

6. Conclusion

We propose a novel and effective framework of layerwise semantic reconstruction to enhance reconstruction autoencoder-based OoD detector by retaining the reconstructive power of autoencoder while maximumly compressing its latent space. The comprehensive experiments demonstrate that our method achieves SOTA performance on multi-class OoD detection with merits including orthogonal to classifier, efficient in computation, and OoD data free. Importantly, our theorized perspectives of quadruplet domain translation and data uncertainty decomposition differ from the paradigms of existing methods, heuristically exploring rooms for future work.

References

- [1] Joint detection and recounting of abnormal events by learning deep generic knowledge. In *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017. 1
- [2] D. Abati, A. Porrello, S. Calderara, and R. Cucchiara. Latent space autoregression for novelty detection. 2018. 1, 2, 7
- [3] J Ballé, V. Laparra, and E. P. Simoncelli. End-to-end optimized image compression. In *ICLR 2017*, 2016. 1
- [4] E. Candes, X. Li, M. Yi, and J. Wright. Robust principal component analysis?: Recovering low-rank matrices from sparse errors. *IEEE*, 2010. 2
- [5] R. Chalapathy and S. Chawla. Deep learning for anomaly detection: A survey. 2019. 1
- [6] R. Chalapathy, A. K. Menon, and S. Chawla. Robust, deep and inductive anomaly detection. *Springer, Cham*, 2017. 2
- [7] Taylor Denouden, Rick Salay, Krzysztof Czarnecki, Vahdat Abdelzad, Buu Phan, and Sachin Vernekar. Improving reconstruction autoencoder out-of-distribution detection with mahalabis distance. *arXiv preprint arXiv:1812.02765*, 2018. 1, 2
- [8] A. Dosovitskiy and T. Brox. Inverting visual representations with convolutional networks. In *IEEE Conference on Computer Vision & Pattern Recognition*, 2015. 3
- [9] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. 1
- [10] Matthias Hein, Maksym Andriushchenko, and Julian Bitterwolf. Why relu networks yield high-confidence predictions far away from the training data and how to mitigate the problem. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 41–50, 2019. 1
- [11] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016. 1, 7
- [12] Yen-Chang Hsu, Yilin Shen, Hongxia Jin, and Zsolt Kira. Generalized odin: Detecting out-of-distribution image without learning from out-of-distribution data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10951–10960, 2020. 1, 6
- [13] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. 5
- [14] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. *Handbook of Systemic Autoimmune Diseases*, 1(4), 2009. 6
- [15] Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *Advances in Neural Information Processing Systems*, pages 7167–7177, 2018. 1
- [16] Shiyu Liang, Yixuan Li, and Rayadurgam Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. *arXiv preprint arXiv:1706.02690*, 2017. 5, 6
- [17] Q. Liu, J. Xu, R. Jiang, and W. H. Wong. Roundtrip: A deep generative neural density estimator. 2020. 1
- [18] O. Lyudchik. Outlier detection using autoencoders. 2016. 2
- [19] Andrey Malinin and Mark Gales. Reverse kl-divergence training of prior networks: Improved uncertainty and adversarial robustness. In *Advances in Neural Information Processing Systems*, pages 14547–14558, 2019. 1
- [20] J. Mccarthy. Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence*, 4, 1969. 1
- [21] Anh Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 427–436, 2015. 1
- [22] Marco AF Pimentel, David A Clifton, Lei Clifton, and Lionel Tarassenko. A review of novelty detection. *Signal Processing*, 99:215–249, 2014. 1
- [23] M. Rosenblatt. Remarks on some nonparametric estimates of a density function. *Annals of Mathematical Statistics*, 27(3):832–837, 1956. 1
- [24] A. Shafaei, M. Schmidt, and J. J. Little. A less biased evaluation of out-of-distribution sample detectors. 2018. 1
- [25] Gabi Shalev, Yossi Adi, and Joseph Keshet. Out-of-distribution detection using multiple semantic label representations. In *Advances in Neural Information Processing Systems*, pages 7375–7385, 2018. 2, 4
- [26] Q. Song. Deep autoencoding gaussian mixture model for unsupervised anomaly detection. 2018. 2
- [27] Yaniv Taigman, Ming Yang, Marc’Aurelio Ranzato, and Lior Wolf. Web-scale training for face identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2746–2754, 2015. 2, 4
- [28] S. Thulasidasan, Sushil Thapa, Sayera Dhaubhadel, Gopinath Chennupati, Tanmoy Bhattacharya, and J. Bilmes. An effective baseline for robustness to distributional shift. *ArXiv*, abs/2105.07107, 2021. 1, 6
- [29] A. Tong, G. Wolf, and S. Krishnaswamy. Fixing bias in reconstruction-based anomaly detection with lipschitz discriminators. In *2020 IEEE 30th International Workshop on Machine Learning for Signal Processing (MLSP)*, 2020. 2
- [30] Apoorv Vyas, Nataraj Jammalamadaka, Xia Zhu, Dipankar Das, Bharat Kaul, and Theodore L Willke. Out-of-distribution detection using an ensemble of self supervised leave-out classifiers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 550–564, 2018. 1, 6
- [31] Qing Yu and Kiyoharu Aizawa. Unsupervised out-of-distribution detection by maximum classifier discrepancy. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9518–9526, 2019. 1
- [32] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016. 5