

This WACV 2021 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

TrustMAE: A Noise-Resilient Defect Classification Framework using Memory-Augmented Auto-Encoders with Trust Regions

Daniel Stanley Tan^{1,2}, Yi-Chun Chen², Trista Pei-Chun Chen², Wei-Chao Chen^{2,3} ¹National Taiwan University of Science and Technology, Taiwan ²Inventec Corporation, Taiwan ³Skywatch Innovation Inc., Taiwan

d10515805@mail.ntust.edu.tw, yichun8447@gmail.com chen.trista@inventec.com, weichao.chen@skywatch24.com

Abstract

In this paper, we propose a framework called Trust-MAE to address the problem of product defect classification. Instead of relying on defective images that are difficult to collect and laborious to label, our framework can accept datasets with unlabeled images. Moreover, unlike most anomaly detection methods, our approach is robust against noises, or defective images, in the training dataset. Our framework uses a memory-augmented auto-encoder with a sparse memory addressing scheme to avoid overgeneralizing the auto-encoder, and a novel trust-region memory updating scheme to keep the noises away from the memory slots. The result is a framework that can reconstruct defect-free images and identify the defective regions using a perceptual distance network. When compared against various state-of-the-art baselines, our approach performs competitively under noise-free MVTec datasets. More importantly, it remains effective at a noise level up to 40% while significantly outperforming other baselines.

1. Introduction

For manufacturing facilities, product appearance assessment is an essential step in quality assurance. Undetected defects, such as scratches, bumps, and discolorations, can result in costly product returns and customer trust loss. Today, humans still perform most appearance inspection tasks because of the difficulty of describing various defects using traditional computer vision algorithms in the automatic optical inspection machines (AOIs). However, managing human inspectors has been a significant management problem because it is difficult to maintain a consistent inspection standard across different product lines.

At first, it appears that one could address the problem with the ever-popular object detector networks [9, 10, 24,

37]. However, these fully-supervised models [21, 30, 31] require datasets with clearly annotated bounding boxes, which can be laborious and equally tricky to label with consistency. Additionally, because these methods tend to perform poorly for defects not present in the dataset, it can take an indefinite amount of time to collect sufficient training data. In many cases, we have experienced months of delays between data collection and model acceptance by the facilities, which is unacceptable for products with relatively short life cycles.

Instead of relying on explicitly labeled defects, it is possible to learn the distribution from normal samples and treat those deviating too far as defects, thus enabling the models to detect previously unseen defects [8, 13, 19]. A popular candidate is an auto-encoder that can erase defects from the input images upon trained with normal images [3, 15, 25, 27, 44, 46]. However, in practice, the autoencoders can become overly general and learn to reconstruct the defects [12]. In particular, when the surfaces contain lots of texture [2, 3], the reconstructions can be erratic, leading to many false-positives.

Also, while the generative approaches do not require detailed labeling of the images, they often assume the input data are free of defective images. As a result, the algorithms can become overly sensitive to noise when defective images accidentally leak into the dataset, which frequently occurs in many manufacturing facilities. Furthermore, many input images tend to contain some imperfections, and if we need to detect these as defects, the percentage of normal images would undoubtedly drop. Thus, it becomes essential for our approach to be resilient to data noise, as the defect-free image regions can also be helpful for training purposes.

To address these issues using auto-encoders, we need to limit the latent space effectively such that it can still reconstruct normal image regions without over-generalizing to defects. For this purpose, we devise a scheme inspired by MemAE [12] with several significant differences. We adopt a memory store for the latent space, but during the update phase, we increase the sparsity such that the updated information focus only on essential memory slots. Additionally, we add trust regions, which essentially classify defective latent space samples and avoid noisy samples from polluting the memory slots. To solve the textured surface problem, we further employ a perceptually-based distance approach to effectively separate defects from normal. Figure 1 illustrates the proposed TrustMAE method. We can achieve competitive results on the MVTec [2] dataset across all product types against baselines. Our framework is resilient to noise, achieving good performance even when the training data contain over 40% of defective images. We further include an ablation study in the results to demonstrate the effectiveness of these ideas. Our contributions are as follows:

- A defect classification framework resilient to noise in the training dataset,
- A novel sparse memory addressing scheme to avoid over-generalization of memory slots for the autoencoder,
- A memory update scheme using trust regions to avoid noise contamination, and
- A perceptual-based classification method to effectively detect defective high-level regions for textured images.

2. Trust Memory Auto-Encoder (TrustMAE)

Given a dataset containing both normal and defective images, we aim to learn a model to separate normal from defective images without accessing labels that differentiate the two. Moreover, by treating small imperfections as defects, we reduce the percentage of defect-free images in the dataset. Consequently, it would be desirable to utilize both the normal images and the good image patches within the defective images to increase the available data for training, which means the model should be robust against noise (defect images).

Figure 1 shows an overview of our proposed method, which contains the following three components:

- 1. A memory-augmented auto-encoder (Section 2.1) with sparse memory addressing. It stores normal image features in a memory bank and reconstructs a normal version of the input image from memory. Large differences between the input and its reconstruction indicate potential defects.
- Trust region memory updates (Section 2.2). This update process prevents the memory from storing unwanted defect features.

3. Spatial perceptual distance (Section 2.3). It produces patch-level differences using deep features between the input image with its reconstructed normal, making our methods more robust to slight shifts and color inaccuracies than shallow pixel-wise distances such as mean square error.

Next, we describe these components and how they can increase robustness against data noise.

2.1. Memory-Augmented Auto-Encoder

The standard auto-encoder follows an encoder-decoder structure wherein an encoder E projects the input image $x \in \mathbb{R}^{H \times W}$ into a lower-dimensional latent vector z, and the decoder D produces an output image $\hat{x} \in \mathbb{R}^{H \times W}$ which is a reconstruction of the input image x. To prevent defects from being reconstructed, we incorporate an external memory [12, 14] module **M** that stores a set of normal feature prototypes during training. We can think of these feature prototypes as puzzle pieces that the network uses to reconstruct the input image. At inference time, the set of feature prototypes are fixed, which makes it harder for the autoencoder to reconstruct defects because the memory module only contains normal features, assuming it was trained on only normal images.

2.1.1 Reconstructing from Memory

The memory module is implemented as a tensor \mathbf{M} \in $\mathbb{R}^{M \times Z}$ where M denotes the number of memory slots, and Z denotes the dimensions of the latent vector z. Given an input image x, we first compute its latent representation $\mathbf{Z} = E(x)$. Note that in our setting, we want to preserve the spatial information so we designed our encoder to produce a feature map $\mathbf{Z} \in \mathbb{R}^{h \times w \times Z}$ with a lower resolution than the original input image. For convenience, we use $z_i \in \mathbb{R}^Z$ to refer to the i^{th} element of **Z**. Each vector z_i in **Z** represents features of a patch in the input image. Instead of passing the features map \mathbf{Z} directly to the decoder D, we compute approximate features \hat{z}_i for every z_i using a convex combination of the feature prototypes stored in the memory module. Eq. 1 defines the memory addressing operation, where w is a weight vector indicating how similar z is with each of the feature prototypes stored in the memory module M. The decoder then outputs a reconstruction $\hat{x} = D(\mathbf{Z})$ using only approximates \hat{z} derived from the memory entries.

$$\hat{z} = w\mathbf{M}, \ \sum_{i=1}^{M} w_i = 1 \tag{1}$$

The weight vector w acts as a soft-addressing mechanism that retrieves the closest feature prototypes from the memory that are necessary for reconstruction. We measured



Figure 1: Framework of TrustMAE. Given a defective image, TrustMAE reconstructs a normal version of the input from memory. Then, it computes the perceptual difference between the input and its reconstruction, generating a distance map where large errors are indicative of defects.

the similarity between the feature vector z and the memory items \mathbf{M}_i using negative euclidean distance and apply a softmax function to normalize the weights, as shown in Eq. 2.

$$w_i = \frac{\exp\left(-\|z - \mathbf{M}_i\|_2\right)}{\sum_{j=1}^{M} \exp\left(-\|z - \mathbf{M}_j\|_2\right)}, i = \{1, ..., M\}.$$
 (2)

2.1.2 Sparse Memory Addressing

Enforcing sparsity in memory addressing forces the model to approximate the feature vector z using fewer but more relevant memory items. It effectively prevents the model from unexpectedly combining several unrelated memory items to reconstruct defects. Moreover, it implicitly performs memory selection [29], allowing us to save computation by removing items in the memory that were never accessed when reconstructing the training data.

Let superscript $w^{(i)}$ denote an ordinal rank indexing of the elements of w, where $w^{(1)} > w^{(2)} > ... > w^{(M)}$. We compute a sparse approximation \hat{w} of the weight vector wthat corresponds to getting the k closest memory items followed by a re-nomalization step, as shown in Eq. 3, where $\mathbb{1}$ is the indicator function that returns a value of 1 if the condition inside is true and 0 otherwise.

$$\hat{w}^{(i)} = \frac{w^{(i)} \mathbb{1}\{i \le k\}}{\sum_{i} w^{(j)} \mathbb{1}\{j \le k\}}$$
(3)

Since we are only using a select few memory items for reconstruction, it is desirable to prevent the model from learning redundant memory items. Therefore, we impose a margin between the closest memory item $\mathbf{M}^{(1)}$ and second closest memory item $\mathbf{M}^{(2)}$ with respect to the input latent vector z.



Figure 2: Illustration of the trust region memory updates.

$$L_{margin} = \left[\|z - \mathbf{M}^{(1)}\|_2 - \|z - \mathbf{M}^{(2)}\|_2 + 1 \right]_+ \quad (4)$$

2.2. Trust Region Memory Updates

Without the assumption that the training data set only contains normal samples, the memory auto-encoder will treat defective samples as normal and learn to store defect features into the memory, leading to poor defect detection performance. We leverage on two key observations to prevent defective samples from contaminating the memory: (1) Defects are rare and they do not always appear in the same location, which means that the proportion of defects at the patch level will be significantly smaller than the proportion of defects at the image level. (2) Normal data have regularity in appearance, making it easier for the memory autoencoder to reconstruct normal as compared to defects during the early stages of training. This implies that normal features are initially mapped closer to the memory items than defective features.

Based on these insights, we update memory items based on a specified trust region that pulls the features within the region towards the closest memory item and push the features outside the region away from the memory item, as illustrated in Figure 2. More specifically, we split the feature space based on a distance radius with respect to each memory item in M, as shown in Eq. 5. Everything within δ_1 radius are considered normal features that we want to pull closer, while everything outside are considered potential defects that we want to push farther. To prevent the model from pushing the defective features to infinity, we ignore everything farther than a predefined trust threshold δ_2 .

$$r(z, \mathbf{M}^{(1)}) = \begin{cases} 1 & \|z - \mathbf{M}^{(1)}\|_2 \le \delta_1 \\ -1 & \delta_1 < \|z - \mathbf{M}^{(1)}\|_2 \le \delta_2 \\ 0 & \text{otherwise} \end{cases}$$
(5)

We implement the trust region updates as an additional loss function defined in Eq. 6, where $\mathbf{M}^{(1)}$ denotes the closest memory item in \mathbf{M} with respect to z.

$$L_{trust} = r(z, \mathbf{M}^{(1)}) \| z - \mathbf{M}^{(1)} \|_2$$
(6)

Since patches that are easy to reconstruct tend to to have smaller distances to the memory slots than patches that are harder to reconstruct, we require δ_1 to be adaptive to each of these cases. We set δ_1 to be the average distances between memory item \mathbf{M}_i with the features z_i in the current batch that retrieves \mathbf{M}_i as its closest memory item. Our main intuition is that, assuming normal features are abundant and are similar to each other, then normal features will mostly be pulled closer to the memory and only occasionally pushed out, but defect features will always be pushed out since they would always be farther than the average distance, thus, preventing contamination of the memory items.

2.2.1 Loss Functions

Following recent works on image synthesis [26], we employ several loss functions to encourage sharper images: Reconstruction loss L_{rec} [26], SSIM loss L_{sm} [3], VGG feature loss L_{vgg} [18], GAN loss L_{GAN} [13], and GAN feature loss L_{feat} [42, 43]. Please refer to the supplementary materials for more details on these loss functions. The total loss function is then defined in Eq. 7, where the λ co-efficients are hyper-parameters that control the relative weighting of each term. We used similar weightings as [26].

$$L_{total} = \lambda_{rec} L_{rec} + \lambda_{sm} L_{sm} + \lambda_{vgg} L_{vgg} + \lambda_{GAN} L_{GAN} + \lambda_{feat} L_{feat} + \lambda_{margin} L_{margin} + \lambda_{trust} L_{trust}$$
(7)

2.3. Spatial Perceptual Distance

The difference between the input and its normal reconstruction dictates whether it will be considered a defect or not. Ideally, we want a distance measure that assigns high scores on defective regions while assigning low scores on normal regions. Previous approaches [2, 3, 12, 25] use shallow pixel-wise distances such as mean squared error, which is sensitive to slight global shifts and small reconstruction inaccuracies. As a result, they do not work well when our focus is on texture pattern similarity rather than exact pixel alignments. Therefore, we use a spatial perceptual distance [45] that computes differences on a patch level using deep features extracted from a pre-trained convolutional neural network. Computing deep features allows us to incorporate high-level information in measuring distance, which has been shown to correlate well with human perception [45], making it a more suitable distance measure. Moreover, operating on a patch level allows us to capture texture information that are impossible to get from shallow pixel-wise distances such as the mean squared error that only considers pixel intensities.

Let $\phi_l(x) \in \mathbb{R}^{H_l \times W_l \times C_l}$ denote the features extracted from the *l*th layer of a convolutional neural network. The perceptual distance $PD(x, \hat{x}) \in \mathbb{R}^{H \times W}$ is defined as

$$PD(x,\hat{x}) = \sum_{l} \gamma_{l}^{T} g\big(\phi_{l}(x) - \phi_{l}(\hat{x})\big)^{2}, \qquad (8)$$

where g is a bilinear up-sampling function that resizes the feature tensors to $H \times W$ and $\gamma_l \in \mathbb{R}^{C_l}$ is a weight vector used to aggregate the channel-wise feature differences at layer l. Note that, for brevity, we abuse the notation of dot product to refer to the channel-wise aggregation, which can be implemented efficiently as a 1×1 convolution. We adopt the weights of [45] for γ_l and ϕ , which were optimized for the Berkeley Adobe Perceptual Patch Similarity (BAPPS) dataset.

We define the reconstruction error as a multiplicative combination of the mean squared error and the perceptual distance, which further highlights defective regions and suppress the non-defective regions.

3. Experimental Results

We demonstrate the effectiveness of our proposed method on a real-world **MVTec** [2] dataset consisting of 15 different industrial products. It contains over 70 different types of defects labeled with both defect types and segmentation masks. We use the original train-test splits to compare against prior works in a noise-free setting. But for the noisy setting, we re-sample half of the defect images as part of our training data since the original training set does not contain any defective image. The number of normal and defective images in the new training set are 3629 and 610, respectively. In the testing set, there are 467 normal images and 648 defective images.

Table 1: Image-based comparison in terms of AUC with several state-of-the-art baselines on the MVTec Dataset. (Best score in bold and the second best score in bold italic.)

Method	mean AUC	boule	olde?	capsule	cal per	^{8njd}	hazelhur	leather	metal nut	piţ	screw.	tije	toothbush	transistor	Doo ₄₁	² ipper
GeoTrans [11]	67.23	74.40	67.00	61.90	84.10	63.00	41.70	86.90	82.00	78.30	43.70	35.90	81.30	50.00	97.20	61.10
GANomaly [1]	76.15	89.20	73.20	70.80	84.20	74.30	79.40	79.20	74.50	75.70	69.90	78.50	70.00	74.60	65.30	83.40
ARNet [16]	83.93	94.10	68.10	88.30	86.20	78.60	73.50	84.30	87.60	83.20	70.60	85.50	66.70	100	100	92.30
f-Ano-GAN [33]	65.85	91.36	76.37	72.79	56.57	59.63	63.15	62.50	59.70	64.07	50.00	61.34	67.31	77.92	75.00	50.00
MemAE [12]	81.85	86.19	56.65	75.71	74.55	98.91	97.39	95.48	53.12	77.88	83.59	79.11	96. 78	71.62	97.11	83.71
TrustMAE-noise free (max)	90.78	96.98	85.06	78.82	97.43	99.08	98.50	95.07	76.10	83.31	82.37	97.29	96.94	87.50	99.82	87.45
TrustMAE-noise free (mean)	88.04	98.57	91.64	79.34	87.40	92.73	93.36	87.33	/4.14	92.12	61.94	89.50	90.28	91.71	96.32	94.22

Table 2: Pixel-based comparison in terms of Area Under the Receiver Operating Curve (AUC) with several state-of-the-art baselines on the MVTec Dataset. (Best score in bold and the second best score in bold italic.)

			_	ue	ž		hur	4	nut		۵		brush	Stor.		~
Method	mean AUC	Polil	cable	Caps,	Carpo	Brid	hard	leath	merda,	llid	Screy	tile	tooth	trans,	000 ⁴¹	iq.
AE-L2 [3]	80.40	90.90	73.20	78.60	53.90	96.00	97.60	75.10	88.00	88.50	97.90	47.60	97.10	90.60	63.00	68.00
AE-SSIM [3]	81.83	93.30	79.00	76.90	54.50	96.00	96.60	71.00	88.10	89.50	98.30	49.60	97.30	90.40	64.10	82.80
MemAE [12]	85.74	85.04	71.17	92.95	81.16	95.56	97.15	92.91	78.99	93.25	95.57	70.76	94.77	66.73	85.44	84.65
Towards Visually Explaining [22]	86.07	87.00	90.00	74.00	78.00	73.00	98.00	95.00	94.00	83.00	97.00	80.00	94.00	93.00	77.00	78.00
CNN Feature Dictionary [23]	78.07	78.00	79.00	84.00	72.00	59.00	72.00	87.00	82.00	68.00	87.00	93.00	77.00	66.00	91.00	76.00
AnoGAN [34]	74.27	86.00	78.00	84.00	54.00	58.00	87.00	64.00	76.00	87.00	80.00	50.00	90.00	80.00	62.00	78.00
AE-SSIM Grad [6]	86.38	95.10	85.90	88.40	77.40	98.00	96.60	60.20	92.00	92.70	92.50	62.60	98.40	93.40	73.80	88.70
γ -VAE Grad [6]	88.77	93.10	88.00	91.70	72.70	97.90	98.80	89.70	91.40	93.50	97.20	58.10	98.30	93.10	80.90	87.10
AE-L2 Grad [6]	88.77	91.60	86.40	95.20	73.40	98.10	98.40	92.10	89.90	91.20	98.00	57.50	98.30	92.10	80.50	88.90
VAE Grad [6]	89.29	92.20	91.00	91.70	73.50	96.10	97.60	92.50	90.70	93.00	94.50	65.40	98.50	91.90	83.80	86.90
TrustMAE-noise free	93.94	93.39	92.85	87.42	98.53	97.45	98.51	98.05	91.76	89.90	97.63	82.48	98.10	92.72	92.62	97.76

3.1. Implementation Details

We train our model on unlabeled training data for 100 epochs using Adam optimizer with a learning rate of 0.0001 for the generator and 0.0004 for the discriminator. We set the batch size to 16, the number of memory slots to M = 64, and the latent dimension Z = 512. As shown by Gong et al. [12], the number of memory slots M have little effect in performance as long as it is large enough. Based on trial and error, using top k = 3 memory slots work well. We set δ_1 to be the mean distances between the features and memory item and $\delta_2 = 20$. As we show in Section 3.6, the performance of the model is robust to the setting of δ_2 . We adopt a symmetric architecture for our encoder-decoder networks with three residual blocks in the bottleneck layer on both the encoder and decoder networks. We use five down-sampling layers in the encoder, except on grid, screw, pill, and capsule where we only use three down-sampling layers. We defer the discussion on this design choice to Section 3.7. Following previous works, we evaluate the models using the area under the receiver operating characteristic curve (AUC).

3.2. Varying Noise Levels

We conduct experiments for defect classification on increasing noise levels by introducing more defects into the training data. For defect classification, we need to define a single defect score for an input image to indicate the likelihood of it being defective. We define our defect score as the spatial pooling of the reconstruction errors computed with the perceptual distance. Since there are limited number of defective images, we reduced the number of normal data instead to match the desired noise percentage for experiments with higher noise (defect) level (30% - 40%). Figure 4 shows defect classification results at varying noise levels. For auto-encoders trained using mean squared error (AE-MSE) and SSIM (AE-SSIM), their performance significantly drop as we increase the noise levels. Autoencoders are sensitive to noise since defect information can be directly passed from the encoder to the decoder, thus,



Figure 3: Visual comparison of defect segmentation results with state-of-the-art baselines. Ground-truth is represented by the red contour, and predicted defect segmentation is represented by the green overlay.



Figure 4: Comparison with baselines at different noise percentages.

easily reconstructing defects. MemAE [12] also does not perform well, because without enough sparsity and filtering mechanism, it could still memorize and reconstruct defects. Building on top of MemAE, our model enhances the noise robustness of the memory through our trust region memory updates, and uses a more suitable perceptual distance, thus, improving the overall performance, even at high noise levels.

3.3. Comparison with Baselines

To better compare with prior works, we report the performance of our model using spatial max pooling and mean pooling of the reconstruction errors as our defect score, which are the two most common pooling methods. We use the original train-test splits and train our model on only normal data (noise-free) to maintain a fair comparison. Table 1 shows our defect classification performance measured in terms of image level AUC in comparison with several recent works on defect detection [1,11,12,16,33]. Our model with max pooling achieves the best AUC score on majority of the classes as well as the best mean AUC overall.

We analyzed the classes with lower performance (e.g. metal nut, pill, screw, and capsule), and observed that the model have difficulties modeling large variations in the normal data due to large rotations or randomly appearing prints. This is because the sparse memory addressing combined with the trust region memory updates restricts the variations that the memory can capture. Note that we can improve the performance by adding a warping module to align the data but it is outside the scope of this paper.

We also compared against prior works in the defect segmentation setting, which is evaluated using pixel-level AUC. To get a segmentation map, we can directly impose a threshold on the reconstruction errors. Table 2 shows our defect segmentation performance in comparison with several recent works on defect segmentation [3, 5, 6, 12, 19, 22, 23, 34]. We can observe that our method consistently performs competitively across all classes, achieving the best mean AUC among all the baselines.

We also show visual results on defect segmentation in Figure 3. Our model is able to capture a more complete segmentation of the defects as compared with the baselines, especially on textured objects. This is due to our spatial perceptual distance that can better measure texture similarity as opposed to shallow distance metrics that the baselines used.

3.4. Ablation Study

We conduct an ablation study to evaluate how each of our design choices affect the performance of the model. Figure 5 shows the defect classification results in terms of mean AUC across all classes at varying noise levels.

Trust Region Memory Updates. In the low noise regime (5% - 10% noise levels), we can observe similar



Figure 5: Ablation studies of our model trained with different levels of noisy data.



Figure 6: When the model is trained without trust region (TR), it could hallucinate defects when reconstructing from either normal data (left) or defective (right) data when training data are polluted with noises.

performance between the models with and without trust region memory updates (TR). However, the performance gap starts to widen as the noise levels increase. We observe that at higher noise levels, there are more chances for the memory slots to be contaminated with defective features. This causes the model to hallucinate defects (as shown in Figure 6 left) or reconstruct defects well (as shown in Figure 6 right), leading to poorer performance. Our proposed trust region memory updates (TR) makes the memory more robust to defect contamination, leading to better normal reconstructions, thus, better defect detection performance.

Perceptual Distance. Without the perceptual distance, we can observe a significant drop in performance across all noise levels. As we show in Figure 7, shallow pixel-wise distances, such as MSE, are inadequate for defect detection since they produce lots of noise across the whole image due to slight shifts and differences in textures between the input image and its reconstruction, leading to an increased number of false positives. Structural similarity index measure (SSIM) accounts for patch information but it normalizes distances independently for each patch, which makes it hard to separate defects since the error values are close to each other. Moreover, SSIM only accounts for the low



Figure 7: Visual comparison of the spatial perceptual distance (PD) with commonly used distance metrics such as mean squared error (MSE), and structural similarity index measure (SSIM).



Figure 8: Memory Access with and without sparse addressing.

order statistics (means and covariances) of patch intensities, which are inadequate for defects that have very similar patch statistics with normal images such as the tile and carpet example shown in the bottom two rows of Figure 7. In contrast, we use deep features to compute the differences, allowing us to capture textures and other high level information. We can observe that the perceptual distance assigns noticeably higher values on defective regions and lower values on normal regions, making defects more pronounced and easier to separate.

3.5. Memory Access Patterns

We visualize memory access patterns with and without sparse addressing in Figure 8 for a single item z_i (left), and aggregated access frequencies across all the data (right). Without any restriction, the model tends to spread the information across all memory slots. This is because using



Figure 9: Effect of varying the trust threshold δ_2 .

more memory slots gives more flexibility for reconstruction. This in itself is not detrimental if the model is trained on only normal data. However, if the training data contains defective images then defect features will also spread to all the memory slots, making it harder to remove them. This is reflected in the model's poorer performance without sparse addressing (SP), as shown in Figure 5. In contrast, sparse memory addressing forces the model to rely only a few memory slots. As a consequence, defect features can only affect a few memory slots, making it easier to prevent the memory from being contaminated with defect features.

3.6. Effect of Trust Threshold δ_2

Figure 9 shows the results of varying δ_2 on cable at 30% noise level. We can observe that the model is quite robust in the range $\delta_2 = [15, 40]$, achieving approximately 88% AUC. Setting $\delta_2 = 0$ is equivalent to ignoring everything outside the trust region, which performs worse since there is nothing encouraging the model to map defect features far from the memory items. We also tried removing δ_2 , but the model becomes unstable since it will just keep pushing the feature vectors to infinity.

3.7. Effect of Down-Sampling Layers

The number of down-sampling layers control the receptive field, and thus, the patch size that the memory slots operate on. More down-sampling corresponds to larger patch sizes, and vice versa. We observe that using fewer downsampling achieves better results for classes that have many small details that greatly vary across different instances (e.g. grid, screw, pill, and capsule). This is because the memory combined with the trust region updates has the effect of limiting the variations of patch features that the memory can store. Without enough variations captured by the memory slots, it becomes difficult for the network to reconstruct varying small details accurately using limited set of large patch features.

4. Related Works

Techniques for defect detection can be broadly grouped into: classification-based [11,24,32,35,36], detection-based [10, 37], segmentation-based [7, 10, 17, 20, 28, 38–40], and reconstruction-based [3, 5, 6, 23, 25, 34, 41, 46].

Our method belongs to the reconstruction-based, which relies on auto-encoders trained on normal data to reconstruct an image that is close to the normal data. It then classifies defects based on their reconstruction error. Schlegl et al. [34] and Chen et al. [4] combined auto-encoders with generative adversarial networks (GAN) [13] to better model the distribution of normal images. Likewise, Akcay et al. [1] used a conditional GAN and computed defect scores using the discriminator. Schlegl et al. [33] detected defects by combining discriminator feature error and image reconstruction error.

To prevent the model from reconstructing defects, Haselmann et al. [15] cropped out holes and inferred the pixels in the hole. Gong et al. [12] augmented auto-encoders with memory and reconstructs inputs based on features stored in its memory. Ye et al. [16] restored original data by forcing the network to learn high-level feature embeddings of selected attributes. While existing approaches work well under the assumption that training data only contains normal samples, they tend to perform poorer when training data are polluted with noise, which common in reality. We propose trust region memory updates to further increase model's resistance to pollution in training data.

Other than modeling normal images, different ways to compute distances are also gaining focus. Several works [2,3,12,25] used shallow pixel-wise distances such as MSE. However, they are sensitive to slight shifts, making it unsuitable for texture similarity. Bergmann et al. [3] used SSIM, while Liu et al. [22] used gradient attention maps computed from image embeddings to segment defects. In contrast, we propose to use spatial perceptual distance that computes patch-wise differences of deep features.

5. Conclusion

We presented TrustMAE, a noise-resilient and unsupervised framework for product defect detection. It addresses the common defect data collection problem using unlabeled images while enabling a mixture of defective images within the training dataset. Our framework compares favorably to state-of-the-art baselines, and we showed the effectiveness of its components via an extensive ablation study.

We believe the noise-resilience property is a novel and useful one because the normal-defective classification of the training images can change according to practical application requirements. While most anomaly detection frameworks require noise-free datasets, we hope to gain more insights into these methods and create a noise-robust framework that performs well across all product classes soon.

References

- Samet Akcay, Amir Atapour-Abarghouei, and Toby P Breckon. Ganomaly: Semi-supervised anomaly detection via adversarial training. In *Asian Conference on Computer Vision*, pages 622–637. Springer, 2018.
- [2] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Mvtec ad–a comprehensive real-world dataset for unsupervised anomaly detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9592–9600, 2019.
- [3] Paul Bergmann, Sindy Löwe, Michael Fauser, David Sattlegger, and Carsten Steger. Improving unsupervised defect segmentation by applying structural similarity to autoencoders. In Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP), pages 372–380, 2019.
- [4] Xiaoran Chen and Ender Konukoglu. Unsupervised detection of lesions in brain mri using constrained adversarial auto-encoders. In *International Conference on Medical Imaging with Deep Learning, MIDL 2018, 4-6 July 2018, Amsterdam, Netherlands*, 2018.
- [5] Bin Dai and David Wipf. Diagnosing and enhancing vae models. *arXiv preprint arXiv:1903.05789*, 2019.
- [6] David Dehaene, Oriel Frigo, Sébastien Combrexelle, and Pierre Eline. Iterative energy-based projection on a normal data manifold for anomaly localization. *arXiv preprint arXiv:2002.03734*, 2020.
- [7] Xinghui Dong, Chris J Taylor, and Tim F Cootes. Small defect detection using convolutional neural network features and random forests. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018.
- [8] Vincent Dumoulin, Ishmael Belghazi, Ben Poole, Alex Lamb, Martín Arjovsky, Olivier Mastropietro, and Aaron C. Courville. Adversarially learned inference. In 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings, 2017.
- [9] Max Ferguson, Ronay Ak, Yung-Tsun Tina Lee, and Kincho H Law. Automatic localization of casting defects with convolutional neural networks. In 2017 IEEE International Conference on Big Data (Big Data), pages 1726–1735. IEEE, 2017.
- [10] Max K Ferguson, Ak Ronay, Yung-Tsun Tina Lee, and Kincho H Law. Detection and segmentation of manufacturing defects with convolutional neural networks and transfer learning. *Smart and sustainable manufacturing systems*, 2, 2018.
- [11] Izhak Golan and Ran El-Yaniv. Deep anomaly detection using geometric transformations. In Advances in Neural Information Processing Systems, pages 9758–9769, 2018.
- [12] Dong Gong, Lingqiao Liu, Vuong Le, Budhaditya Saha, Moussa Reda Mansour, Svetha Venkatesh, and Anton van den Hengel. Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection. arXiv preprint arXiv:1904.02639, 2019.
- [13] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and

Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

- [14] Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. arXiv preprint arXiv:1410.5401, 2014.
- [15] Matthias Haselmann, Dieter P Gruber, and Paul Tabatabai. Anomaly detection using deep learning based image completion. In 2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA), pages 1237– 1242. IEEE, 2018.
- [16] Chaoqin Huang, Fei Ye, Jinkun Cao, Maosen Li, Ya Zhang, and Cewu Lu. Attribute Restoration Framework for Anomaly Detection. *arXiv e-prints*, page arXiv:1911.10676, Nov. 2019.
- [17] Yibin Huang, Congying Qiu, Yue Guo, Xiaonan Wang, and Kui Yuan. Surface defect saliency of magnetic tile. In 2018 IEEE 14th International Conference on Automation Science and Engineering (CASE), pages 612–617. IEEE, 2018.
- [18] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.
- [19] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings, 2014.
- [20] Jie Lei, Xin Gao, Zunlei Feng, Huamou Qiu, and Mingli Song. Scale insensitive and focus driven mobile screen defect detection in industry. *Neurocomputing*, 294:72–81, 2018.
- [21] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [22] Wenqian Liu, Runze Li, Meng Zheng, Srikrishna Karanam, Ziyan Wu, Bir Bhanu, Richard J Radke, and Octavia Camps. Towards visually explaining variational autoencoders. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8642–8651, 2020.
- [23] Paolo Napoletano, Flavio Piccoli, and Raimondo Schettini. Anomaly detection in nanofibrous materials by cnn-based self-similarity. *Sensors*, 18(1):209, 2018.
- [24] Vidhya Natarajan, Tzu-Yi Hung, Sriram Vaikundam, and Liang-Tien Chia. Convolutional networks for voting-based anomaly classification in metal surface inspection. In 2017 IEEE International Conference on Industrial Technology (ICIT), pages 986–991. IEEE, 2017.
- [25] Duc Tam Nguyen, Zhongyu Lou, Michael Klar, and Thomas Brox. Anomaly detection with multiple-hypotheses predictions. In *International Conference on Machine Learning*, pages 4800–4809, 2019.
- [26] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [27] Yongzhi Qu, Miao He, Jason Deutsch, and David He. Detection of pitting in gears using a deep sparse autoencoder. *Applied Sciences*, 7(5):515, 2017.

- [28] Domen Racki, Dejan Tomazevic, and Danijel Skocaj. A compact convolutional neural network for textured surface anomaly detection. In 2018 IEEE Winter Conference on Applications of Computer Vision (WACV), pages 1331–1339. IEEE, 2018.
- [29] Jack Rae, Jonathan J Hunt, Ivo Danihelka, Timothy Harley, Andrew W Senior, Gregory Wayne, Alex Graves, and Timothy Lillicrap. Scaling memory-augmented neural networks with sparse reads and writes. In Advances in Neural Information Processing Systems, pages 3621–3629, 2016.
- [30] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [31] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [32] Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel Müller, and Marius Kloft. Deep one-class classification. In *International Conference on Machine Learning*, pages 4393–4402, 2018.
- [33] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Georg Langs, and Ursula Schmidt-Erfurth. f-anogan: Fast unsupervised anomaly detection with generative adversarial networks. *Medical image analysis*, 54:30–44, 2019.
- [34] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *International Conference on Information Processing in Medical Imaging*, pages 146–157. Springer, 2017.
- [35] Bernhard Schölkopf, Robert C Williamson, Alex J Smola, John Shawe-Taylor, and John C Platt. Support vector method for novelty detection. In Advances in neural information processing systems, pages 582–588, 2000.
- [36] Benjamin Staar, Michael Lütjen, and Michael Freitag. Anomaly detection with convolutional neural networks for industrial surface inspection. *Proceedia CIRP*, 79:484–489, 2019.
- [37] Xiaohong Sun, Jinan Gu, Rui Huang, Rong Zou, and Benjamin Giron Palomares. Surface defects recognition of wheel hub based on improved faster r-cnn. *Electronics*, 8(5):481, 2019.
- [38] Domen Tabernik, Samo Šela, Jure Skvarč, and Danijel Skočaj. Segmentation-based deep-learning approach for surface-defect detection. *Journal of Intelligent Manufacturing*, pages 1–18, 2019.
- [39] Daniel Stanley Tan, Robert Neil Leong, Ann Franchesca Laguna, Courtney Anne Ngo, Angelyn Lao, Divina Amalin, and Dionisio Alvindia. A framework for measuring infection level on cacao pods. In 2016 IEEE Region 10 Symposium (TENSYMP), pages 384–389. IEEE, 2016.
- [40] Daniel Stanley Tan, Robert Neil Leong, Ann Franchesca Laguna, Courtney Anne Ngo, Angelyn Lao, Divina M Amalin, and Dionisio G Alvindia. Autodidac: Automated tool for dis-

ease detection and assessment for cacao black pod rot. *Crop Protection*, 103:98–102, 2018.

- [41] Xian Tao, Dapeng Zhang, Wenzhi Ma, Xilong Liu, and De Xu. Automatic metallic surface defect detection and recognition with convolutional neural networks. *Applied Sciences*, 8(9):1575, 2018.
- [42] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018.
- [43] Xiangyu Xu, Deqing Sun, Jinshan Pan, Yujin Zhang, Hanspeter Pfister, and Ming-Hsuan Yang. Learning to superresolve blurry face and text images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 251–260, 2017.
- [44] Yong-Ho Yoo, Ue-Hwan Kim, and Jong-Hwan Kim. Convolutional recurrent reconstructive network for spatiotemporal anomaly detection in solder paste inspection. arXiv preprint arXiv:1908.08204, 2019.
- [45] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018.
- [46] Bo Zong, Qi Song, Martin Renqiang Min, Wei Cheng, Cristian Lumezanu, Daeki Cho, and Haifeng Chen. Deep autoencoding gaussian mixture model for unsupervised anomaly detection. In *International Conference on Learning Repre*sentations, 2018.