

Natural Synthetic Anomalies for Self-Supervised Anomaly Detection and Localization

Hannah M. Schlüter¹, Jeremy Tan¹, Benjamin Hou¹, and
Bernhard Kainz^{1,2}

¹ Imperial College London

{hannah.schlueter17,j.tan17,benjamin.hou1,b.kainz}@imperial.ac.uk

² Friedrich-Alexander-Universität Erlangen-Nürnberg

Abstract. We introduce a simple and intuitive self-supervision task, Natural Synthetic Anomalies (NSA), for training an end-to-end model for anomaly detection and localization using only normal training data. NSA integrates Poisson image editing to seamlessly blend scaled patches of various sizes from separate images. This creates a wide range of synthetic anomalies which are more similar to natural sub-image irregularities than previous data-augmentation strategies for self-supervised anomaly detection. We evaluate the proposed method using natural and medical images. Our experiments with the MVTec AD dataset show that a model trained to localize NSA anomalies generalizes well to detecting real-world a priori unknown types of manufacturing defects. Our method achieves an overall detection AUROC of **97.2** outperforming all previous methods that learn without the use of additional datasets. Code available at <https://github.com/hmsch/natural-synthetic-anomalies>.

Keywords: image anomaly localization, self-supervised learning.

1 Introduction

Anomaly detection is a binary classification task where the aim is to separate normal data from anomalous examples. There are different types of anomaly detection depending on what training data and labels are available. A difficult yet realistic setting is to detect and localize unknown types of anomalies while only having access to normal data during training. To be useful in real applications, an automated system must be able to detect subtle and rare anomalies; irregularities that are either impossible to spot for humans because of contextual uniformity or get lost due to task-remote stimuli that lead to inattentional blindness [16].

Attempting to detect rare anomalies often means that it is impossible to acquire sufficient amounts of human-annotated training data for a supervised method. Obtaining precise ground-truth annotations is time-consuming and requires expert knowledge depending on the application domain. Anomaly detection based on only normal data has applications in many areas including defect detection in industrial production pipelines [2,5,6,14,23,32], unsupervised lesion detection in medical images [17,20,27,28,29,38], or finding unusual events in surveillance videos [1,5,6].

The main challenge of unsupervised approaches is designing a training setup that will encourage the model to learn features relevant to anomaly detection without having any prior knowledge of the types of anomalies to expect. A common theme among top-performing approaches [6,11,23,24] for anomaly detection in natural images, specifically on the MVTec AD benchmark, is to sidestep this challenge by relying on deep features from pre-trained ImageNet models.

We find approaches that learn from scratch more interesting. They are more widely applicable to other domains where the usefulness of ImageNet pre-training is limited, such as medical imaging [22]. Many rely on learning a compressed representation of normal data and use these embeddings or reconstructions derived from them to define an anomaly score. Self-supervised learning is thus becoming a prominent strategy in anomaly detection. By designing an appropriate task, self-supervision can be an effective proxy for supervised learning, bypassing the need for labeled data. While various self-supervised tasks, such as context prediction [7] or estimating geometric transformations [8,9], can be used to learn a compressed representation of the data, recent works [14,28,29,36] show that data-augmentation strategies mimicking real defects are particularly effective for sub-image anomaly detection. These methods create synthetic anomalies by replacing or blending image patches with content from other images or image locations. However, recently proposed synthesis strategies [14,36] feature obvious discontinuities. This is raising concerns that the model may overfit to prior assumptions that are inherently encoded in synthetic manipulations. To prevent this, one can either resort to simpler encoders [14], which impedes end-to-end localization, or add additional networks for reconstruction [36], which greatly increases model size, computational costs, and training time. Even methods that linearly interpolate similar patches to create more subtle irregularities can suffer from the same problem [28]. [29] solves this problem by using Poisson image editing [21], but the resulting anomalies are so subtle that they may represent variations of the normal class rather than true anomalies.

Contribution. We introduce a simple and intuitive self-supervised method for sub-image anomaly detection and localization. Our Natural Synthetic Anomalies (NSA), are a) more natural than the current state-of-the-art CutPaste [14], FPI [28], or DRAEM [36] due to the use of Poisson image editing, b) more diverse than CutPaste, FPI, or PII anomalies due to rescaling, shifting and a new Gamma-distribution-based patch shape sampling strategy, and c) more relevant to the task by imposing background constraints and using pixel-level labels derived from the resulting difference to the normal image rather than interpolation factors as used in FPI and PII. Like FPI and PII, NSA can be used to train an end-to-end model for anomaly detection and localization rather than generating compressed representations for a multi-stage pipeline.

We evaluate the proposed method on the MVTec AD dataset [2] which contains normal training data and both normal and anomalous test data for a wide range of natural and manufacturing defects for 10 object and 5 texture classes. NSA achieves the new state-of-the-art localization (96.3 AUROC) and detection (97.2 AUROC) performance among methods that do not use additional datasets.

It also performs comparably to the best methods that use additional data and much bigger models [6,36]. [6] uses a model pre-trained on ImageNet. Compared to the large amount of data in ImageNet, our method only uses MVTec AD data, which contains between 60 and 391 training images per class.

NSA is a very general method for creating diverse and realistic synthetic anomalies in images. Since it does not rely on pre-training with ImageNet or any other dataset, it can easily be adapted to domains beyond natural images. Thus, we also evaluate NSA using a curated subset of a public chest X-ray dataset [33] where it outperforms other state-of-the-art self-supervised methods for disease detection.

2 Related Work

Reconstruction-based anomaly detection establishes pixel-level and image-level anomaly scores from the pixel-wise reconstruction error using variational autoencoders (VAE) [38], Bayesian autoencoders [20], generative adversarial networks (GAN) [27], or the restoration distance using a vector-quantized VAE (VQ-VAE) [17,32] trained with normal data. Anomaly scores can be improved by leveraging additional information derived from the model, such as the discriminator output when using a GAN [27], the KL-divergence of the latent representation of a VAE for image-level scores or its gradient for pixel-level scores [38], or the likelihood of the latent representation under a learnt prior using a VQ-VAE [17,32] or latent space autoregression [1]. A downside of these approaches is that it is difficult to control the capacity of the model. Depending on regularization, the model cannot reconstruct all details of normal examples well or may be able to reconstruct anomalous regions too.

Embedding-based anomaly detection derives an anomaly score from the distance between embedding vectors of normal training images and test examples. An embedding-similarity metric can be defined using any method for one-class classification such as support vector data description (SVDD) used in [34], Gaussian distributions used in [6,14,23], or nearest neighbour search used in [5,24]. Features for the embedding vectors are often extracted from pre-trained deep neural networks [5,6,23,24], but can also be learned from scratch using a self-supervised task [14] or together with the one-class classification objective in Deep-SVDD [26] and Deep One-Class Classification (DOCC) [25] or a combination thereof [34]. When using embeddings of the entire image, embedding-based approaches can perform detection but not localization and are hence less interpretable. To circumvent this issue, [6,14,24,34] work with patch-level embeddings to create an anomaly map while [5] compares test images to their nearest neighbors from the training set at the pixel-level. In a similar vein, flow-based methods can perform density estimation [10]. However, these methods can sometimes assign higher likelihood to outlier samples [18] and typically need pre-trained feature extractors to achieve better performance [35].

Self-supervised learning. A supervisory signal from a proxy task defined based on unlabeled data, such as predicting the relative position of patches [7]

or estimating geometric transformations [8,9], can help the model learn useful features for a downstream task. While [7,8,9] use features learned from the proxy task to discover different object classes, self-supervised learning has also been successfully applied to sub-image anomaly detection [14,28,29,34,36]. In [28,29], the output for the self-supervised task is a prediction of the interpolation factor where a foreign patch from another observation in the training distribution has been blended into the current image. This output is used directly as an anomaly score without any further training step. We also employ this general setup for our method. In [36], the foreign patches are first removed by a reconstructive sub-network and then a discriminative sub-network produces an anomaly map by comparing the input image with foreign patches to the reconstruction.

Poisson image editing. Pasting part of one image into another causes obvious discontinuities. [21] developed a method to seamlessly clone an object from one image into another image. For a source image given by g and a destination image given by f^* , we seek an interpolant f over the interior of a region Ω with boundary $\partial\Omega$ that solves the minimization problem given by (1). According to [21], this has the unique solution of the Poisson partial differential equations (2) with Dirichlet boundary conditions given by the destination image.

$$f = \arg \min_f \iint_{\Omega} |\nabla f - \mathbf{v}|^2 \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega} \quad (1)$$

$$\Delta f = \operatorname{div} \mathbf{v} \text{ over } \Omega, \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega} \quad (2)$$

[21] gives two options for defining the guidance field \mathbf{v} : a) use the source image gradient (3) or b) a mix of source and destination gradients (4).

$$\mathbf{v} = \nabla g \quad (3)$$

$$\forall \mathbf{x} \in \Omega, \mathbf{v}(\mathbf{x}) = \begin{cases} \nabla f^*(\mathbf{x}), & \text{if } |\nabla f^*(\mathbf{x})| > |\nabla g(\mathbf{x})|, \\ \nabla g(\mathbf{x}), & \text{otherwise} \end{cases} \quad (4)$$

In practice, a finite difference discretization of (2) is solved numerically. Seamless cloning is implemented in the OpenCV library [3] which we use in our self-supervised task. A Poisson image editing approach has also recently been used for anomaly detection in medical images [29].

3 NSA Self-supervised task

As only normal data is available at training time, the model needs to be trained using a proxy task. In our case, the task is to localize synthetic anomalies created from normal data by blending a patch from a source image into the destination image as follows:

1. Select a random rectangular patch in the source image.
2. Randomly resize the patch and select a different destination location.
3. Seamlessly blend the patch into the destination image.

4. Optionally, repeat steps 1-3 to add multiple patches to the same image.
5. Create a pixel-wise label mask.

Patch sampling and constraints: Given two normal $W \times H$ training images x_s and x_d , we select a random rectangular patch p_s with width w and height h sampled from a truncated Gamma distribution, and center (c_x, c_y) in the source image x_s sampled from a uniform distribution:

$$w = W \min(\max(w_{\min}, 0.06 + r_w), w_{\max}), \quad \text{with } r_w \sim \text{Gamma}(2, 0.1) \quad (5)$$

$$h = H \min(\max(h_{\min}, 0.06 + r_h), h_{\max}), \quad \text{with } r_h \sim \text{Gamma}(2, 0.1) \quad (6)$$

$$c_x \sim U\left(W \frac{w_{\min}}{2}, W - W \frac{w_{\min}}{2}\right), \quad c_y \sim U\left(H \frac{h_{\min}}{2}, H - H \frac{h_{\min}}{2}\right) \quad (7)$$

Sampling the width and height from a truncated Gamma distribution means we assume anomalies are local (small) but want the model to be able to recognize larger irregularities too. Hence, some long slim rectangles and occasionally large patches are generated. The width and height bounds are selected based on the dimensions of the object. For images containing an object and a plain background, we calculate object masks m_s and m_d by thresholding the pixel-wise absolute difference to the background brightness b . For each pixel i the masks are given by:

$$m_s^{(i)} = |x_s^{(i)} - b| < t_{\text{brightness}}, \quad m_d^{(i)} = |x_d^{(i)} - b| < t_{\text{brightness}} \quad (8)$$

We apply (7) repeatedly until $(p_s \cap m_s)/(wh) > t_{\text{object}}$ to ensure the patch contains part of the object. Then we resize the patch to obtain p'_s with width $w' = sw$ and height $h' = sh$. We select a destination patch p_d in the destination image x_d with the same dimensions and center (c'_x, c'_y) where:

$$s = \max\left(\frac{w_{\min}}{w}, \frac{h_{\min}}{h}, \min\left(r_s, \frac{w_{\max}}{w}, \frac{h_{\max}}{h}\right)\right) \quad \text{with } r_s \sim N(1, 1/4) \quad (9)$$

$$c'_x \sim U\left(W \frac{w'}{2}, W - W \frac{w'}{2}\right), \quad c'_y \sim U\left(H \frac{h'}{2}, H - H \frac{h'}{2}\right) \quad (10)$$

To prevent creating many examples of patches floating in the background, we apply (10) repeatedly until

$$(p_d \cap m_d)/(w'h') > t_{\text{object}} \quad (\text{contains part of the object}), \quad (11)$$

$$(m_{p_d} \cap m_{p'_s})/|m_{p'_s}| > t_{\text{overlap}} \quad (\text{object portions of patch and destination image overlap}) \quad (12)$$

where m_{p_d} and $m_{p'_s}$ are the object masks of the source and destination patches. We seamlessly blend p'_s into x_d at location (c'_x, c'_y) to obtain the training sample \tilde{x} . After blending the first patch, we add up to $n - 1$ further patches by flipping $n - 1$ coins whether to add another patch or not. Fig. 1 shows a simplified outline of how the synthetic anomalies are created.




Fig. 1: NSA anomalies are created by seamlessly cloning a patch from a normal training image into another normal training image.

Hyperparameters as assumptions: The patch selection procedure has a number of hyperparameters which characterize our synthetic out-distribution. Ideally, we would want this synthetic out-distribution to match the real out-distribution as closely as possible but since the real distribution is unknown at training time, we cannot use it to select the hyperparameters and sampling distributions. As [37] proved for generative models, no test statistic is useful for all possible out-distributions and some assumptions must be made to increase the likelihood that the test is useful for relevant or probable out-distributions. For our method, we try to keep these assumptions as broad as possible by generating many different sizes, locations, quantities, and shapes of anomalies and only disregarding those which are obviously irrelevant due to extreme sizes compared to the object or little overlap with the object.

Labels: We use the local intensity differences where a foreign patch has been introduced to create a pixel-wise label \tilde{y} which is either a) binary: whether there is a difference or not, b) continuous based on the mean absolute intensity difference across C color channels, or c) a logistic function of the previous. All labels are median filtered to be more coherent. Before filtering, the label values at each pixel i are calculated as follows:

$$\tilde{y}_{\text{binary}}^{(i)} = \begin{cases} 1, & \text{if } \tilde{x}^{(i)} \neq x_d^{(i)} \\ 0, & \text{otherwise} \end{cases}, \quad \tilde{y}_{\text{continuous}}^{(i)} = \frac{1}{C} \sum_{c=1}^C |\tilde{x}^{(i,c)} - x_d^{(i,c)}| \quad (13)$$

$$\tilde{y}_{\text{logistic}}^{(i)} = \frac{\tilde{y}_{\text{binary}}^{(i)}}{1 + \exp\left(-k\left(\tilde{y}_{\text{continuous}}^{(i)} - y_0\right)\right)} \quad (14)$$

In contrast, FPI [28] and PII [29] use the patch interpolation factor as a label. This is somewhat ill-posed because the interpolation factor cannot be determined without knowing the pixel intensities of both the source and destination patches.

Our labels are directly related to the change in intensity (created by the patch blending) and therefore provide a more consistent training signal.

Loss: When using bounded labels ($\tilde{y}_{\text{binary}}$ or $\tilde{y}_{\text{logistic}}$) we define our pixel-wise regression objective using binary cross-entropy loss. For unbounded labels ($\tilde{y}_{\text{cont.}}$) we use mean squared error loss. The loss is given in (15)–(16) where $\hat{y} = f(\hat{x})$ is the output of a deep convolutional encoder-decoder.

$$\mathcal{L}_{\text{bce}} = \frac{1}{W \times H} \sum_i -\tilde{y}_{\text{bounded}}^{(i)} \log \hat{y}^{(i)} - (1 - \tilde{y}_{\text{bounded}}^{(i)}) \log (1 - \hat{y}^{(i)}) \quad (15)$$

$$\mathcal{L}_{\text{mse}} = \frac{1}{W \times H} \sum_i \left(\tilde{y}_{\text{continuous}}^{(i)} - \hat{y}^{(i)} \right)^2 \quad (16)$$

By varying size, aspect ratio, source and destination location, and resizing the




Fig. 2: Synthetic anomalies created with CutPaste, FPI, PII, and NSA.

scale of the patches, this method dynamically creates a wide range of synthetic anomalies during training. The examples feature changes in size, shape, texture, location, and color of local image components as well as missing components by blending in a patch containing some background, while staying true to the overall distribution of the images and avoiding obvious discontinuities. Hence, these examples are a more realistic approximation of natural sub-image anomalies than CutPaste augmentations constructed by simply pasting patches at different locations [14] and more diverse than interpolating patches from two separate images at corresponding locations as in FPI [28] and PII [29] although still noticeably artificial to a human observer (Fig. 2).

4 Experiments

We compare end-to-end detection and localization models trained using our self-supervised task to end-to-end models trained using our implementations of FPI [28], PII [29], and CutPaste augmentation [14] on the MVTec AD dataset [2] and a curated subset of a public chest X-ray dataset [33]. We assess performance using the area under the receiver operating characteristic curve (AUROC).

Datasets: MVTec AD [2] contains normal training data and normal and anomalous test data featuring various types of natural and manufacturing defects for 10 object and 5 texture classes.

The NIH chest X-ray dataset [33] contains normal images as well as 14 different types of pathological patterns. There is a lot of natural variation in the normal class which is challenging for unsupervised methods. However, the most obvious differences are easily explained by the different views and the gender of the patients. We reduce this variation by reducing the curated subset defined in [31] further to only posteroanterior (back-to-front) view images of patients aged over 18 and separating them by gender. This leaves us with 1973 normal training images, 299 normal and 139 abnormal test images of male patients. For female patients, we have 1641 normal training, 244 normal and 123 abnormal test images. We call this dataset re-curated chest X-ray (rCXR) in the following. Note that the authors of PII [29] used the full NIH chest X-ray rather than the curated subset defined in [31].

4.1 Network architecture and training setup

We use an encoder-decoder architecture with ResNet-18 [12] without the classification layers as the encoder, two 1x1 convolutions in the bottleneck to reduce the number of channels and a simpler ResNet-based decoder. The final activation is sigmoid and we use binary-crossentropy loss for all models besides NSA (continuous) for which we use ReLU activation and mean squared error loss as the labels are unbounded. The models are trained on batches of size 64 using Adam [13] with a cosine-annealing learning rate [15] that decays from 10^{-3} to 10^{-6} over 320 epochs. For non-aligned objects, the loss takes longer to converge, so we use 560 epochs for the hazelnut, metal nut, and screw classes in the MVTec AD dataset. For rCXR, we use 240 epochs. The same training hyperparameters are used for all variants of the self-supervised task. Hyperparameters for the self-supervised task are given in the supplementary material. Note that in our implementation of FPI, PII, and CutPaste we use object masks and the patch sizes are sampled from a truncated Gamma distribution rather than a uniform distribution [14,28,29] to allow for a more fair comparison with NSA. We call our CutPaste baseline CutPaste (end-to-end) to distinguish it from the multi-stage framework used in [14].

The MVTec AD images have high resolutions of up to 1024×1024 pixels and use the RGB color scheme. We resize object images to 256×256 pixels, apply a random rotation of up to 5 degrees for non-aligned and rotation invariant objects (bottle, hazelnut, metal nut, screw), center-crop to 230×230 pixels and crop

a random 224×224 part of the image before creating self-supervised training examples to achieve slight rotation and translation invariance. When testing we use 224×224 center-crops of 256×256 object images. For texture classes, we use random 256×256 crops of 264×264 images for training and 256×256 images for testing. Intensities are normalized using the mean and standard deviation of ImageNet as commonly used before feeding them into the model.

The rCXR images have a resolution of 1024×1024 pixels in grayscale. We resize them to 256×256 pixels for training and apply a random rotation of up to 3 degrees, center-crop to 230×230 pixels and take a random crop of 224×224 pixels. For testing we use 224×224 center-crops of 256×256 resampled images. *Implementation:* We use PyTorch [19] V1.8.1 and train each model on an Nvidia GeForce GTX 1080 GPU while the self-supervised examples are created in parallel using 8 processes on an Intel Core i7-7700K CPU. The code is available at <https://github.com/hmsch/natural-synthetic-anomalies>.

4.2 Results and evaluation

Defect detection. In Tab. 1 we compare the detection performance of our models trained using variations of NSA, FPI [28], PII [29], and CutPaste [14] to CutPaste (3-way) from [14] which was previously the top-performing approach for defection detection in the MVTec AD dataset [2] without using additional datasets. For NSA, we report the mean and standard error of the detection AUROC for each class as well as the object, texture, and overall averages across five different random seeds. Our best method, NSA (logistic), achieves an overall image-level AUROC of **97.2** outperforming CutPaste (3-way) [14] by 2.0 which is well outside of the standard error range. A single NSA (logistic) model is also better than an ensemble of 5 CutPaste (3-way) models [14] (96.1 AUROC) and comparable to EfficientNet [30] pre-trained with ImageNet and finetuned with CutPaste (3-way) [14] (97.1 AUROC).

Methods using pre-trained ImageNet models, such as PaDiM [6] (97.9 AUROC), do not provide a fair comparison to our from-scratch approach. DRAEM [36] is a more relevant competitor. It uses the describable textures dataset (DTD) [4] not for pre-training but for creating its synthetic anomalies. DRAEM’s training objective consists of localizing and correcting the synthetic anomalies, while our model performs anomaly localization. DRAEM (98.0 AUROC) outperforms our approach slightly overall, however both approaches give comparable results for many classes despite the fact that our models only have around 11 million parameters while DRAEM’s two components add up to over 97 million parameters. We note that no standard errors were reported for DRAEM’s results. The authors of DRAEM claim that realism of the synthetic anomalies is not important for their method. However, their own ablation study indicates that using plain solid colors instead of real textures from DTD, causes localization AUROC and AP to drop from 97.1 and 68.4 down to 92.6 and 56.5, respectively [36]. As such, external data from DTD, which is designed to span a range of textures found in the wild [4], may help to produce synthetic anomalies that overlap more with real anomalies. Furthermore, when using a similar training setup as ours, *i.e.*,

without its reconstructive sub-network, DRAEM performs worse (93.9 AUROC) than NSA which uses more realistic synthetic anomalies.

Table 1: Image-level AUROC % for MVTec AD and standard error across five different random seeds. For our models, the image-level score is the average pixel score across the image. Best scores between DRAEM [36], CutPaste (3-way) [14], and NSA within standard error are bold-faced. Note that DRAEM uses additional data.

		SOTA		Our Experiments						
		DRAEM [36]	CutPaste (3-way) [14]	CutPaste (end-to-end)	FPI	PII	NSA (binary)	NSA (continuous)	NSA (logistic)	
object	bottle	99.2	98.3 ± 0.5	100.0	90.2	97.6	97.6 ± 0.2	97.5 ± 0.2	97.7 ± 0.3	
	cable	91.8	80.6 ± 0.5	75.4	68.0	68.9	92.1 ± 2.4	90.2 ± 3.0	94.5 ± 1.0	
	capsule	98.5	96.2 ± 0.5	89.2	87.5	84.9	93.2 ± 0.8	92.8 ± 2.2	95.2 ± 1.7	
	hazelnut	100.0	97.3 ± 0.3	81.4	86.0	82.7	93.5 ± 1.9	89.3 ± 4.9	94.7 ± 1.1	
	metal nut	98.7	99.3 ± 0.2	70.6	88.4	98.9	99.4 ± 0.3	94.6 ± 2.1	98.7 ± 0.7	
	pill	98.9	92.4 ± 1.3	90.3	71.8	86.3	97.0 ± 0.9	94.3 ± 1.1	99.2 ± 0.6	
	screw	93.9	86.3 ± 1.0	65.5	61.2	74.7	90.3 ± 1.2	90.1 ± 0.9	90.2 ± 1.4	
	toothbrush	100.0	98.3 ± 0.9	96.7	85.8	93.1	100.0 ± 0.0	99.6 ± 0.5	100.0 ± 0.0	
	transistor	93.1	95.5 ± 0.5	88.2	79.6	90.1	93.5 ± 0.9	92.8 ± 2.2	95.1 ± 0.2	
	zipper	100.0	99.4 ± 0.2	98.7	97.7	99.8	99.8 ± 0.1	99.5 ± 0.7	99.8 ± 0.1	
average		97.4	94.3 ± 0.6	85.6	81.6	87.7	95.6 ± 0.5	94.1 ± 1.0	96.5 ± 0.3	
texture	carpet	97.0	93.1 ± 1.1	53.1	56.0	65.6	85.6 ± 7.6	90.9 ± 2.2	95.6 ± 0.6	
	grid	99.9	99.9 ± 0.1	99.7	99.5	100.0	99.9 ± 0.1	98.5 ± 3.3	99.9 ± 0.1	
	leather	100.0	100.0 ± 0.0	86.6	91.7	100.0	99.9 ± 0.1	100.0 ± 0.0	99.9 ± 0.1	
	tile	99.6	93.4 ± 1.0	87.8	90.2	98.4	99.7 ± 0.2	100.0 ± 0.0	100.0 ± 0.0	
	wood	99.1	98.6 ± 0.5	84.6	74.4	91.9	96.7 ± 1.2	97.8 ± 0.8	97.5 ± 1.5	
	average	99.1	97.0 ± 0.5	82.4	82.4	91.2	96.4 ± 1.4	97.5 ± 0.9	98.6 ± 0.3	
overall average		98.0	95.2 ± 0.6	84.5	81.9	88.9	95.9 ± 0.7	95.2 ± 0.5	97.2 ± 0.3	

Synthetic anomalies should be as diverse and realistic as possible. In our experiments, models trained with self-supervised examples created using Poisson blending clearly outperform models trained with simpler data-augmentation strategies like CutPaste [14] and FPI [28]. Examples created with Poisson blending are more visually similar to real-world defects as they do not have artificial discontinuities (Fig. 2) and according to Tab. 1, the corresponding models generalize better to real defects. Although PII [29] performs well for most texture classes, NSA, which also shifts and resizes the patches, performs much better for objects. Since PII and FPI use the same source and destination location for the patches, their synthetic anomalies are very subtle for aligned object classes and much subtler than the real defects in the MVTec AD dataset. In classes with lower AUROC, the synthetic training anomalies may have less similarity to the real test anomalies. These classes also tend to have higher standard error. In contrast, classes with high AUROC have low standard error. This could indicate that for these classes the self-supervised task can sensitize the network to the

distribution of real anomalies reliably. It may also be possible to use the variance between random seeds to gauge the reliability of predictions.

Labels should approximate the degree of abnormality. Aside from abnormal irregularities there can also be natural variation between and within the images of each class. When training a model with binary labels, the final activations tend to saturate and the predictions do not give any measure of how anomalous the regions with high scores are. Training a model with continuous labels teaches the model to differentiate between different degrees of anomalies. When using unbounded continuous labels, training is less stable and the AUROC scores have a high standard error. Models trained with bounded continuous labels outperform the binary ones most in classes with high inherent variation such as cable, hazelnut, transistor, carpet, and wood (Tab. 1) when using a threshold independent metric such as AUROC. PII [29] also uses continuous labels; pixels corresponding to the foreign patch are assigned to a uniform value equivalent to the interpolation factor. However, NSA (logistic) outperforms PII for some textures and most objects, including unaligned objects. For aligned objects, NSA creates more diverse anomalies than PII because the location of the source and destination patches can be different. But for unaligned objects, this advantage is negligible. Despite the similarity in generated anomalies, NSA still yields higher performance in these classes. A possible explanation is that the labels for NSA (logistic) are inhomogeneous and based on the outcome of the blending rather than its setup and hence more accurately represent the degree of abnormality.

Defect localization. In Tab. 2 we report the pixel-level performance of the models from Tab. 1. Although NSA performs well for objects on the image-level when trained with unbounded continuous labels, pixel-wise performance is much better for NSA with bounded binary or continuous labels. NSA (logistic) achieves a **96.3** average pixel-level AUROC performing similarly to CutPaste (3-way) [14] (96.0 AUROC). It also reaches similar performance to DRAEM [36] for many classes although DRAEM achieves a higher overall score (97.3 AUROC). Fig. 3 shows that NSA (logistic) can localize a wide range of real-world defects accurately including anomalies that are very different from the synthetic anomalies seen during training (*e.g.*, white writing on hazelnuts, misplaced transistors, stained tiles and carpet).

Medical imaging. In Tab. 3, we compare the performance of NSA and end-to-end models trained using other self-supervised tasks for the task of binary classification of rCXR images into healthy (normal) and pathological (abnormal) categories. Models trained using NSA clearly outperform end-to-end models trained with FPI [28] or CutPaste [14]. NSA also outperforms PII [29]. However, the type of label used for NSA is less important for this dataset. Since there is high inter-sample variability in the normal data, it is possible that all synthetic anomalies created by NSA would be considered abnormal. So, approximating the degree of abnormality with a continuous label does not improve over the binary labels when evaluating the model with real anomalies.

We do not report pixel-level metrics, as we only have very rough bounding boxes for less than 10% of the test set. Fig. 4 shows example predictions for

Table 2: Pixel-level AUROC % with standard error for the MVTec AD models from Tab. 1. Note that DRAEM uses additional data.

		SOTA		Our Experiments					
		DRAEM [36]	CutPaste (3-way) [14]	CutPaste (end-to-end)	FPI	PII	NSA (binary)	NSA (continuous)	NSA (logistic)
object	bottle	99.1	97.6 \pm 0.1	97.7	91.8	93.1	98.4 \pm 0.2	97.3 \pm 0.5	98.3 \pm 0.1
	cable	94.7	90.0 \pm 0.2	81.0	66.5	70.2	93.3 \pm 3.4	91.0 \pm 2.8	96.0 \pm 1.4
	capsule	94.3	97.4 \pm 0.1	97.5	95.9	90.2	98.1 \pm 0.2	91.6 \pm 5.6	97.6 \pm 0.9
	hazelnut	99.7	97.3 \pm 0.1	94.8	89.8	97.0	97.2 \pm 0.6	97.7 \pm 0.6	97.6 \pm 0.6
	metal nut	99.5	93.1 \pm 0.4	68.1	96.2	95.4	98.2 \pm 0.2	97.3 \pm 0.3	98.4 \pm 0.2
	pill	97.6	95.7 \pm 0.1	98.1	62.3	95.3	98.5 \pm 0.2	97.1 \pm 2.7	98.5 \pm 0.3
	screw	97.6	96.7 \pm 0.1	90.7	90.4	92.8	96.7 \pm 0.4	92.3 \pm 5.3	96.5 \pm 0.1
	toothbrush	98.1	98.1 \pm 0.0	95.7	81.8	81.3	95.6 \pm 0.6	94.5 \pm 0.7	94.9 \pm 0.7
	transistor	90.9	93.0 \pm 0.2	85.9	78.5	86.9	87.8 \pm 1.9	80.2 \pm 3.3	88.0 \pm 1.8
	zipper	98.8	99.3 \pm 0.0	92.9	91.8	93.8	94.2 \pm 0.2	90.7 \pm 1.5	94.2 \pm 0.3
	average	97.0	95.8 \pm 0.1	90.2	84.5	89.6	95.8 \pm 0.4	93.0 \pm 1.7	96.0 \pm 0.4
texture	carpet	95.5	98.3 \pm 0.0	83.3	70.8	97.2	94.5 \pm 4.1	81.8 \pm 6.8	95.5 \pm 2.3
	grid	99.7	97.5 \pm 0.1	97.6	94.2	98.9	99.1 \pm 0.0	98.0 \pm 0.3	99.2 \pm 0.1
	leather	98.6	99.5 \pm 0.0	96.4	88.3	99.2	99.6 \pm 0.0	99.5 \pm 0.2	99.5 \pm 0.1
	tile	99.2	90.5 \pm 0.2	72.7	65.0	98.0	99.0 \pm 0.2	97.4 \pm 0.7	99.3 \pm 0.0
	wood	96.4	95.5 \pm 0.1	84.0	71.1	91.1	94.0 \pm 0.8	90.6 \pm 3.7	90.7 \pm 1.9
	average	97.9	96.3 \pm 0.1	86.8	77.9	96.9	97.3 \pm 0.7	93.5 \pm 0.9	96.8 \pm 0.7
	overall average	97.3	96.0 \pm 0.1	89.1	82.3	92.0	96.3 \pm 0.2	93.1 \pm 1.1	96.3 \pm 0.4

several healthy and pathological cases. The localization predictions are good for examples 1–5, 7, and 8, but the model fails to detect any abnormal findings in the 6th example and disagrees with the bounding box annotation for bottom examples 9 and 10. We did not compare to CutPaste (3-way) [14] here as there is no reference implementation available.

Limitations. The 9th example in Fig. 4 shows cardiomegaly, *i.e.*, an enlarged heart, for which the bounding box contains the entire heart. But the model only activates for portions of the heart that exceed the normal size found in healthy patients. In these cases, the model lacks the semantic understanding that radiologists use to categorize diseases. In the 10th example, the model activates more for the tubes on the patient’s right side than for the finding

Table 3: Image-level AUROC % for rCXR and standard error across five different random seeds. Best scores per row within standard error are bold-faced.

	CutPaste (end-to-end)	FPI	PII	NSA (binary)	NSA (continuous)	NSA (logistic)
male	59.8	73.7	91.7 \pm 0.6	94.0 \pm 0.5	93.4 \pm 0.3	94.0 \pm 0.6
female	56.2	67.4	92.8 \pm 0.4	94.3 \pm 0.6	93.0 \pm 0.4	94.0 \pm 0.5

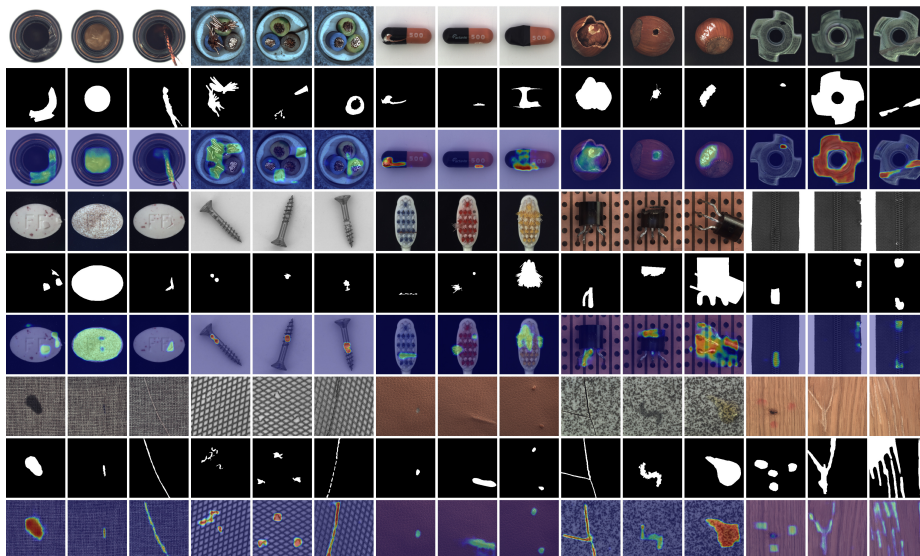


Fig. 3: Examples of defect localization in the MVTec AD dataset using models trained with NSA (logistic). From top to bottom: input images, human annotation, heatmap of pixel-level predictions. Best viewed in a digital version.

inside of the bounding box. Unlike a human radiologist, an anomaly detection model cannot be expected to automatically classify clinically correctly placed lines, tubes, or cardiac devices as normal if they are not expected in the healthy training distribution.

Our method sometimes fails to detect very small defects (see capsule, hazelnut, screw, and wood examples in Fig. 5), predicts too large regions or has false positives (see bottle, cable, zipper, carpet, leather in Fig. 5). The transistor class has several examples of misplaced or missing transistors. These are detected but the predicted localization does not match the large human annotation (Fig. 5) resulting in a low localization AUROC (Tab. 2). Patch-wise localization, as used for CutPaste [14] and PaDiM [6] among others, can detect the missing transistor for each patch and hence produce a segmentation map closer to the human annotation. However, for these large anomalies the type of defect is immediately obvious to a human inspector once detected so we argue that precise localization would not be necessary for most applications.

Since the model does not see any real anomalies during training, it may predict statistically unlikely normal variations as abnormal and fails to recognize subtle anomalies that are very different from the synthetic anomalies seen during training. Hence, predictions from a self-supervised anomaly detection model should not be used on their own for decision making. Such models can however be useful as an instant second observer and for quality control. As far as we are aware, there are no further potential negative societal impacts of this work.

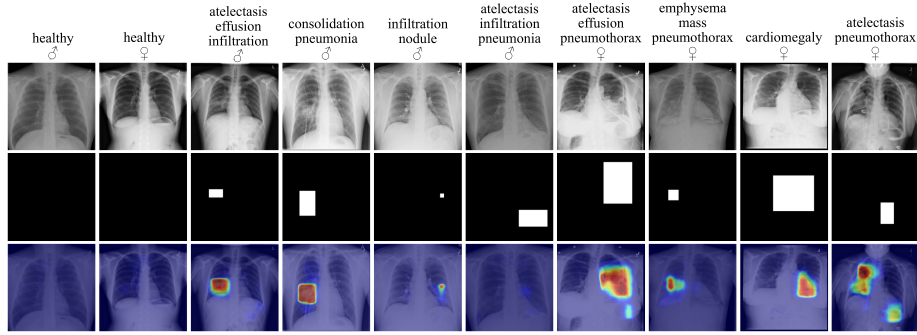


Fig. 4: Example localization predictions for chest X-ray disease detection using models trained with NSA (binary). From top to bottom: input images, rough bounding box from a radiologist, heatmap of pixel-level predictions. Each case is labeled with pathology keywords that [33] mined from radiologist reports.




Fig. 5: Failure cases for MVTec AD defect localization using models trained with NSA (logistic). Examples include false pos./negatives and incorrect localization. From left to right: input, human annotation, heatmap of pixel-level predictions.

5 Conclusion

We propose a self-supervised task that creates diverse and realistic synthetic anomalies. These training examples are generated under controlled conditions that help to produce relevant and subtle anomalies. This provides a more consistent training signal and results in better detection of real anomalies. The formulation of the loss and synthetic labels yields an effective and computationally efficient training task. This helps NSA outperform state-of-the-art methods on both natural and medical imaging datasets, demonstrating its generalizability. In the future, additions such as quantifying uncertainty or exploiting classes of known anomalies could help facilitate the use of NSA in critical applications.

Acknowledgments. This work was supported by the UK Research and Innovation London Medical Imaging and Artificial Intelligence Centre for Value Based Healthcare and the iFind project, Wellcome Trust IEH Award [102431].

References

1. Abati, D., Porrello, A., Calderara, S., Cucchiara, R.: Latent space autoregression for novelty detection. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019. pp. 481–490. Computer Vision Foundation / IEEE (2019)
2. Bergmann, P., Batzner, K., Fauser, M., Sattlegger, D., Steger, C.: The mvtec anomaly detection dataset: A comprehensive real-world dataset for unsupervised anomaly detection. *International journal of computer vision* **129**(4), 1038–1059 (2021)
3. Bradski, G.: The OpenCV Library. *Dr. Dobb's Journal of Software Tools* (2000)
4. Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., Vedaldi, A.: Describing textures in the wild. In: 2014 IEEE Conference on Computer Vision and Pattern Recognition. pp. 3606–3613 (2014)
5. Cohen, N., Hoshen, Y.: Sub-image anomaly detection with deep pyramid correspondences (2021)
6. Defard, T., Setkov, A., Loesch, A., Audigier, R.: Padim: A patch distribution modeling framework for anomaly detection and localization. In: Pattern Recognition. ICPR International Workshops and Challenges - Virtual Event, January 10-15, 2021, Proceedings, Part IV. Lecture Notes in Computer Science, vol. 12664, pp. 475–489. Springer (2020)
7. Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: 2015 IEEE International Conference on Computer Vision (ICCV). pp. 1422–1430 (2015)
8. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. In: International Conference on Learning Representations (2018)
9. Golan, I., El-Yaniv, R.: Deep anomaly detection using geometric transformations. In: Proceedings of the 32nd International Conference on Neural Information Processing Systems. p. 9781–9791. NIPS'18, Curran Associates Inc., Red Hook, NY, USA (2018)
10. Grathwohl, W., Chen, R.T.Q., Bettencourt, J., Sutskever, I., Duvenaud, D.: FFJORD: free-form continuous dynamics for scalable reversible generative models. In: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019. OpenReview.net (2019)
11. Gudovskiy, D.A., Ishizaka, S., Kozuka, K.: Cflow-ad: Real-time unsupervised anomaly detection with localization via conditional normalizing flows. In: 2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 1819–1828 (2022)
12. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (2016)
13. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings (2015)
14. Li, C.L., Sohn, K., Yoon, J., Pfister, T.: Cutpaste: Self-supervised learning for anomaly detection and localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9664–9674 (June 2021)
15. Loshchilov, I., Hutter, F.: SGDR: stochastic gradient descent with warm restarts. In: 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings. OpenReview.net (2017)

16. Mack, A., Rock, I., et al.: Inattention blindness. MIT press (1998)
17. Marimont, S.N., Tarroni, G.: Anomaly detection through latent space restoration using vector quantized variational autoencoders. In: 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI). pp. 1764–1767 (2021)
18. Nalisnick, E.T., Matsukawa, A., Teh, Y.W., Görür, D., Lakshminarayanan, B.: Do deep generative models know what they don't know? In: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019. OpenReview.net (2019)
19. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: Advances in Neural Information Processing Systems 32, pp. 8024–8035. Curran Associates, Inc. (2019)
20. Pawlowski, N., Lee, M.J., Rajchl, M., McDonagh, S.G., Ferrante, E., Kamnitsas, K., Cooke, S., Stevenson, S., Khetani, A., Newman, T., Zeiler, F.A., Digby, R., Coles, J., Rueckert, D., Menon, D., Newcombe, V., Glocker, B.: Unsupervised lesion detection in brain ct using bayesian convolutional autoencoders. In: OpenReview (2018)
21. Pérez, P., Gangnet, M., Blake, A.: Poisson image editing. In: ACM SIGGRAPH 2003 Papers. p. 313–318. SIGGRAPH '03, Association for Computing Machinery, New York, NY, USA (2003)
22. Raghu, M., Zhang, C., Kleinberg, J., Bengio, S.: Transfusion: Understanding transfer learning for medical imaging. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 32. Curran Associates, Inc. (2019)
23. Rippel, O., Mertens, P., Merhof, D.: Modeling the distribution of normal data in pre-trained deep features for anomaly detection. In: 25th International Conference on Pattern Recognition, ICPR 2020, Virtual Event / Milan, Italy, January 10-15, 2021. pp. 6726–6733. IEEE (2020)
24. Roth, K., Pemula, L., Zepeda, J., Schölkopf, B., Brox, T., Gehler, P.: Towards total recall in industrial anomaly detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 14318–14328 (June 2022)
25. Ruff, L., Kauffmann, J.R., Vandermeulen, R.A., Montavon, G., Samek, W., Kloft, M., Dietterich, T.G., Müller, K.R.: A unifying review of deep and shallow anomaly detection. Proceedings of the IEEE (2021)
26. Ruff, L., Vandermeulen, R., Goernitz, N., Deecke, L., Siddiqui, S.A., Binder, A., Müller, E., Kloft, M.: Deep one-class classification. In: Proceedings of the 35th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 80, pp. 4393–4402. PMLR (07 2018)
27. Schlegl, T., Seeböck, P., Waldstein, S.M., Langs, G., Schmidt-Erfurth, U.: f-anogan: Fast unsupervised anomaly detection with generative adversarial networks. Medical Image Analysis **54**, 30 – 44 (2019)
28. Tan, J., Hou, B., Batten, J., Qiu, H., Kainz, B.: Detecting outliers with foreign patch interpolation. Machine Learning for Biomedical Imaging **1** (2022)
29. Tan, J., Hou, B., Day, T., Simpson, J., Rueckert, D., Kainz, B.: Detecting outliers with poisson image interpolation. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2021. pp. 581–591. Springer International Publishing, Cham (2021)

30. Tan, M., Le, Q.: EfficientNet: Rethinking model scaling for convolutional neural networks. In: Proceedings of the 36th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 97, pp. 6105–6114. PMLR (06 2019)
31. Tang, Y.X., Tang, Y.B., Peng, Y., Yan, K., Bagheri, M., Redd, B.A., Brandon, C.J., Lu, Z., Han, M., Xiao, J., Summers, R.M.: Automated abnormality classification of chest radiographs using deep convolutional neural networks. *npj Digital Medicine* **3**(1), 1–8 (2020)
32. Wang, L., Zhang, D., Guo, J., Han, Y.: Image anomaly detection using normal data only by latent space resampling. *Applied Sciences* **10**(23) (2020)
33. Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M.: Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2097–2106 (2017)
34. Yi, J., Yoon, S.: Patch svdd: Patch-level svdd for anomaly detection and segmentation. In: Proceedings of the Asian Conference on Computer Vision (ACCV) (11 2020)
35. Yu, J., Zheng, Y., Wang, X., Li, W., Wu, Y., Zhao, R., Wu, L.: Fastflow: Unsupervised anomaly detection and localization via 2d normalizing flows. *CoRR abs/2111.07677* (2021)
36. Zavrtanik, V., Kristan, M., Skocaj, D.: Draem-a discriminatively trained reconstruction embedding for surface anomaly detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8330–8339 (2021)
37. Zhang, L., Goldstein, M., Ranganath, R.: Understanding failures in out-of-distribution detection with deep generative models. In: Meila, M., Zhang, T. (eds.) Proceedings of the 38th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 139, pp. 12427–12436. PMLR (18–24 Jul 2021)
38. Zimmerer, D., Isensee, F., Petersen, J., Kohl, S., Maier-Hein, K.H.: Unsupervised anomaly localization using variational auto-encoders. In: Medical Image Computing and Computer Assisted Intervention - MICCAI 2019 - 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part IV. Lecture Notes in Computer Science, vol. 11767, pp. 289–297. Springer (2019)