
ANOMALY DETECTION BY CONTEXT CONTRASTING

Anonymous authors

Paper under double-blind review

ABSTRACT

Anomaly detection focuses on identifying samples that deviate from the norm. When working with high-dimensional data such as images, a crucial requirement for detecting anomalous patterns is learning lower-dimensional representations that capture concepts of normality. Recent advances in self-supervised learning have shown great promise in this regard. However, many successful self-supervised anomaly detection methods assume prior knowledge about anomalies to create synthetic outliers during training. Yet, in real-world applications, we often do not know what to expect from unseen data, and we can solely leverage knowledge about normal data. In this work, we propose CON₂, which learns representations through context augmentations that allow us to observe samples from two distinct perspectives while keeping the invariances of normal data. CON₂ learns rich representations of context-augmented samples by clustering them according to their context while simultaneously aligning their positions across clusters. At test time, representations of anomalies that do not adhere to the invariances of normal data then deviate from their respective context cluster. Learning representations in such a way thus allows us to detect anomalies without making assumptions about anomalous data.

1 INTRODUCTION

Reliably detecting anomalies is essential in many safety-critical fields such as healthcare (Schlegl et al., 2017; Ryser et al., 2022), finance (Golmohammadi & Zaiane, 2015), industrial fault detection (Atha & Jahanshahi, 2018; Zhao et al., 2019), or cyber-security (Xin et al., 2018). A common real-world example of anomaly detection is the standard screening scenario, where doctors regularly examine the general population for anomalies that would indicate a health risk. Standard screening datasets thus predominantly comprise samples from healthy people, as most screened individuals do not exhibit any diseases. Detecting anomalies in this setting is challenging, as anomalies can arise from an arbitrary set of potentially rare diseases or measurement errors, while we predominantly encounter normal samples from healthy people in the dataset. The field of anomaly detection tackles such problems by learning representations that reflect normality during training and, at test time, detecting anomalies as deviations from the learned normal structure (Ruff et al., 2021).

Recent works have demonstrated that learning a representation space containing features that tightly represent normality is essential for anomaly detection (Ruff et al., 2018; Oza & Patel, 2018; Sabokrou et al., 2020). Current state-of-the-art methods carefully design synthetic anomalies and explicitly encourage anomalous representations to be different from normal ones (Tack et al., 2020; Wang et al., 2023). However, anomalies can be diverse and unexpected, making it difficult to simulate them in real-world settings.

This work presents a novel anomaly detection objective, CON₂¹, which learns informative, tightly clustered representations of normal samples. We illustrate an overview of the algorithm in Figure 1. Unlike previous works, which focus on prior knowledge about anomalies, the proposed CON₂ models properties of normal samples, which is particularly useful in more specialized data, such as in the medical domain, which we demonstrate in our experiments. CON₂ leverages *context augmentations* that let us observe samples in different contexts while preserving their normal content. Our new CON₂ objective clusters representations according to these contexts while encouraging similar rep-

¹The code is attached to this submission and will be made publicly available upon acceptance.

054
055
056
057
058
059
060
061
062
063
064
065
066
067
068
069
070
071
072
073
074
075
076
077
078
079
080
081
082
083
084
085
086
087
088
089
090
091
092
093
094
095
096
097
098
099
100
101
102
103
104
105
106
107

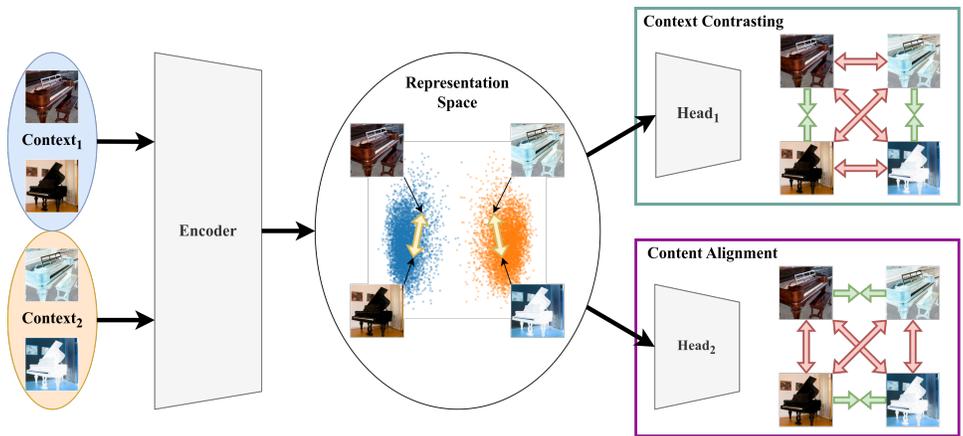


Figure 1: Overview of CON₂. Samples get context augmented and passed through an encoder. The *context contrasting* loss (Equation (2)) ensures context-specific representations (■ and ■ clusters) while the *content alignment* loss (Equation (3)) encourages a context independent structure (↔) within each context cluster. We learn representations in a contrastive fashion, matching corresponding positive (⇒⇐) and discriminating between negative (⇐⇒) pairs of representations separately for *context contrasting* and *content alignment*.

representations within each cluster. Consequently, CON₂ ensures a highly informative structure within each cluster by preserving the relative normality of samples independent of their context.

Our main contributions include the definition of context augmentations to model invariances in normal data and the introduction of CON₂, which uses context augmentations to learn informative, tightly clustered representations of normal data. We further present the anomaly score function \mathcal{S}_{NND} that measures the anomalousness of new samples given representations from CON₂. Additionally, we propose the \mathcal{S}_{LH} anomaly score, which offers a more compute efficient alternative to our initial anomaly score. Finally, we demonstrate the advantage of modeling invariances of normal data in our experiments, where we present strong results when performing anomaly detection on specialized medical and more general natural image datasets.

In the next section, we provide an overview of related work and draw a comparison to our approach before proceeding to introduce our method.

2 RELATED WORK

Learning useful normal representations of high-dimensional data to perform anomaly detection has recently become a popular line of research. Prior work has tackled the problem from various angles, for instance, using hypersphere compression (Ruff et al., 2018). Other popular methods define pretext tasks such as learning reconstruction models (Chen et al., 2017; Zong et al., 2018; You et al., 2019) or predicting data transformations (Golan & El-Yaniv, 2018; Hendrycks et al., 2019b; Bergman & Hoshen, 2019). While these approaches had some success in the past, the learned representations are not very informative. On the other hand, methods learning more informative representations through self-supervised learning have recently been shown to improve over prior work (Sun et al., 2022; Schwag et al., 2021).

Another line of work focuses on estimating the training density with the help of generative models, detecting anomalies as samples from low probability regions (An & Cho, 2015; Schlegl et al., 2019; Nachman & Shih, 2020; Mirzaei et al., 2022). However, these methods tend to generalize better to unseen distributions than to the observed training distribution (Nalisnick et al., 2018), which proves problematic for anomaly detection.

In addition to the traditional setting, where we assume training data without any labels, some recent works weaken this restriction and assume access to a limited number of labeled samples. This setting is called anomaly detection with Outlier Exposure (OE) (Hendrycks et al., 2019a), and it has



117 Figure 2: Examples of the context augmentations used throughout our experiments. *Flip* denotes
118 vertical flipping, *Invert* denotes the transformation that replaces each pixel value x with $1 - x$, and
119 *Equalize* stands for histogram equalization. In our experiments, *Flip* and *Invert* fulfill alignment and
120 distinctiveness for almost all datasets, while *Equalize* can sometimes violate distinctiveness.

121
122
123 been shown that already a few labeled samples can greatly boost performance over an unlabeled
124 dataset (Ruff et al., 2020; Qiu et al., 2022; Liznerski et al., 2022). Using large, pretrained models
125 as feature extractors is a special case of OE, as additional data is not explicitly accessible. Some
126 approaches have been introduced that use representations from pretrained models directly in zero-
127 shot fashion (Bergman et al., 2020; Liznerski et al., 2022; Jeong et al., 2023; Zhou et al., 2024),
128 while others demonstrate the benefit of fine-tuning (Cohen & Avidan, 2022; Reiss & Hoshen, 2023;
129 Li et al., 2023). OE has been very successful in the past, often outperforming traditional anomaly
130 detection settings across many benchmarks, though at the cost of either requiring labeled samples
131 or vast amounts of data for pretraining, which are both often not available or hard to obtain in more
132 specialized domains.

133 Another setting that has recently gained popularity is out-of-distribution (OOD) detection. In OOD
134 detection, we have additional information about our dataset in the form of labels. Anomaly detection
135 is a special case of OOD detection with only a single label. While the problem is similar, most
136 approaches that tackle OOD detection make specific use of a classifier trained on the dataset labels
137 (Hendrycks & Gimpel, 2017; Lee et al., 2018; Wang et al., 2022), which cannot directly be applied
138 in the anomaly detection setting, as training a classifier on a single class is not straightforward.

139 In comparison, our method operates in the traditional anomaly detection setting and can be applied
140 to datasets without knowledge about anomalies. Further, while we assume access to a dataset con-
141 taining normal samples, our method does not rely on additional labels, as they can be difficult and
142 expensive to obtain, particularly in more specialized settings.

143 3 METHODS

144
145
146 In the following, we introduce the notion of context augmentations and then present our CON_2 ob-
147 jective, which leverages these augmentations to learn tightly clustered, informative representations.
148 We then explain how to use these representations to detect anomalies at test time.

149 3.1 CONTEXT AUGMENTATION

150
151
152 The intuition behind context augmentation comes from the observation that certain transformations
153 can augment a sample into another context, creating a distinct new view without altering its infor-
154 mation content. For example, inverting an image, i.e., exchanging every pixel value x with the value
155 $1 - x$, neither adds nor destroys any information (*alignment*) but instead allows us to observe the same
156 sample from a different perspective. In the following, we want to learn these invariances while still
157 being able to distinguish between the two transformations to learn symmetric representation clusters
158 for both the original and the augmented sample space. In the previous example, this only works if
159 the inverted version of an image does not naturally appear in the dataset already (*distinctiveness*).
160 Otherwise, it is impossible to distinguish between original and transformed instances properly.

161 We define two requirements that let us determine whether a transformation is suitable as a context
augmentation for a given dataset. Let $X \subset \mathcal{X}$ be our dataset, let $t_C : \mathcal{X} \rightarrow \mathcal{X}$ be a data augmentation,

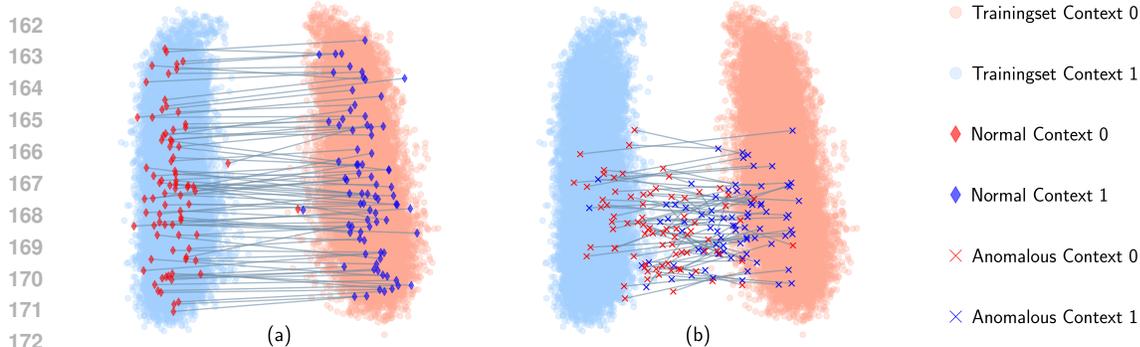


Figure 3: Two-dimensional PCA embedding of the train, normal test (a), and anomalous test samples (b) after training CON_2 on the *car* class of CIFAR10. The lines connecting representations mark embeddings corresponding to the same sample in different contexts. The parallel lines indicate that sample representations are positioned approximately at the same location across context clusters for the normal test samples, while anomalies do not exhibit the same invariances as normal samples and thus fail to adhere well to the learned structure.

and let $t_C(X) = \{t_C(\mathbf{x}) \mid \mathbf{x} \in X\}$ be the dataset transformed by t_C . The function t_C is a *context augmentation* if it fulfills the following two properties:

Distinctiveness For two samples $\mathbf{x} \sim p_X$ from the original and $\mathbf{x}^C \sim p_{t_C(X)}$ from the context augmented data distribution, we have that $p_{t_C(X)}(\mathbf{x}) \approx 0$ and $p_X(\mathbf{x}^C) \approx 0$, i.e., there is a clear distinction between the original and the context augmented distribution after applying t_C . For instance, if our normal class consists of images of melanoma, flipping the image violates distinctiveness, as melanoma can be photographed from any angle. Conversely, histogram equalizing or color inversion of the image satisfies distinctiveness, as the resulting color distribution is distinct from the original samples of such a dataset.

Alignment Let $\mathbf{x}, \mathbf{x}' \in X$, and let $d(\mathbf{x}, \mathbf{x}')$ denote an appropriate similarity measure for samples in the input space. Then, we require that $d(\mathbf{x}, \mathbf{x}') \approx d(t_C(\mathbf{x}), t_C(\mathbf{x}'))$, i.e., originally similar normal samples should stay just as similar in the new context, meaning that the original and the context-augmented normal distributions should align. For instance, masking part of a torso x-ray image would violate alignment, as we could potentially remove important regions, such as the lungs, from the image altogether. On the other hand, two vertically flipped x-rays are as similar to each other as their original counterparts.

While it may be dataset-dependent whether a transformation, such as histogram equalization (*Equalize*), fulfills these conditions, there are transformations, such as vertical flipping (*Flip*) or color inversion (*Invert*), that seem to fulfill distinctiveness and alignment across a broad range of datasets. We present some examples of context augmentations in Figure 2.

3.2 CONTEXT CONTRASTING

Our method learns representations in a contrastive fashion (van den Oord et al., 2019). Contrastive learning is a popular approach for self-supervised representation learning. Typically, it relies on the definition of positive and negative pairs of samples and learns to maximize the similarity of representations of positive pairs while pushing apart representations of negative pairs. Popular contrastive approaches, such as SimCLR (Chen et al., 2020) or SupCon (Khosla et al., 2020), achieve this by incorporating a form of instance discrimination in their loss function. Here, we define the instance discrimination loss as

$$\ell(\mathbf{x}, \mathbf{x}', X) = -\log \frac{\exp(\text{sim}(\mathbf{x}, \mathbf{x}')/\tau)}{\sum_{\mathbf{x}'' \in X: \mathbf{x}'' \neq \mathbf{x}} \exp(\text{sim}(\mathbf{x}, \mathbf{x}'')/\tau)}, \quad (1)$$

where we consider $\text{sim}(\mathbf{x}, \mathbf{x}')$ to be the cosine similarity between two samples $\mathbf{x}, \mathbf{x}' \in X$. We refer to Appendix A.1 for more background about contrastive learning.

In the following, we present our novel CON_2 objective, which leverages the distinctiveness and alignment assumptions of context augmentations to learn informative representations, which we will later use for anomaly detection. Specifically, we apply distinctiveness to learn context-specific representation clusters. Alignment further allows us to distinguish samples from each other while encouraging a similar relative location of a sample across clusters. We present an example underlining this intuition in Figure 3, where we show our representation space after training a model with CON_2 on the samples of the *car* class of CIFAR10.

Assume a set of samples X_{train} , a context augmentation t_C , a set of augmentations \mathcal{T} that models invariances of the dataset like in Chen et al. (2020), and let

$$X_C = \{(\mathbf{x}, 0) \mid \mathbf{x} \in X_{\text{train}}\} \cup \{(t_C(\mathbf{x}), 1) \mid \mathbf{x} \in X_{\text{train}}\}$$

denote the dataset after applying context augmentation t_C , labeling each sample with its context. For $t, t' \sim \mathcal{T}$ and $\mathbf{x}_i^C \in X_C$, let $\tilde{\mathbf{x}}_{2i} = t(\mathbf{x}_i^C)$ and $\tilde{\mathbf{x}}_{2i+1} = t'(\mathbf{x}_i^C)$ denote two transformations of the same context-augmented sample using random augmentations from \mathcal{T} and let the set of all such pairs be

$$\tilde{X}_C = \{(t(\mathbf{x}^C), y), (t'(\mathbf{x}^C), y) \mid (\mathbf{x}^C, y) \in X_C \wedge t, t' \sim \mathcal{T}\}.$$

Further, we denote $f(\tilde{X}_C) := \{(f(\mathbf{x}), y) \mid \mathbf{x} \in \tilde{X}_C\}$ for any function f . CON_2 then consists of two parts, *context contrasting* and *content alignment*.

Context Contrasting By leveraging the distinctiveness property of context augmentations, we can learn tightly concentrated, context-specific representation clusters with our *context contrasting* loss. For a given sample \mathbf{x} , we derive its representation $g_\theta(\mathbf{x})$ using an encoder g_θ . We then define the *context contrasting* loss as

$$\mathcal{L}_{\text{Context}}(\tilde{X}_C) = \frac{1}{4N} \sum_{(\tilde{\mathbf{x}}_i, y_i) \in \tilde{X}_C} \frac{1}{2N-1} \sum_{\substack{(\tilde{\mathbf{x}}_j, y_j) \in \tilde{X}_C \\ \tilde{\mathbf{x}}_j \neq \tilde{\mathbf{x}}_i \wedge y_i = y_j}} \ell(f_\Phi(\tilde{\mathbf{x}}_i), f_\Phi(\tilde{\mathbf{x}}_j), f_\Phi(\tilde{X}_C)), \quad (2)$$

where $f_\Phi = h_\phi(g_\theta(\mathbf{x}))$ and h_ϕ is a projection head that gets discarded after training similar to Chen et al. (2020). Intuitively, context contrasting encourages representations of the same context to be clustered together while pushing other context clusters away, similar to the class representations in supervised contrastive learning (SupCon) (Khosla et al., 2020).

Content Alignment While $\mathcal{L}_{\text{Context}}$ allows us to learn context-dependent representation clusters, it does not enforce a specific structure within each cluster. To make the cluster structure more informative, CON_2 leverages the alignment property of context augmentations to align representations across clusters through context-independent instance discrimination. More specifically, let $\Lambda(i) = \{2i, 2i+1, 4i, 4i+1\}$, i.e., $\Lambda(i)$ corresponds to all indices² of samples in \tilde{X}_C which are augmentations of the original sample $\mathbf{x}_i \in X$. We then define the *content alignment* loss as

$$\mathcal{L}_{\text{Content}}(\tilde{X}_C) = \frac{1}{N} \sum_{k=1}^N \frac{1}{12} \sum_{i \in \Lambda(k)} \sum_{j \in \Lambda(k) \setminus i} \ell(f_\Psi(\tilde{\mathbf{x}}_i), f_\Psi(\tilde{\mathbf{x}}_j), f_\Psi(\tilde{X}_C)), \quad (3)$$

where $f_\Psi(\mathbf{x}) = h_\psi(g_\theta(\mathbf{x}))$, and h_ψ denotes a projection head that is independent of h_ϕ . Content alignment ensures that all representations of the same normal sample can be matched across different contexts, encouraging alignment of the representations within each context cluster.

Finally, we combine context contrasting and content alignment to our loss function CON_2 , which enables us to learn *context-specific, content-aligned* representations of normality:

$$\mathcal{L}_{\text{CON}_2}(\tilde{X}_C) = \mathcal{L}_{\text{Context}}(\tilde{X}_C) + \alpha \mathcal{L}_{\text{Content}}(\tilde{X}_C) \quad (4)$$

To account for the different scalings of $\mathcal{L}_{\text{Context}}$ and $\mathcal{L}_{\text{Content}}$, we introduce a weighting factor $\alpha \in \mathbb{R}^+$. Figure 1 provides a visual overview of how we apply CON_2 to learn representations.

²Indexing in correspondence to previous section. Strict ordering is not necessary.

3.3 ANOMALY DETECTION

In the anomaly detection setting, we typically assume an unlabeled training set containing predominantly normal samples, whereas we want to discriminate between normal and anomalous samples at test time (Ruff et al., 2021). To detect anomalies, we typically define an anomaly score function \mathcal{S} that maps a given sample’s representation onto a scalar, determining its anomalousness. We can then define a threshold on this anomaly score, predicting *anomaly* for samples above the threshold and *normal* for samples below. We provide additional background about the anomaly detection setting in Appendix A.2.

To detect anomalies using the representations of CON_2 , we define two anomaly score functions that measure how well a test sample adheres to the context representation clusters. One of the most popular and straightforward approaches to achieve this is a non-parametric nearest neighbor approach (Bergman et al., 2020; Sun et al., 2022). Our first score adopts a similar procedure using the cosine similarity. Specifically, let us define the cosine distance between the training set X_{train} and a given test sample \mathbf{x} with transformation t as

$$s_{\text{NND}}(\mathbf{x}; t) = - \max_{\mathbf{x}' \in X_{\text{train}}} \frac{\langle g_{\theta}(t(\mathbf{x})), g_{\theta}(t(\mathbf{x}')) \rangle}{\|g_{\theta}(t(\mathbf{x}))\| \|g_{\theta}(t(\mathbf{x}'))\|}. \quad (5)$$

Intuitively, the better a new sample aligns with the context cluster given by augmentation t , the more likely we are to consider it to be normal. In turn, for samples with a lower cosine similarity, it seems to either be difficult to assign the correct context cluster, or they do not share much of the normal information within the correct context cluster. While this approach works well in practice, it is rather memory-inefficient, as we need to store the representations of all samples in X_{train} .

To adapt our approach to settings with resource constraints, we further introduce a likelihood-based score function s_{LH} . To make this score function as lightweight as possible, we assume that representations within each context cluster are distributed according to a multivariate Gaussian, which allows us to efficiently estimate the empirical mean and covariance from the training set and evaluate the probability density to derive an anomaly score. Contrastive approaches typically tend to learn representations with relatively large norms, which may lead to numerical instabilities when estimating the covariance matrix. Our s_{LH} thus estimates the empirical mean and covariance on the normalized representations. In particular, let

$$Z_{\text{train}}^{(t)} = \left\{ \frac{g_{\theta}(t(\mathbf{x}))}{\|g_{\theta}(t(\mathbf{x}))\|} \mid \mathbf{x} \in X_{\text{train}} \right\} \quad (6)$$

be the normalized representations of the training set augmented with some augmentation t . We then compute the density of a multivariate normal distribution based on the empirical mean and covariance,

$$\bar{\mu} \left(Z_{\text{train}}^{(t)} \right) \text{ and } \bar{\Sigma} \left(Z_{\text{train}}^{(t)} \right). \quad (7)$$

We then define

$$s_{\text{LH}}(\mathbf{x}; t) = - \log \left(\mathcal{N} \left(\frac{g_{\theta}(t(\mathbf{x}))}{\|g_{\theta}(t(\mathbf{x}))\|} \mid \bar{\mu} \left(Z_{\text{train}}^{(t)} \right), \bar{\Sigma} \left(Z_{\text{train}}^{(t)} \right) \right) \right). \quad (8)$$

We further leverage that our model can differentiate between the two contexts and learns invariances across different augmentations from \mathcal{T} by applying test-time augmentations, similar to previous works (Tack et al., 2020; Wang et al., 2023), which further improves our anomaly detection performance. More specifically, let $\mathcal{T}_{\text{test}} = \{t_1, \dots, t_A\}$ be a set of A test time augmentations. We define our final anomaly score functions $\mathcal{S}_{\{\text{NND}, \text{LH}\}} : \mathcal{X} \rightarrow \mathbb{R}$ as

$$\mathcal{S}_{\{\text{NND}, \text{LH}\}}(\mathbf{x}) = \frac{1}{A} \left(\sum_{i=1}^{A/2} s_{\{\text{NND}, \text{LH}\}}(\mathbf{x}; t_i) + \sum_{i=A/2}^A s_{\{\text{NND}, \text{LH}\}}(\mathbf{x}; t_i \circ t_C) \right), \quad (9)$$

where \circ defines the composition of two functions (Peirce, 1852).

4 EXPERIMENTS

In the following, we present how CON_2 allows us to learn highly informative representations of normality by incorporating prior knowledge about invariances of normal data. After briefly introducing our baselines, we demonstrate how we can leverage this knowledge in a realistic medical

setting, showcasing the applicability of our method to a specialized domain where prior knowledge about anomalies is typically hard to obtain. Further, we present the generality of our method by comparing it to baselines on popular natural image datasets in the one-class classification setting, where anomaly detection with CON₂ consistently exhibits strong performance across various settings. We refer to Appendices C and D for more details regarding the choice of hyperparameters and our datasets.

Baselines We compare our work to various recent contrastive anomaly detection baselines, including SSD (Sehwag et al., 2021), CSI (Tack et al., 2020), and UniCon-HA (Wang et al., 2023). SSD works by learning representations using SimCLR and detecting anomalies with a Mahalanobis distance-based anomaly score. Similarly, CSI and UniCon-HA learn representations with SimCLR but additionally design synthetic anomalies using rotation transformations. CSI leverages these synthetic anomalies using an additional classifier to discriminate between normal and synthetic anomaly samples and detects anomalies at test time with a score that combines nearest neighbor distance, sample norm, and classifier confidence. UniCon-HA does not require an additional classifier but instead clusters all normal samples close to each other while minimizing the similarity of synthetic anomaly representations and normal training samples. UniCon-HA also modifies the instance discrimination loss to weight positive and negative pairs according to the distance between representations. It further introduces a hierarchical augmentation scheme that lets them apply their loss on different layers of their neural network architecture using layer-specific augmentation strategies. We also compare against a baseline that learns SimCLR embeddings and detects samples in nearest neighbor fashion similar to KNN+ (Sun et al., 2022), which was originally developed for out-of-distribution detection. Finally, we also compare to anomaly detection using CLIP (Radford et al., 2021; Liznerski et al., 2022), which detects anomalies by using a pretrained CLIP model and comparing image embeddings with text embeddings describing the normal class. Apart from CLIP-AD, we conduct all experiments with the ResNet18 architecture (He et al., 2016) to ensure comparability between methods.

4.1 MEDICAL ANOMALY DETECTION

In this experiment, we demonstrate how incorporating prior knowledge about invariances of normal data through context augmentations with CON₂ leads to strong anomaly detection performance on two challenging medical imaging datasets. We compare the performance of CON₂ with recent unsupervised anomaly detection methods. Additionally, we also compare to CLIP-AD (Liznerski et al., 2022), which relies on a pretrained CLIP model (Radford et al., 2021) and thus incorporates a form of outlier exposure as explained in Section 2.

Table 1: Anomaly detection results on two real-world medical imaging datasets. We train each model with three different seeds and report the mean \pm standard deviation.

Method	Score \mathcal{S}	Pneumonia	Melanoma
CLIP-AD	$\mathcal{S}_{\text{CLIP}}$	71.2	77.2
SimCLR	\mathcal{S}_{NND}	91.0 \pm 0.9	72.9 \pm 2.8
SSD	$\mathcal{S}_{\text{Mahalanobis}}$	90.9 \pm 0.2	79.0 \pm 2.2
CSI	\mathcal{S}_{CSI}	73.9 \pm 1.6	92.3 \pm 0.2
UniCon-HA	$\mathcal{S}_{\text{UniCon}}$	86.4 \pm 0.1	91.1 \pm 0.8
CON ₂ (Equalize)		93.0 \pm 0.3	94.0 \pm 0.3
CON ₂ (Invert)	\mathcal{S}_{LH}	90.6 \pm 1.0	93.0 \pm 0.4
CON ₂ (Flip)		91.5 \pm 0.6	92.9 \pm 0.5
CON ₂ (Equalize)		93.9 \pm 0.3	94.5 \pm 0.2
CON ₂ (Invert)	\mathcal{S}_{NND}	91.1 \pm 0.7	94.1 \pm 0.4
CON ₂ (Flip)		92.8 \pm 1.1	93.4 \pm 1.1

We train CON₂ on the healthy samples of a real-world medical chest x-ray dataset (Kermayn et al., 2018) and a melanoma imaging dataset (Javid, 2022), discriminating between unseen healthy and anomalous samples at test time. Here, we model invariances of normal samples with the three context augmentations *Flip*, *Invert*, and *Equalize* mentioned in Section 3.1. We run each experiment across three seeds, train on healthy samples, and apply our anomaly score functions to the representations of test samples to detect anomalies. We report the mean and standard deviation of the resulting area under the receiver operating characteristics curves (AUROC) in Table 1.

We can see that our CON₂ consistently exhibits a strong performance with both \mathcal{S}_{LH} and \mathcal{S}_{NND} across all three context augmentations. However, we note a significant performance decrease with *Flip* on

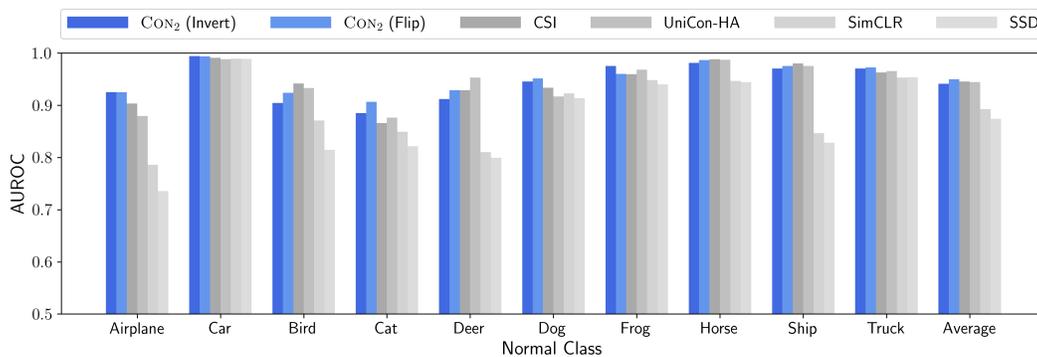


Figure 4: AUROCs of CIFAR10 when setting one class as normal and detecting the rest as anomalous. We compare CON₂ with the *Invert* and *Flip* context augmentations with \mathcal{S}_{NND} to other contrastive anomaly detection methods. Both the *Invert* and *Flip* context augmentations fulfill our assumptions, resulting in good performances across all classes. Our method further outperforms our baselines in most classes. CON₂ with *Flip* has the highest average across all methods considered.

the Melanoma dataset. This performance decrease most likely stems from the fact that *Flip* violates distinctiveness on melanoma images as they could be taken from any angle. Apart from CON₂ (Flip) on Melanoma, our method outperforms all baselines, confirming that modeling invariances of normal data offers a clear advantage in specialized settings. We further note that the CLIP-AD method, which exhibits impressive performance on natural image datasets (see Appendix E.2), lacks behind most of our baselines, indicating that, even in the age of foundation models, unsupervised anomaly detection methods are still important in specialized domains.

4.2 NATURAL IMAGE BENCHMARKS

In addition to the results on the more specialized medical imaging domain, our method also exhibits robust performance on more traditional natural imaging benchmark datasets. In this experiment, we train CON₂ on the CIFAR10, CIFAR100 (Krizhevsky et al., 2009), ImageNet30 (Russakovsky et al., 2015; Hendrycks et al., 2019b), Dogs vs. Cats (Cukierski, 2013), and Muffin vs. Chihuahua (Cortinhas, 2023) datasets in the one-class classification setting (Ruff et al., 2021). In the one-class classification setting, we typically work on multi-class classification datasets where we consider one of the classes as the normal class and the rest as anomalies. In particular, we train our model on the training samples of the normal class and want to differentiate between unseen samples of this normal class and all other classes at test time. Here, we train each model across three seeds for each class of each dataset, reporting the mean and standard deviation of the resulting AUROCs.

On natural images, the *Equalize* context augmentation does not satisfy distinctiveness, as this transformation often results in scenes that seem slightly differently illuminated (see Figure 2 for some examples). We thus only present results of CON₂ with *Flip* and *Invert* context augmentations. In Section 4.1, we saw that the more efficient \mathcal{S}_{LH} anomaly score exhibits relatively strong performance, however, \mathcal{S}_{NND} typically performs slightly better and we thus only report \mathcal{S}_{NND} in this section. Further, we note that CLIP-AD exhibits a strong performance on natural image datasets as, during training, CLIP has been exposed to samples that are similar to anomalies of this one-class classification setting. It is thus hard to compare CLIP-AD to our method and baselines, which were all trained without outlier exposure, on natural image datasets and we thus do not compare to CLIP-AD in this section. For completeness, we report the full results including including both scores, CON₂ with *Equalize*, and CLIP-AD results for all datasets in Appendix E.2.

In Figure 4, we compare the performance of CON₂ and our baselines across the different one-class settings of CIFAR10. CON₂ outperforms our baselines on almost all classes, where CON₂ (Flip) with an average AUROC of 95.3 performs better than CON₂ (Invert), which exhibits an average AUROC of 94.6. We suspect that *Invert* exhibits similar issues as *Equalize* in some instances, i.e., it may not always fully satisfy the distinctiveness assumption.

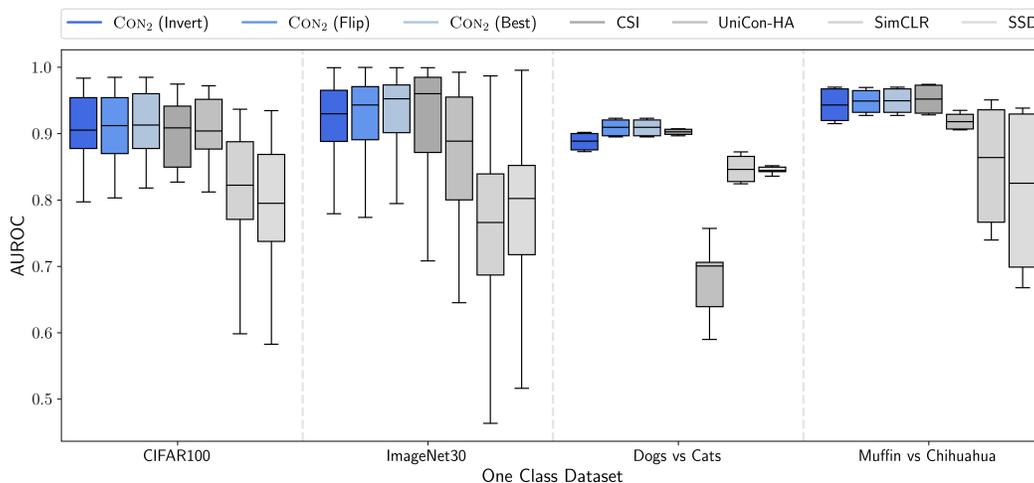


Figure 5: One class classification results for CIFAR100, ImageNet30, Dogs vs. Cats, and Muffin vs. Chihuahua. Our method consistently outperforms our baselines on CIFAR100 and Dogs vs. Cats while exhibiting more robust performance across different normal classes with a similar average performance to CSI on ImageNet30 and Muffin vs. Chihuahua. Additionally, we provide results including CON₂ (Best), which demonstrates how carefully selecting context augmentations satisfying the assumptions of Section 3.1 further improves the capabilities of anomaly detection with CON₂.

We further provide results on one-class CIFAR100, ImageNet30, Dogs vs. Cats, and Muffin vs. Chihuahua in Figure 5. There, in addition to the *Invert* and *Flip* context augmentation, we also provide results for CON₂ (Best), which selects the context augmentation individually for each class, depending on which satisfies alignment and distinctiveness better for the current normal class. Our method compares well against established baselines on natural images, matching or improving the state-of-the-art. Similar to what we saw on CIFAR10, CON₂ displays a robust performance across the board. Our approach outperforms baselines on CIFAR100 and Dogs vs. Cats while matching the performance on ImageNet30 and Muffin vs. Chihuahua while exhibiting much more consistent performance across different normal classes as can be seen from the much lower variance displayed in Figure 5. We can also see that selecting the context augmentation that best fits the normal class can improve the performance. However, we also achieve strong performance if the context augmentation violates alignment and distinctiveness on only some samples of the dataset. We provide further ablations, including experiments on applying multiple context augmentations, demonstrating that a single context augmentation is sufficient, and additional one-class classification results in Appendix E.

5 CONCLUSION

In this work, we presented a novel approach to anomaly detection, focusing on learning representations of normality by leveraging prior knowledge about invariances in the normal data rather than simulating anomalous data as in previous works. Employing knowledge about the invariances of normal data is more realistic and provides a stronger foundation for anomaly detection, particularly in specialized domains such as healthcare, where anomalous data is rare or hard to simulate accurately.

CON₂ learns dense, highly informative context clusters that capture the properties of normal data. These clusters provide rich representations and ensure that a sample’s relative positioning is consistent across clusters, strengthening the model’s ability to differentiate between normal and anomalous data. This results in a more structured representation space, making our approach well-suited for anomaly detection tasks with our anomaly score functions.

We demonstrated the efficacy of our approach on two real-world medical imaging datasets, where our method achieved impressive results. This highlights the applicability of CON₂ in safety-critical applications where robust anomaly detection is essential. Additionally, our approach exhibited strong performance on natural imaging datasets, consistently outperforming baseline methods,

486 demonstrating its versatility across different domains. Further, our method highlights the impor-
487 tance of domain-specific approaches in specialized fields like healthcare, where tailored models can
488 outperform foundation model-based approaches such as CLIP-AD, despite their success in more
489 general settings.

490 In conclusion, CON₂ represents a significant advancement in anomaly detection by learning struc-
491 tured representations of normal data without relying on anomalous data. This approach is partic-
492 ularly valuable in specialized, high-stakes settings, offering robust and effective solutions across
493 various application domains.

494 REFERENCES

495
496
497 Jinwon An and Sungzoon Cho. Variational autoencoder based anomaly detection using reconstruc-
498 tion probability. *Special lecture on IE*, 2(1):1–18, 2015. [2](#), [15](#)

499
500 Jason Ansel, Edward Yang, Horace He, Natalia Gimelshein, Animesh Jain, Michael Voznesensky,
501 Bin Bao, Peter Bell, David Berard, Evgeni Burovski, Geeta Chauhan, Anjali Chourdia, Will
502 Constable, Alban Desmaison, Zachary DeVito, Elias Ellison, Will Feng, Jiong Gong, Michael
503 Gschwind, Brian Hirsh, Sherlock Huang, Kshiteej Kalambarkar, Laurent Kirsch, Michael Lazos,
504 Mario Lezcano, Yanbo Liang, Jason Liang, Yinghai Lu, C. K. Luk, Bert Maher, Yunjie Pan, Chris-
505 tian Puhersch, Matthias Reso, Mark Saroufim, Marcos Yukio Siraichi, Helen Suk, Shunting Zhang,
506 Michael Suo, Phil Tillet, Xu Zhao, Eikan Wang, Keren Zhou, Richard Zou, Xiaodong Wang, Ajit
507 Mathews, William Wen, Gregory Chanan, Peng Wu, and Soumith Chintala. PyTorch 2: Faster
508 Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation.
509 In *Proceedings of the 29th ACM International Conference on Architectural Support for Program-*
510 *ming Languages and Operating Systems, Volume 2, volume 2 of ASPLOS '24*, pp. 929–947, New
511 York, NY, USA, April 2024. Association for Computing Machinery. ISBN 9798400703850. doi:
10.1145/3620665.3640366. [16](#)

512 Deegan J Atha and Mohammad R Jahanshahi. Evaluation of deep learning approaches based on
513 convolutional neural networks for corrosion detection. *Structural Health Monitoring*, 17(5):1110–
514 1128, September 2018. ISSN 1475-9217. doi: 10.1177/1475921717737051. [1](#)

515
516 Liron Bergman and Yedid Hoshen. Classification-Based Anomaly Detection for General Data. In
517 *International Conference on Learning Representations*, 2019. [2](#)

518 Liron Bergman, Niv Cohen, and Yedid Hoshen. Deep nearest neighbor anomaly detection. *arXiv*
519 *preprint arXiv:2002.10445*, 2020. [3](#), [6](#)

520
521 C. M. Bishop. Novelty detection and neural network validation. *IEEE Proceedings - Vision, Image*
522 *and Signal Processing*, 141(4):217–222, August 1994. ISSN 1359-7108. doi: 10.1049/ip-vis:
523 19941330. [16](#)

524 Markus M. Breunig, Hans-Peter Kriegel, Raymond T. Ng, and Jörg Sander. LOF: identifying
525 density-based local outliers. In *Proceedings of the 2000 ACM SIGMOD international conference*
526 *on Management of data, SIGMOD '00*, pp. 93–104, New York, NY, USA, May 2000. Association
527 for Computing Machinery. ISBN 978-1-58113-217-5. doi: 10.1145/342009.335388. [16](#)

528
529 Jinghui Chen, Saket Sathe, Charu Aggarwal, and Deepak Turaga. Outlier Detection with Autoen-
530 coder Ensembles. In *Proceedings of the 2017 SIAM International Conference on Data Mining*
531 *(SDM)*, Proceedings, pp. 90–98. Society for Industrial and Applied Mathematics, June 2017. [2](#)

532 Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A Simple Framework
533 for Contrastive Learning of Visual Representations. *37th International Conference on Machine*
534 *Learning, ICML 2020, Part F168147-3:1575–1585*, February 2020. doi: 10.48550/arxiv.2002.
535 05709. [4](#), [5](#), [15](#), [16](#), [18](#)

536
537 Matan Jacob Cohen and Shai Avidan. Transformally-two (feature spaces) are better than one. In *Pro-*
538 *ceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4060–
539 4069, 2022. [3](#)

Samuel Cortinhas. Muffin vs. Chihuahua, 2023. [8](#), [18](#)

540 Will Cukierski. Dogs vs. Cats, 2013. [8](#), [17](#)

541

542 Imant Daunhawer, Alice Bizeul, Emanuele Palumbo, Alexander Marx, and Julia E Vogt. Identifi-
543 ability results for multimodal contrastive learning. In *The Eleventh International Conference on*
544 *Learning Representations*, 2023. [15](#)

545 William Falcon and The PyTorch Lightning team. PyTorch Lightning, March 2019. [16](#)

546

547 Izhak Golan and Ran El-Yaniv. Deep Anomaly Detection Using Geometric Transformations. In
548 *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
549 [2](#)

550 Koosha Golmohammadi and Osmar R. Zaiane. Time series contextual anomaly detection for detect-
551 ing market manipulation in stock market. In *2015 IEEE International Conference on Data Science*
552 *and Advanced Analytics (DSAA)*, pp. 1–10, October 2015. doi: 10.1109/DSAA.2015.7344856. [1](#)

553 Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen,
554 David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern,
555 Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime
556 Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler
557 Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Ar-
558 ray programming with NumPy. *Nature*, 585(7825):357–362, September 2020. doi: 10.1038/
559 s41586-020-2649-2. [16](#)

560

561 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-
562 nition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp.
563 770–778, 2016. [7](#)

564 Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution
565 examples in neural networks. In *International Conference on Learning Representations*, 2017. [3](#)

566 Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier
567 exposure. In *International Conference on Learning Representations*, 2019a. [2](#)

568 Dan Hendrycks, Mantas Mazeika, Saurav Kadavath, and Dawn Song. Using self-supervised learn-
569 ing can improve model robustness and uncertainty. *Advances in neural information processing*
570 *systems*, 32, 2019b. [2](#), [8](#), [17](#)

571

572 Muhammad Hasnain Javid. Melanoma skin cancer dataset of 10000 images, 2022. [7](#), [17](#)

573

574 Jongheon Jeong, Yang Zou, Taewan Kim, Dongqing Zhang, Avinash Ravichandran, and Onkar
575 Dabeer. Winclip: Zero-/few-shot anomaly classification and segmentation. In *Proceedings of the*
576 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19606–19616, 2023. [3](#)

577 Daniel S. Kermany, Michael Goldbaum, Wenjia Cai, Carolina C. S. Valentim, Huiying Liang,
578 Sally L. Baxter, Alex McKeown, Ge Yang, Xiaokang Wu, Fangbing Yan, Justin Dong, Made K.
579 Prasadha, Jacqueline Pei, Magdalene Y. L. Ting, Jie Zhu, Christina Li, Sierra Hewett, Jason
580 Dong, Ian Ziyar, Alexander Shi, Runze Zhang, Lianghong Zheng, Rui Hou, William Shi, Xin Fu,
581 Yaou Duan, Viet A. N. Huu, Cindy Wen, Edward D. Zhang, Charlotte L. Zhang, Oulan Li, Xi-
582 aobo Wang, Michael A. Singer, Xiaodong Sun, Jie Xu, Ali Tafreshi, M. Anthony Lewis, Huimin
583 Xia, and Kang Zhang. Identifying Medical Diagnoses and Treatable Diseases by Image-Based
584 Deep Learning. *Cell*, 172(5):1122–1131.e9, February 2018. ISSN 0092-8674, 1097-4172. doi:
585 10.1016/j.cell.2018.02.010. [7](#), [17](#)

586 Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron
587 Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in neural*
588 *information processing systems*, 33:18661–18673, 2020. [4](#), [5](#), [15](#), [16](#)

589 Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images,
590 2009. [8](#), [17](#)

591

592 Kimin Lee, Honglak Lee, Kibok Lee, and Jinwoo Shin. Training confidence-calibrated classifiers for
593 detecting out-of-distribution samples. In *International Conference on Learning Representations*,
2018. [3](#)

-
- 594 Jingyao Li, Pengguang Chen, Zexin He, Shaozuo Yu, Shu Liu, and Jiaya Jia. Rethinking out-of-
595 distribution (ood) detection: Masked image modeling is all you need. In *Proceedings of the*
596 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 11578–11589, 2023. 3
597
- 598 Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation Forest. In *2008 Eighth IEEE Interna-*
599 *tional Conference on Data Mining*, pp. 413–422, December 2008. doi: 10.1109/ICDM.2008.17.
600 16
- 601 Philipp Liznerski, Lukas Ruff, Robert A. Vandermeulen, Billy Joe Franks, Klaus Robert Muller,
602 and Marius Kloft. Exposing Outlier Exposure: What Can Be Learned From Few, One, and Zero
603 Outlier Images. *Transactions on Machine Learning Research*, August 2022. ISSN 2835-8856. 3,
604 7
- 605 Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. In *Inter-*
606 *national Conference on Learning Representations*, 2017. 18
607
- 608 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Confer-*
609 *ence on Learning Representations*, 2019. 18
- 610 Wes McKinney. Data Structures for Statistical Computing in Python. In Stéfan van der Walt and
611 Jarrod Millman (eds.), *Proceedings of the 9th Python in Science Conference*, pp. 56 – 61, 2010.
612 doi: 10.25080/Majora-92bf1922-00a. 16
613
- 614 Hossein Mirzaei, Mohammadreza Salehi, Sajjad Shahabi, Efstratios Gavves, Cees GM Snoek, Mo-
615 hammad Sabokrou, and Mohammad Hossein Rohban. Fake It Until You Make It: Towards Accu-
616 rate Near-Distribution Novelty Detection. In *The Eleventh International Conference on Learning*
617 *Representations*, 2022. 2
- 618 Benjamin Nachman and David Shih. Anomaly detection with density estimation. *Physical Review*
619 *D*, 101(7):075042, April 2020. doi: 10.1103/PhysRevD.101.075042. 2
620
- 621 Eric Nalisnick, Akihiro Matsukawa, Yee Whye Teh, Dilan Gorur, and Balaji Lakshminarayanan.
622 Do Deep Generative Models Know What They Don’t Know? *7th International Conference on*
623 *Learning Representations, ICLR 2019*, October 2018. doi: 10.48550/arxiv.1810.09136. 2, 16
- 624 Poojan Oza and Vishal M. Patel. One-class convolutional neural network. *IEEE Signal Processing*
625 *Letters*, 26(2):277–281, 2018. 1
626
- 627 Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier
628 Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, and others. Scikit-
629 learn: Machine learning in Python. *Journal of machine learning research*, 12(Oct):2825–2830,
630 2011. 16
- 631 B. Peirce. *An Elementary Treatise on Curves, Functions, and Forces*. Number Bd. 1 in An Ele-
632 mentary Treatise on Curves, Functions, and Forces. J. Munroe, 1852. URL [https://books.](https://books.google.ch/books?id=wVhtAAAAMAAJ)
633 [google.ch/books?id=wVhtAAAAMAAJ](https://books.google.ch/books?id=wVhtAAAAMAAJ). 6
634
- 635 Chen Qiu, Aodong Li, Marius Kloft, Maja Rudolph, and Stephan Mandt. Latent outlier exposure for
636 anomaly detection with contaminated data. In *International Conference on Machine Learning*,
637 pp. 18153–18167. PMLR, 2022. 3
- 638 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
639 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
640 models from natural language supervision. In *International conference on machine learning*, pp.
641 8748–8763. PMLR, 2021. 7
- 642 Tal Reiss and Yedid Hoshen. Mean-Shifted Contrastive Loss for Anomaly Detection. *Proceedings*
643 *of the AAAI Conference on Artificial Intelligence*, 37(2):2155–2162, June 2023. ISSN 2374-3468.
644 doi: 10.1609/aaai.v37i2.25309. 3
645
- 646 Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexan-
647 der Binder, Emmanuel Müller, and Marius Kloft. Deep one-class classification. In *International*
conference on machine learning, pp. 4393–4402. PMLR, 2018. 1, 2

648 Lukas Ruff, Robert A. Vandermeulen, Nico Görnitz, Alexander Binder, Emmanuel Müller, Klaus-
649 Robert Müller, and Marius Kloft. Deep semi-supervised anomaly detection. In *International*
650 *Conference on Learning Representations*, 2020. [3](#)
651

652 Lukas Ruff, Jacob R. Kauffmann, Robert A. Vandermeulen, Grégoire Montavon, Wojciech Samek,
653 Marius Kloft, Thomas G. Dietterich, and Klaus-Robert Müller. A unifying review of deep and
654 shallow anomaly detection. *Proceedings of the IEEE*, 109(5):756–795, 2021. [1](#), [6](#), [8](#), [15](#), [18](#)
655

656 Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng
657 Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-
658 Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Com-*
659 *puter Vision*, 115(3):211–252, December 2015. ISSN 0920-5691, 1573-1405. doi: 10.1007/
660 s11263-015-0816-y. [8](#), [17](#)

661 Alain Ryser, Laura Manduchi, Fabian Laumer, Holger Michel, Sven Wellmann, and Julia E. Vogt.
662 Anomaly Detection in Echocardiograms with Dynamic Variational Trajectory Models. In *Pro-*
663 *ceedings of the 7th Machine Learning for Healthcare Conference*, pp. 425–458. PMLR, Decem-
664 ber 2022. [1](#)

665 Mohammad Sabokrou, Mahmood Fathy, Guoying Zhao, and Ehsan Adeli. Deep end-to-end one-
666 class classifier. *IEEE transactions on neural networks and learning systems*, 32(2):675–684,
667 2020. [1](#)
668

669 Thomas Schlegl, Philipp Seeböck, Sebastian M. Waldstein, Ursula Schmidt-Erfurth, and Georg
670 Langs. Unsupervised Anomaly Detection with Generative Adversarial Networks to Guide Marker
671 Discovery. In Marc Niethammer, Martin Styner, Stephen Aylward, Hongtu Zhu, Ipek Oguz,
672 Pew-Thian Yap, and Dinggang Shen (eds.), *Information Processing in Medical Imaging*, pp.
673 146–157, Cham, 2017. Springer International Publishing. ISBN 978-3-319-59050-9. doi:
674 10.1007/978-3-319-59050-9_12. [1](#)

675 Thomas Schlegl, Philipp Seeböck, Sebastian M. Waldstein, Georg Langs, and Ursula Schmidt-
676 Erfurth. f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks.
677 *Medical Image Analysis*, 54:30–44, May 2019. ISSN 1361-8415. doi: 10.1016/j.media.2019.01.
678 010. [2](#), [15](#)

679 Bernhard Schölkopf, John C. Platt, John Shawe-Taylor, Alex J. Smola, and Robert C. Williamson.
680 Estimating the support of a high-dimensional distribution. *Neural computation*, 13(7):1443–1471,
681 2001. [16](#)
682

683 Vikash Sehwal, Mung Chiang, and Prateek Mittal. SSD: A unified framework for self-supervised
684 outlier detection. In *9th International Conference on Learning Representations, ICLR 2021, Vir-*
685 *tual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. [2](#), [7](#), [16](#)
686

687 Kihyuk Sohn. Improved Deep Metric Learning with Multi-class N-pair Loss Objective. In *Advances*
688 *in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. [15](#)
689

690 Yiyun Sun, Yifei Ming, Xiaojin Zhu, and Yixuan Li. Out-of-Distribution Detection with Deep
691 Nearest Neighbors. In *Proceedings of the 39th International Conference on Machine Learning*,
692 pp. 20827–20840. PMLR, June 2022. [2](#), [6](#), [7](#), [16](#)

693 Jihoon Tack, Sangwoo Mo, Jongheon Jeong, and Jinwoo Shin. CSI: Novelty Detection via Con-
694 trastive Learning on Distributionally Shifted Instances. In *Advances in Neural Information Pro-*
695 *cessing Systems*, volume 33, pp. 11839–11852. Curran Associates, Inc., 2020. [1](#), [6](#), [7](#), [15](#), [16](#)
696

697 David M.J. Tax and Robert P.W. Duin. Support Vector Data Description. *Machine Learning*, 54(1):
698 45–66, January 2004. ISSN 1573-0565. doi: 10.1023/B:MACH.0000008084.60811.49. [16](#)
699

700 The pandas development team. pandas-dev/pandas: Pandas, February 2020. [16](#)
701

701 Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predic-
702 tive coding, 2019. [4](#), [15](#)

702 Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau,
703 Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der
704 Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson,
705 Eric Jones, Robert Kern, Eric Larson, C J Carey, Ilhan Polat, Yu Feng, Eric W. Moore,
706 Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero,
707 Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt,
708 and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing
709 in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2. 16

710 Julius von Kügelgen, Yash Sharma, Luigi Gresele, Wieland Brendel, Bernhard Schölkopf, Michel
711 Besserve, and Francesco Locatello. Self-Supervised Learning with Data Augmentations Provably
712 Isolates Content from Style. In *Advances in Neural Information Processing Systems*, volume 34,
713 pp. 16451–16467. Curran Associates, Inc., 2021. 15

714 Guodong Wang, Yunhong Wang, Jie Qin, Dongming Zhang, Xiuguo Bao, and Di Huang. Unilaterally
715 aggregated contrastive learning with hierarchical augmentation for anomaly detection. In
716 *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*,
717 pp. 6865–6874. IEEE, 2023. doi: 10.1109/ICCV51070.2023.00634. 1, 6, 7, 15, 16

718 Haoqi Wang, Zhizhong Li, Litong Feng, and Wayne Zhang. Vim: Out-of-distribution with virtual-
719 logit matching. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pp. 4911–4920. IEEE, 2022. doi: 10.1109/CVPR52688.2022.00487. 3

720 Zhirong Wu, Yuanjun Xiong, Stella X. Yu, and Dahua Lin. Unsupervised feature learning via non-
721 parametric instance discrimination. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pp. 3733–3742. Computer
722 Vision Foundation / IEEE Computer Society, 2018. doi: 10.1109/CVPR.2018.00393. 15

723 Yang Xin, Lingshuang Kong, Zhi Liu, Yuling Chen, Yanmiao Li, Hongliang Zhu, Mingcheng Gao,
724 Haixia Hou, and Chunhua Wang. Machine Learning and Deep Learning Methods for Cybersecurity. *IEEE Access*, 6:35365–35381, 2018. ISSN 2169-3536. doi: 10.1109/ACCESS.2018.2836950. 1

725 Suhang You, Kerem C. Tezcan, Xiaoran Chen, and Ender Konukoglu. Unsupervised Lesion Detection
726 via Image Restoration with a Normative Prior. In *Proceedings of The 2nd International Conference on Medical Imaging with Deep Learning*, pp. 540–556. PMLR, May 2019. 2

727 Rui Zhao, Ruqiang Yan, Zhenghua Chen, Kezhi Mao, Peng Wang, and Robert X. Gao. Deep learning
728 and its applications to machine health monitoring. *Mechanical Systems and Signal Processing*, 115:213–237, January 2019. ISSN 0888-3270. doi: 10.1016/j.ymssp.2018.05.050. 1

729 Qihang Zhou, Guansong Pang, Yu Tian, Shibo He, and Jiming Chen. AnomalyCLIP: Object-
730 agnostic prompt learning for zero-shot anomaly detection. In *The Twelfth International Conference on Learning Representations*, 2024. 3

731 Bo Zong, Qi Song, Martin Renqiang Min, Wei Cheng, Cristian Lumezanu, Dae-ki Cho, and Haifeng
732 Chen. Deep autoencoding gaussian mixture model for unsupervised anomaly detection. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. 2

733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

756 A BACKGROUND

757
758 In this section, we provide some terminology for contrastive learning and background about the
759 anomaly detection setting.
760

761 A.1 CONTRASTIVE LEARNING

762
763 Recently, contrastive learning has emerged as a popular approach for representation learning
764 (van den Oord et al., 2019; Chen et al., 2020). By design, contrastive learning has the capability to
765 learn representations that are agnostic to certain invariances (von Kügelgen et al., 2021; Daunhawer
766 et al., 2023), which makes contrastive learning a particularly interesting choice to learn informative
767 representations of normal samples (Tack et al., 2020; Wang et al., 2023), as it allows us to incor-
768 porate prior knowledge about our data into the representing learning process in the form of data
769 augmentations. More specifically, invariances are learned by forming positive and negative pairs
770 over the training dataset by applying data augmentations that should retain the relevant content of a
771 sample.

772 The goal of contrastive learning is to learn an encoding function $g_\theta(\mathbf{x})$, where representations of
773 positive pairs of samples are close and negative pairs are far from each other. For a given pair of
774 samples $\mathbf{x}, \mathbf{x}' \in X$, we can define the instance discrimination loss as (Sohn, 2016; Wu et al., 2018;
775 van den Oord et al., 2019)

$$776 \ell(\mathbf{x}, \mathbf{x}', X) = -\log \frac{\exp(\text{sim}(\mathbf{x}, \mathbf{x}')/\tau)}{\sum_{\mathbf{x}'' \in X: \mathbf{x}'' \neq \mathbf{x}} \exp(\text{sim}(\mathbf{x}, \mathbf{x}'')/\tau)} .$$

779 As mentioned in Section 3.2, we consider the function $\text{sim}(\mathbf{x}, \mathbf{x}')$ to correspond to the cosine sim-
780 ilarity between the two input vectors, as this is one of the most popular choices in the contrastive
781 learning literature.

782 One of the most prominent contrastive methods is SimCLR (Chen et al., 2020), which creates posi-
783 tive pairs through sample augmentations. There exists a supervised extension called SupCon (Khosla
784 et al., 2020), which incorporates class labels into the SimCLR loss. For a given set of augmenta-
785 tions T , a dataset $X = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$, and an augmented dataset \tilde{X} where $|\tilde{X}| = 2N$ and $(\tilde{\mathbf{x}}_{2i}, y_i)$,
786 $(\tilde{\mathbf{x}}_{2i+1}, y_i) \in \tilde{X}$ denote two transformations of the same sample using random augmentations from
787 T , SimCLR and SupCon introduce the following loss functions:

$$788 \mathcal{L}_{\text{SimCLR}}(\tilde{X}) = \frac{1}{2N} \sum_{i=1}^N (\ell(f_\Theta(\tilde{\mathbf{x}}_{2i}), f_\Theta(\tilde{\mathbf{x}}_{2i+1}), f_\Theta(\tilde{X})) + \ell(f_\Theta(\tilde{\mathbf{x}}_{2i+1}), f_\Theta(\tilde{\mathbf{x}}_{2i}), f_\Theta(\tilde{X}))) ,$$

$$789 \mathcal{L}_{\text{SupCon}}(\tilde{X}) = \sum_{(\tilde{\mathbf{x}}_i, y_i) \in \tilde{X}} \frac{1}{N(y_i) - 1} \sum_{\substack{(\tilde{\mathbf{x}}_j, y_j) \in \tilde{X}: \\ \tilde{\mathbf{x}}_j \neq \tilde{\mathbf{x}}_i \wedge y_i = y_j}} \ell(f_\Theta(\tilde{\mathbf{x}}_i), f_\Theta(\tilde{\mathbf{x}}_j), f_\Theta(\tilde{X})) .$$

794 Here, we denote $f_\Theta(\mathbf{x}) = h_{\theta'}(g_\theta(\mathbf{x}))$, where $\mathbf{z} = g_\theta(\mathbf{x})$ is a feature extractor and $h_{\theta'}(\mathbf{z})$ is a
795 projection head that is typically only used during training (Chen et al., 2020). Further, we define
796 $f_\Theta(\tilde{X}) = \{f_\Theta(\tilde{\mathbf{x}}) \mid (\tilde{\mathbf{x}}, y) \in \tilde{X}\}$ and $N(y) = |\{(\tilde{\mathbf{x}}_i, y_i) \mid (\tilde{\mathbf{x}}_i, y_i) \in \tilde{X} \wedge y_i = y\}|$ is the number of
797 samples in \tilde{X} with label y .
798

800 A.2 ANOMALY DETECTION

801 In the anomaly detection setting, we are given an unlabeled dataset $\{\mathbf{x}_1, \dots, \mathbf{x}_n\} = X \subset \mathcal{X}$, while
802 assuming that most samples are normal, i.e., the dataset is practically free of outliers (Ruff et al.,
803 2021). The goal is to learn a model from the given dataset that discriminates between normal and
804 anomalous data at test time.

805 In this work, we assume the challenging case where our dataset is completely free of anomalies.
806 Hence, we aim to discriminate between the normal class and a completely unobserved set of anoma-
807 lies at test time. This setting is sometimes called one-class classification or novelty detection.
808

809 To achieve this goal, one straightforward approach is to approximate the distribution $p_{\mathcal{X}}(\mathbf{x})$ directly
using generative models (An & Cho, 2015; Schlegl et al., 2019). Because we assume normal data to

Table 2: Average compute hours for the experiments for each dataset and method per run. SimCLR and SSD use the same representations, so we can evaluate both methods in one go and list their compute hours together.

Method \ Dataset	CIFAR10	CIFAR100	ImageNet30	Dogs vs. Cats	Muffin vs. Chihuahua	Pneumonia	Melanoma
SimCLR/SSD	4	2	3	9	3	3	5
CSI	8	4	4	14	5	8	6
UniCon-HA	16	16	8	18	7	12	18
CON ₂	5	3	4	11	4	5	6

lie in high-density regions of $p_{\mathcal{X}}$, we can discriminate between normal and anomalous samples by applying a threshold function $p_{\mathcal{X}}(\mathbf{x}) \leq \tau$, where $\tau \in \mathbb{R}$ is an often task-specific threshold (Bishop, 1994). As density-based approaches are often difficult to apply to high-dimensional data directly (Nalisnick et al., 2018), we follow a slightly different line of work.

In this paper, we focus on learning a function $g_{\theta} : \mathcal{X} \rightarrow \mathcal{Z}$ that provides us with representations that capture the normal attributes of samples in the dataset (Sehwag et al., 2021; Tack et al., 2020; Wang et al., 2023), by mapping normal samples close to each other in representation space. On the other hand, anomalies that lack the learned normal structure should be mapped to a different part of the representation space.

Given $g_{\theta}(\mathbf{x})$, a popular approach to detect anomalies is by defining a scoring function $\mathcal{S} : \mathcal{Z} \rightarrow \mathbb{R}$ (Breunig et al., 2000; Schölkopf et al., 2001; Tax & Duin, 2004; Liu et al., 2008). The score function maps a representation onto a metric that estimates the anomalousness of a sample. To identify anomalies at test time, we can use \mathcal{S} similarly to the density $p_{\mathcal{X}}$, i.e., we consider a new sample \mathbf{x} to be normal if $\mathcal{S}(g_{\theta}(\mathbf{x})) \leq \tau$, whereas $\mathcal{S}(g_{\theta}(\mathbf{x})) > \tau$ means \mathbf{x} is an anomaly.

B COMPUTE & CODE

We run all our experiments on single GPUs on a compute cluster using a combination of RTX2080Ti, RTX3090, and RTX4090 GPUs. Each experiment can be run with 4 CPU workers and 16 GB of memory. We provide an overview of the compute for our experiments in Table 2. Our experiments are written using PyTorch (Ansel et al., 2024) with Lightning (Falcon & The PyTorch Lightning team, 2019).

In the following, we list for each of our methods and baselines how we arrive at results and which code we use.

CON₂: We implement CON₂ using PyTorch (Ansel et al., 2024) together with Lightning (Falcon & The PyTorch Lightning team, 2019). To evaluate our method, we use various open-source Python libraries such as NumPy (Harris et al., 2020), scikit-learn (Pedregosa et al., 2011), Pandas (McKinney, 2010; team, 2020), or SciPy (Virtanen et al., 2020). Parts of the implementation of the CON₂ objective are based on code provided by Khosla et al. (2020) (<https://github.com/HobbitLong/SupContrast>).

SimCLR: For this baseline, we implement SimCLR (Chen et al., 2020) and compute anomaly scores in a similar fashion as (Sun et al., 2022). For this baseline, we rely on similar packages as CON₂.

SSD: We use the same representations as for SimCLR but evaluate by following the procedure outlined in Sehwag et al. (2021).

CSI: To run experiments for CSI, we used the code provided in <https://github.com/alinelab/CSI>, implementing new dataloaders for the missing datasets.

UniCon-HA: We conducted experiments by running code provided by Wang et al. (2023) implementing new dataloaders for the missing datasets. We thank the authors for sharing their code with us.

864 C DATASETS

865
866 In the following, we provide details about preprocessing, sources, and licenses of the datasets we
867 use in our experiments.

869 PNEUMONIA

870
871 The Pneumonia dataset was originally published by [Kermany et al. \(2018\)](#) and consists of 5'863 lung
872 xrays, which are labeled with *Pneumonia* and *Normal* labels. We first resize images to 256 and apply
873 center-cropping to feed 224×224 images to our model. We ran all our experiments on the Pneumonia
874 dataset with a batch size of 128. The dataset can be downloaded from <https://www.kaggle.com/datasets/paultimothymooney/chest-xray-pneumonia> and is published under
875 *CC BY 4.0* license.
876
877

878 MELANOMA

879
880 We use the Melanoma dataset of [Javid \(2022\)](#), which consists of 10'600 images of
881 Melanoma labeled with being *benign* or *malignant*. We resize each image to size
882 128×128 before passing them to the model with batch size 128. The dataset is
883 publicly available at [https://www.kaggle.com/datasets/hasnainjaved/](https://www.kaggle.com/datasets/hasnainjaved/melanoma-skin-cancer-dataset-of-10000-images)
884 [melanoma-skin-cancer-dataset-of-10000-images](https://www.kaggle.com/datasets/hasnainjaved/melanoma-skin-cancer-dataset-of-10000-images) and is published under the
885 *CC0: Public Domain* license.
886

887 CIFAR10/CIFAR100

888
889 CIFAR10 and CIFAR100 are natural image datasets with 32×32 samples. Both datasets consist
890 of a total of 60'000 samples, with a total of 10 and 100 samples for CIFAR10 and CIFAR100, re-
891 spectively. As CIFAR100 comes with only 600 samples per class, the dataset authors additionally
892 define a set of 20 superclasses, aggregating 5 labels each. In our one-class classification experi-
893 ments on CIFAR100 we use the superclasses to ensure a manageable number of runs and a sufficient
894 amount of training data. We ran all our experiments on CIFAR10 and CIFAR100 with a batch
895 size of 512. Both datasets were published by [Krizhevsky et al. \(2009\)](#) and can be downloaded
896 from <https://www.cs.toronto.edu/~kriz/cifar.html>. To the best of our knowl-
897 edge, these datasets come without a license.
898

899 IMAGENET30

900
901 The ImageNet30 dataset is a subset of the original ImageNet dataset ([Russakovsky et al., 2015](#)).
902 It was created by [Hendrycks et al. \(2019b\)](#) for the purpose of one-class classification. The dataset
903 consists of 42'000 natural images where each is labeled with one of 30 classes. We preprocess the
904 dataset by resizing the shorter edge to 256 pixels, from which we randomly crop a 224×224 image
905 patch every time we load an image for training. We ran all our experiments on ImageNet with a
906 batch size of 128. The dataset can be downloaded from [https://github.com/hendrycks/](https://github.com/hendrycks/ss-ood)
907 [ss-ood](https://github.com/hendrycks/ss-ood), which comes with the MIT License. Further, while we could not find a license for Im-
908 ageNet, terms of use are provided on <https://image-net.org/>.

909 DOGS VS. CATS

910
911 The Dogs vs. Cats was originally introduced in a Kaggle challenge by Microsoft Research ([Cukier-
912 ski, 2013](#)) and consists of 25'000 images of cats and dogs. We preprocess the dataset by resizing the
913 shorter edge to 128 pixels and then perform center cropping, feeding the resulting 128×128 image
914 to our model. We ran all our experiments on Dogs vs. Cats with a batch size of 256. The dataset
915 can be downloaded from [https://www.kaggle.com/competitions/dogs-vs-cats/](https://www.kaggle.com/competitions/dogs-vs-cats/data)
916 [data](https://www.kaggle.com/competitions/dogs-vs-cats/data). To the best of our knowledge, there is no official license for the dataset, but the Kaggle
917 page points to the Kaggle Competition rules [https://www.kaggle.com/competitions/](https://www.kaggle.com/competitions/dogs-vs-cats/rules)
[dogs-vs-cats/rules](https://www.kaggle.com/competitions/dogs-vs-cats/rules) in the license section.

918 CHIUAHUA VS. MUFFIN

919
920 The Chihuahua vs. Muffin dataset consists of 6'000 images scraped from Google Images. We
921 preprocess the dataset similar to ImageNet30, resizing the shorter edge of the images to 128 pix-
922 els while feeding random 128×128 sized image crops to the model during training. We ran all
923 our experiments on Chihuahua vs. Muffin with a batch size of 256. The dataset was published
924 by Cortinhas (2023) and can be downloaded from [https://www.kaggle.com/datasets/
925 samueltcortinhas/muffin-vs-chihuahua-image-classification/data](https://www.kaggle.com/datasets/samueltcortinhas/muffin-vs-chihuahua-image-classification/data). Ac-
926 cording to the datasets Kaggle page, the dataset is licensed under *CC0: Public Domain*.

927 In addition to the preprocessing mentioned above, we normalize each image with a mean and stan-
928 dard deviation of 0.5 after applying the augmentations of CON_2 .

930 D EXPERIMENTAL DETAILS

931
932 **Setting** We evaluate our method in the so-called one-class classification setting (Ruff et al., 2021).
933 More specifically, during training we assume to have access to only the normal (healthy) class. At
934 test time, the goal is to detect whether a new sample stems from the normal class seen during training
935 or whether it seems anomalous, i.e., deviates from the training distribution.

936
937 **Metrics** Typically, there is a high-class imbalance between normal and anomalous samples in the
938 one-class classification setting. Further, setting an appropriate threshold for the anomaly score is
939 often task-dependent. Therefore, a popular approach to evaluating the performance of anomaly
940 detection methods is to use the area under the receiver operator characteristic curve (AUROC) (Ruff
941 et al., 2021). This metric is threshold agnostic and robust to class imbalance.

942 For all our experiments, we report mean and standard deviation over three seeds per class of the
943 dataset. Note that the average results of a dataset are aggregated over different one-class classifica-
944 tion settings, one per class of the dataset.

945
946 **Hyperparameters** Similar to our method, all baselines make use of test-time augmentations. By
947 default, both CSI and UniCon-HA use 40 test time augmentations, which we adopt for all baselines.
948 In our experiments, we set the augmentation class \mathcal{T} to the set of augmentations introduced by Chen
949 et al. (2020). For the context augmentation, we experiment with vertical flips (Flip), inverting the
950 pixels of an image (Invert), i.e., $t_{\text{invert}}(\mathbf{x}_{ij}) = 1 - x_{ij}$, and histogram equalization (Equalize), see
951 Figure 2 for an illustration.

952 We choose hyperparameters for CON_2 based on their performance on the CIFAR10 dataset and keep
953 them constant across all experiments. We linearly anneal the hyperparameter α in $\mathcal{L}_{\text{CON}_2}$ from 0 to 1
954 over the course of training to encourage the model to first learn the context-specific cluster structure
955 while gradually aligning representations over the course of training. We optimize our loss using the
956 AdamW optimizer (Loshchilov & Hutter, 2019) with $\beta_1 = 0.9$, $\beta_2 = 0.999$, weight decay $\lambda = 0.001$,
957 and using a learning rate of 10^{-3} with a cosine annealing (Loshchilov & Hutter, 2017) schedule. We
958 run all experiments for 2048 epochs.

959 E ABLATIONS

960
961 In this section, we provide some additional experiments going beyond only two context clusters
962 (Appendix E.1) and a more detailed overview of the results on natural images (Section 4.2).

964 E.1 MULTIPLE CONTEXT AUGMENTATIONS

965
966 Our formulation in Section 3.1 can easily be extended beyond only one additional context by slightly
967 adjusting $\mathcal{L}_{\text{Context}}$. However, in addition to a loss in efficiency due to requiring more memory, we
968 did not find additional context augmentations to provide a performance benefit, as can be seen in
969 Figure 6. There, we ran an ablation with different numbers of context augmentations on different
970 classes of CIFAR10 and ImageNet30. In particular, we trained the adapted CON_2 loss for 2, 3, 4,
971 5, 6, 7, and 8 context augmentations, which we derived by combining *Flip*, *Invert*, and *Equalize*
from our previous experiments. Adding more augmentations does not seem to harm cases where we

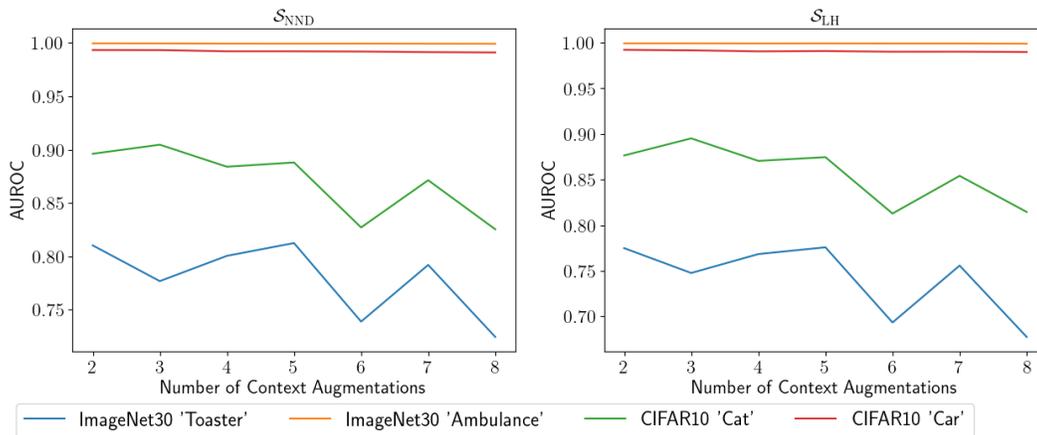


Figure 6: Ablation illustrating the effect of adding more context augmentations. While the performance of well-performing normal classes, such as ImageNet30 *Ambulance* or CIFAR10 *Car* stays consistent when adding more augmentations, we see a decrease for normal classes such as ImageNet30 *Toaster* or CIFAR10 *Cat* that already perform poor to begin with.

Table 3: One class classification results for CIFAR10, CIFAR100, ImageNet30, Dogs vs. Cats, and Muffin vs. Chihuahua. For each dataset, we train models over three different seeds per dataset class. We report mean and standard deviation over all the different one-class settings per dataset.

Method	Score	CIFAR10	CIFAR100	ImageNet30	Dogs vs. Cats	Muffin vs. Chihuahua
CLIP-AD (OE)	$\mathcal{S}_{\text{CLIP}}$	98.5 \pm 1.0	95.1 \pm 2.7	99.9 \pm 0.2	99.7 \pm 0.2	98.6 \pm 2.0
SimCLR	\mathcal{S}_{NND}	89.2 \pm 6.7	81.6 \pm 8.5	74.7 \pm 12.2	84.7 \pm 2.2	85.2 \pm 9.8
SSD	$\mathcal{S}_{\text{Mahalanobis}}$	87.4 \pm 8.1	79.2 \pm 9.4	76.8 \pm 13.0	84.5 \pm 0.6	81.3 \pm 13.1
CSI	\mathcal{S}_{CSI}	94.6 \pm 4.0	90.2 \pm 4.9	92.3 \pm 8.1	90.3 \pm 0.4	95.2 \pm 2.3
UniCon-HA	$\mathcal{S}_{\text{UniCon-HA}}$	94.4 \pm 4.0	90.9 \pm 4.4	85.5 \pm 12.0	67.9 \pm 6.2	91.9 \pm 1.3
CON ₂ (Equalize)	\mathcal{S}_{LH}	91.1 \pm 5.8	86.4 \pm 6.0	85.1 \pm 13.0	79.5 \pm 1.7	85.8 \pm 11.2
	\mathcal{S}_{NND}	91.5 \pm 5.6	87.8 \pm 4.8	86.2 \pm 12.1	83.2 \pm 1.2	88.3 \pm 8.3
CON ₂ (Invert)	\mathcal{S}_{LH}	93.7 \pm 4.3	89.7 \pm 5.2	90.2 \pm 9.4	87.9 \pm 0.6	91.7 \pm 4.3
	\mathcal{S}_{NND}	94.6 \pm 3.6	90.9 \pm 4.7	90.7 \pm 8.9	88.8 \pm 1.4	94.3 \pm 2.7
CON ₂ (Flip)	\mathcal{S}_{LH}	94.7 \pm 3.5	89.8 \pm 5.2	89.4 \pm 11.1	90.3 \pm 0.8	93.4 \pm 2.6
	\mathcal{S}_{NND}	95.3 \pm 2.9	90.8 \pm 4.8	90.5 \pm 10.4	90.9 \pm 1.4	94.9 \pm 2.0

experience good performance in the first place, however, we observe a diminishing performance for slightly more challenging classes.

E.2 ADDITIONAL ONE CLASS CLASSIFICATION RESULTS

In Section 4.2, we present figures with results of CON₂ on CIFAR10, CIFAR100, ImageNet30, Cats vs. Dogs, and Muffin vs. Chihuahua. For completeness, we also present a table containing the full results of CON₂ on all three context augmentations mentioned in Section 3.1 and both scores from Section 3.3 in Table 3. We further present results aggregated over individual one-class classification settings of CIFAR10, CIFAR100, ImageNet30, Cats vs. Dogs, and Muffin vs. Chihuahua for CON₂ on all datasets. We present results for CIFAR10 in Table 4, for all 20 superclasses of CIFAR100 in Table 5, for ImageNet30 in Table 6, for Dogs vs. Cats in Table 7, and for Muffin vs. Chihuahua in Table 8.

E.3 CONTRIBUTIONS OF INDIVIDUAL LOSS PARTS

We perform an ablation study where we evaluate the contribution of the loss components $\mathcal{L}_{\text{Content}}(\cdot)$ (see also Equation (3)) and $\mathcal{L}_{\text{Context}}(\cdot)$ (see also Equation (2)). We see that combining the content loss with the novel context loss leads to performance improvements on both evaluated datasets, Pneumonia and Melanoma.

1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

Table 4: AUROCS for each class of CIFAR10 for both of our scores when applying the Flip context augmentation. For each setting, we evaluated our method across three seeds.

Normal Class	CON ₂ (Equalize)		CON ₂ (Invert)		CON ₂ (Flip)	
	S _{LH}	S _{NNND}	S _{LH}	S _{NNND}	S _{LH}	S _{NNND}
0	89.7 \pm 0.4	91.3 \pm 0.6	90.3 \pm 0.5	92.5 \pm 0.3	90.3 \pm 0.9	92.5 \pm 0.7
1	98.4 \pm 0.0	98.5 \pm 0.2	99.2 \pm 0.1	99.4 \pm 0.0	99.3 \pm 0.0	99.4 \pm 0.0
2	87.0 \pm 0.2	88.7 \pm 0.5	88.3 \pm 0.4	90.5 \pm 0.1	91.6 \pm 0.6	92.4 \pm 0.6
3	77.7 \pm 1.0	78.5 \pm 2.1	86.2 \pm 0.5	88.5 \pm 0.4	88.7 \pm 0.3	90.7 \pm 0.3
4	90.2 \pm 0.8	90.1 \pm 0.6	90.8 \pm 0.3	91.2 \pm 0.5	93.1 \pm 0.5	92.9 \pm 1.1
5	88.5 \pm 0.3	87.7 \pm 0.6	93.9 \pm 0.2	94.5 \pm 0.3	94.7 \pm 0.2	95.1 \pm 0.1
6	96.4 \pm 0.3	96.0 \pm 0.7	97.3 \pm 0.2	97.5 \pm 0.2	96.5 \pm 0.1	96.0 \pm 0.3
7	95.6 \pm 0.3	95.9 \pm 0.3	97.6 \pm 0.1	98.1 \pm 0.1	98.6 \pm 0.0	98.7 \pm 0.0
8	93.2 \pm 0.3	94.4 \pm 0.2	96.3 \pm 0.1	97.1 \pm 0.1	97.0 \pm 0.3	97.5 \pm 0.2
9	94.1 \pm 0.3	93.5 \pm 0.1	96.7 \pm 0.3	97.0 \pm 0.3	97.1 \pm 0.1	97.3 \pm 0.1

Table 5: AUROCS for each superclass of CIFAR100 for both of our scores when applying the Flip context augmentation. For each setting, we evaluated our method across three seeds.

Normal Class	CON ₂ (Equalize)		CON ₂ (Invert)		CON ₂ (Flip)	
	S _{LH}	S _{NNND}	S _{LH}	S _{NNND}	S _{LH}	S _{NNND}
0	85.6 \pm 1.4	86.8 \pm 0.6	84.9 \pm 0.5	86.2 \pm 0.4	84.5 \pm 0.6	86.0 \pm 0.7
1	84.9 \pm 1.2	87.2 \pm 1.7	87.4 \pm 0.5	88.1 \pm 0.3	87.3 \pm 0.7	87.9 \pm 0.9
2	94.4 \pm 0.5	94.8 \pm 0.6	95.8 \pm 0.2	96.4 \pm 0.1	94.8 \pm 0.3	95.0 \pm 0.4
3	80.6 \pm 1.4	83.9 \pm 0.8	88.4 \pm 0.2	89.2 \pm 0.7	90.3 \pm 0.7	90.4 \pm 1.0
4	95.4 \pm 0.8	95.9 \pm 0.9	96.6 \pm 0.1	97.2 \pm 0.1	95.0 \pm 0.1	96.1 \pm 0.2
5	70.7 \pm 2.4	78.5 \pm 3.1	82.8 \pm 0.3	86.2 \pm 0.7	81.8 \pm 1.3	85.9 \pm 1.0
6	80.3 \pm 1.2	80.5 \pm 0.9	89.6 \pm 0.7	90.4 \pm 0.5	90.3 \pm 0.8	90.4 \pm 1.3
7	88.6 \pm 0.8	89.5 \pm 0.5	88.5 \pm 0.1	89.7 \pm 0.2	86.5 \pm 0.8	87.6 \pm 0.6
8	88.9 \pm 0.2	89.7 \pm 0.4	91.3 \pm 0.4	92.3 \pm 0.2	90.8 \pm 0.4	91.6 \pm 0.4
9	90.2 \pm 1.4	91.5 \pm 1.5	94.5 \pm 0.4	95.7 \pm 0.4	94.8 \pm 0.2	95.6 \pm 0.3
10	82.5 \pm 4.2	84.0 \pm 4.5	88.9 \pm 0.3	90.4 \pm 0.5	85.4 \pm 0.9	88.5 \pm 0.7
11	87.3 \pm 1.6	87.6 \pm 1.5	90.3 \pm 0.2	90.9 \pm 0.3	91.1 \pm 0.3	91.1 \pm 0.5
12	86.8 \pm 1.0	87.8 \pm 1.7	88.7 \pm 0.7	89.7 \pm 0.3	91.0 \pm 0.2	91.5 \pm 0.2
13	82.7 \pm 1.7	85.5 \pm 0.7	80.7 \pm 1.1	84.3 \pm 1.2	82.6 \pm 0.6	84.4 \pm 1.1
14	90.9 \pm 0.8	90.3 \pm 0.6	95.7 \pm 0.3	96.2 \pm 0.2	96.7 \pm 0.2	97.2 \pm 0.1
15	81.1 \pm 1.0	82.0 \pm 0.2	79.8 \pm 0.4	80.2 \pm 0.5	80.5 \pm 0.3	81.2 \pm 0.8
16	83.8 \pm 0.6	85.2 \pm 0.4	85.3 \pm 0.5	87.2 \pm 0.5	85.6 \pm 0.6	86.1 \pm 0.8
17	95.3 \pm 1.5	95.8 \pm 1.5	98.0 \pm 0.1	98.3 \pm 0.1	97.7 \pm 0.4	98.3 \pm 0.3
18	91.1 \pm 1.5	90.2 \pm 2.1	94.9 \pm 0.2	95.4 \pm 0.2	95.9 \pm 0.0	96.1 \pm 0.1
19	86.7 \pm 0.3	88.2 \pm 0.6	92.5 \pm 0.3	93.9 \pm 0.2	93.8 \pm 0.3	94.7 \pm 0.3

Table 6: AUROCS for each class of ImageNet30 for both of our scores when applying the Flip context augmentation. For each setting, we evaluated our method across three seeds.

Normal Class	CON ₂ (Equalize)		CON ₂ (Invert)		CON ₂ (Flip)	
	S _{LH}	S _{NNND}	S _{LH}	S _{NNND}	S _{LH}	S _{NNND}
0	91.0 \pm 0.5	92.7 \pm 0.5	94.8 \pm 0.7	94.7 \pm 0.9	92.1 \pm 0.7	92.4 \pm 1.3
1	97.8 \pm 0.2	98.7 \pm 0.2	98.5 \pm 0.2	99.2 \pm 0.1	99.1 \pm 0.1	99.5 \pm 0.1
2	99.6 \pm 0.1	99.5 \pm 0.1	99.9 \pm 0.0	99.9 \pm 0.0	99.9 \pm 0.0	99.9 \pm 0.0
3	82.8 \pm 1.1	82.1 \pm 2.0	82.9 \pm 1.2	79.1 \pm 0.9	82.6 \pm 0.6	82.6 \pm 1.0
4	90.5 \pm 0.3	90.1 \pm 0.7	95.0 \pm 0.1	94.8 \pm 0.2	94.7 \pm 0.2	95.6 \pm 0.5
5	91.0 \pm 2.2	93.3 \pm 1.7	93.6 \pm 0.3	94.4 \pm 0.2	94.8 \pm 0.4	96.7 \pm 0.4
6	94.8 \pm 0.7	95.5 \pm 0.3	97.1 \pm 0.1	98.0 \pm 0.1	96.4 \pm 0.2	96.9 \pm 0.1
7	67.9 \pm 1.7	68.7 \pm 0.3	77.1 \pm 1.0	78.5 \pm 0.8	75.8 \pm 1.5	76.5 \pm 1.6
8	95.4 \pm 0.2	95.7 \pm 0.4	93.4 \pm 0.7	92.7 \pm 1.3	96.6 \pm 0.5	96.9 \pm 0.3
9	74.8 \pm 1.0	77.7 \pm 1.1	86.8 \pm 0.3	88.9 \pm 0.3	84.9 \pm 0.8	86.6 \pm 0.6
10	97.8 \pm 0.2	97.8 \pm 0.1	99.2 \pm 0.1	99.3 \pm 0.0	99.0 \pm 0.2	99.0 \pm 0.2
11	82.4 \pm 1.4	82.4 \pm 1.6	85.4 \pm 0.3	84.4 \pm 0.5	89.2 \pm 1.0	90.2 \pm 0.7
12	90.3 \pm 0.4	93.0 \pm 0.6	95.5 \pm 0.2	97.3 \pm 0.1	96.6 \pm 0.2	97.6 \pm 0.3
13	91.9 \pm 0.7	91.8 \pm 0.5	95.2 \pm 0.3	95.5 \pm 0.2	94.0 \pm 0.5	94.0 \pm 0.3
14	85.1 \pm 0.1	86.6 \pm 0.7	91.6 \pm 0.6	91.9 \pm 0.5	93.3 \pm 0.3	94.2 \pm 0.2
15	90.9 \pm 0.9	89.3 \pm 1.3	93.9 \pm 0.6	92.7 \pm 0.2	93.2 \pm 1.9	93.1 \pm 2.0
16	96.6 \pm 0.4	97.5 \pm 0.3	98.9 \pm 0.1	99.2 \pm 0.1	99.0 \pm 0.1	99.5 \pm 0.2
17	45.7 \pm 1.1	51.2 \pm 2.3	59.0 \pm 1.0	62.9 \pm 1.0	50.9 \pm 1.1	55.2 \pm 0.6
18	78.4 \pm 0.7	80.3 \pm 0.9	89.2 \pm 0.3	89.8 \pm 0.5	92.2 \pm 0.6	93.1 \pm 0.4
19	59.6 \pm 2.4	61.4 \pm 2.7	75.1 \pm 0.7	76.3 \pm 0.5	67.1 \pm 3.4	68.5 \pm 3.7
20	86.3 \pm 1.0	86.2 \pm 0.6	92.2 \pm 0.5	93.0 \pm 0.6	94.2 \pm 0.4	94.9 \pm 0.6
21	86.5 \pm 0.3	87.0 \pm 0.9	95.7 \pm 0.3	96.4 \pm 0.1	95.7 \pm 0.2	96.2 \pm 0.2
22	95.3 \pm 0.5	94.4 \pm 0.8	96.7 \pm 0.5	96.1 \pm 0.3	97.3 \pm 0.3	97.4 \pm 0.2
23	94.1 \pm 0.4	94.5 \pm 0.4	96.3 \pm 0.2	96.7 \pm 0.3	96.4 \pm 0.1	96.9 \pm 0.2
24	72.0 \pm 1.3	73.7 \pm 0.7	90.3 \pm 0.3	92.4 \pm 0.5	88.3 \pm 1.0	90.9 \pm 1.1
25	83.3 \pm 2.2	85.3 \pm 1.8	84.6 \pm 1.3	84.1 \pm 1.5	73.6 \pm 2.3	74.1 \pm 1.0
26	95.1 \pm 0.2	95.2 \pm 0.5	93.3 \pm 0.4	92.3 \pm 0.6	89.1 \pm 0.8	88.7 \pm 0.9
27	91.6 \pm 0.9	91.3 \pm 1.0	96.3 \pm 0.2	96.8 \pm 0.3	97.0 \pm 0.3	97.5 \pm 0.2
28	57.3 \pm 1.5	61.4 \pm 2.5	69.0 \pm 1.3	72.1 \pm 1.8	73.8 \pm 2.8	77.7 \pm 3.4
29	87.0 \pm 1.1	91.4 \pm 1.1	88.5 \pm 1.7	90.8 \pm 2.1	86.1 \pm 1.0	91.3 \pm 1.4

Table 7: AUROCS for the two classes "Dog" and "Cat" for both of our scores when applying the Flip context augmentation. For each setting, we evaluated our method across three seeds.

Normal Class	CON ₂ (Equalize)		CON ₂ (Invert)		CON ₂ (Flip)	
	S _{LH}	S _{NNND}	S _{LH}	S _{NNND}	S _{LH}	S _{NNND}
0	78.4 \pm 1.7	84.1 \pm 0.9	88.3 \pm 0.1	90.0 \pm 0.2	91.0 \pm 0.1	92.1 \pm 0.2
1	80.6 \pm 0.9	82.3 \pm 0.5	87.4 \pm 0.4	87.6 \pm 0.4	89.7 \pm 0.4	89.7 \pm 0.2

Table 8: AUROCS for the two classes "Muffin" and "Chihuahua" for both of our scores when applying the Flip context augmentation. For each setting, we evaluated our method across three seeds.

Normal Class	CON ₂ (Equalize)		CON ₂ (Invert)		CON ₂ (Flip)	
	S _{LH}	S _{NNND}	S _{LH}	S _{NNND}	S _{LH}	S _{NNND}
0	94.0 \pm 0.8	94.4 \pm 1.1	95.6 \pm 0.3	96.8 \pm 0.2	95.8 \pm 0.2	96.7 \pm 0.3
1	73.6 \pm 1.3	79.2 \pm 0.3	87.8 \pm 0.5	91.8 \pm 0.3	91.1 \pm 0.7	93.1 \pm 0.3

1080
 1081
 1082
 1083
 1084
 1085
 1086
 1087
 1088
 1089
 1090
 1091
 1092
 1093
 1094
 1095
 1096
 1097
 1098
 1099
 1100
 1101
 1102
 1103
 1104
 1105
 1106
 1107
 1108
 1109
 1110
 1111
 1112
 1113
 1114
 1115
 1116
 1117
 1118
 1119
 1120
 1121
 1122
 1123
 1124
 1125
 1126
 1127
 1128
 1129
 1130
 1131
 1132
 1133

Method	Score	Pneumonia	Melanoma
$\mathcal{L}_{\text{Content}}$		92.3 ± 0.9	92.8 ± 0.2
$\mathcal{L}_{\text{Context}}$	\mathcal{S}_{LH}	79.9 ± 1.3	92.7 ± 0.5
$\mathcal{L}_{\text{Con}_2}$		93.0 ± 0.3	94.0 ± 0.3
$\mathcal{L}_{\text{Content}}$		89.6 ± 0.4	93.1 ± 0.3
$\mathcal{L}_{\text{Context}}$	\mathcal{S}_{NND}	81.4 ± 1.6	92.5 ± 0.8
$\mathcal{L}_{\text{Con}_2}$		93.9 ± 0.3	94.5 ± 0.2

Table 9: Evaluating the individual loss terms against each other when using the Equalize context augmentation on the medical datasets.