



Article

Autoencoder-Based Visual Anomaly Localization for Manufacturing Quality Control

Devang Mehta and Noah Klarmann *

Faculty of Management and Engineering, Rosenheim Technical University of Applied Sciences,
Hochschulstraße 1, 83024 Rosenheim, Germany

* Correspondence: noah.klarmann@th-rosenheim.de

Abstract: Manufacturing industries require the efficient and voluminous production of high-quality finished goods. In the context of Industry 4.0, visual anomaly detection poses an optimistic solution for automatically controlled product quality with high precision. In general, automation based on computer vision is a promising solution to prevent bottlenecks at the product quality checkpoint. We considered recent advancements in machine learning to improve visual defect localization, but challenges persist in obtaining a balanced feature set and database of the wide variety of defects occurring in the production line. Hence, this paper proposes a defect localizing autoencoder with unsupervised class selection by clustering with k-means the features extracted from a pretrained VGG16 network. Moreover, the selected classes of defects are augmented with natural wild textures to simulate artificial defects. The study demonstrates the effectiveness of the defect localizing autoencoder with unsupervised class selection for improving defect detection in manufacturing industries. The proposed methodology shows promising results with precise and accurate localization of quality defects on melamine-faced boards for the furniture industry. Incorporating artificial defects into the training data shows significant potential for practical implementation in real-world quality control scenarios.

Keywords: anomaly detection; artificial defect simulation; autoencoder; computer vision; defect detection; defect localization; deep learning; deep neural network; deep neural network-based defect detection; feature extraction; Industry 4.0; unsupervised clustering; manufacturing quality control; machine vision; unsupervised class selection; unsupervised learning; visual inspection systems; visual product quality control



Citation: Mehta, D.; Klarmann, N. Autoencoder-Based Visual Anomaly Localization for Manufacturing Quality Control. *Mach. Learn. Knowl. Extr.* **2024**, *6*, 1–17. <https://doi.org/10.3390/make6010001>

Academic Editors: Antonio Fernández-Caballero and Byung-Gyu Kim

Received: 11 September 2023
Revised: 4 December 2023
Accepted: 15 December 2023
Published: 21 December 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Artificial intelligence (AI) promises to be a revolutionary force in the 21st century. It has gained significant attention across various sectors, with extensive research, development, and production of AI-driven products and services. The widespread adoption, ease of use, and flexibility of AI technologies have propelled its evolution. This research aims to contribute this revolutionary wave to visual inspection in furniture manufacturing.

In the manufacturing industry, the visual inspection of products plays a crucial role at different stages of the production process; ensuring the quality of the final product is essential to meet aesthetic and functional requirements. At the University of Applied Sciences Rosenheim, the [proto_lab](https://www.th-rosenheim.de/en/die-hochschule/labore/proto-lab) (<https://www.th-rosenheim.de/en/die-hochschule/labore/proto-lab>, accessed on 1 December 2023) laboratory is an innovative Industry 4.0 platform that produces furniture with state-of-the-art machinery in a fully digitalized way, posing an ideal ecosystem for applying AI use cases. The goal of applying data-based methodologies is to establish a high-quality, efficient, intelligent system to improve the production cycle holistically. In this context, an integrated autonomous system, particularly computer vision (CV) systems, has proven to be an overly promising way. Detecting

product flaws is denoted as anomaly detection, alternatively also known as artifact, novelty, or outlier detection.

Li et al. [1] presents an overview on machine vision applications in furniture manufacturing from 2011 to 2022. Many/most studies rely on classical methods to perform quality checks. These traditional methodologies are well-established, straightforward, and computationally optimized but usually show limitations. Limitations may occur due to variations in the environment, specific manual step-wise feature engineering, heterogeneous images regarding size and quality, and complexity of the image data. While widely utilized for elementary implementations and tasks, classical methods cannot keep pace with technological advancements in camera sensor resolution, optics, and deep learning (DL) techniques [2]. DL methods leverage large amounts of data, requiring minimal expertise and automatic fine-tuning. DL demonstrates flexibility to adapt to different domains and datasets, even when the relevant data are limited—e.g., by making use of transfer learning [3]. The dataset's images are sharp and consistent in feature space in the scope of the problem. Classical CV methods often require algorithmic computation for each type of feature or class, making them expensive to implement. Anomaly localization, however, can be efficiently scaled up by leveraging the advantages of neural networks. Several DL methodologies have shown remarkable predictive performance. In some cases, hybrid approaches might be the best choice when combining traditional CV methods and DL techniques.

As manufacturing systems excel in optimization, productivity, and efficiency, the number of products having defects reduces enormously. Consequently, current research and development efforts are in unsupervised anomaly localization, with generative deep neural networks gaining significant prominence. Initially, a popular approach was to train models using non-anomalous classes and predict classes of anomalous and non-anomalous instances. However, this methodology requires additional information for generating or reconstructing non-anomalous images. Modern challenges include accurately localizing anomalies within a low-variance feature set in an image dataset. Achieving an intelligent, thorough, fast, robust, and reliable CV system necessitates the integration of cutting-edge technologies such as deep learning and 3D point cloud analysis in some cases where depth information is required [1].

This paper presents a hybrid approach for localizing surface defects on melamine-faced boards. In contrast to directly focusing the camera on specific dimensions of boards, our image dataset consists of high-resolution images captured with a camera having a fixed field of view, allowing inspection of boards of various proportions. Achieving such a model is performed by slicing the high-resolution image and selecting classes of interest from an imbalanced dataset with feature extraction and k-means clustering in order to consider the minute variation in the frequency of the features, and finally, simulating anomalies for the autoencoder model to predict all artifacts.

The subsequent sections delve deeper into the literature, methodology, results, and implications of this research, providing an understanding of the approaches utilized in localizing anomalies in furniture manufacturing.

2. Technical Background

In the following, we present an overview of the core methodologies and principles that underpin our research. For readers well-versed in these methodologies, this section may be familiar and can skip to the next section without losing continuity in the paper. However, for those less acquainted with these concepts, this section serves as a short introduction, illuminating the methods and tools utilized throughout our work.

2.1. K-Means Clustering

The unsupervised K-means algorithm, as introduced by Hartigan and Wong [4], divides M points in N dimensions into K clusters in a way that within every cluster, the sum of squares are minimized. Typically, K-means executes as a preprocessing step before the

start of the main algorithm. Generally, clustering has a somewhat simple implementation, and the algorithm is guaranteed to converge. It easily scales up to adapt to new data with mini-batches to save computation time, cluster merging when the new data clusters based on existing centroids are close, centroid initialization by utilizing existing centroids, online k-means for continuous data, or incremental principal component analysis (PCA) when the dimensionality of the data is very high. The k-means algorithm degrades in performance when the data contains several outliers, resulting in incorrect clustering. Because the k-means algorithm works by minimizing the sum of the squared distance between the data point and the centroid, some data would be clustered tightly in high-density regions, while other regions would have data spread farther apart. Hence, it is worth noting that the k-means algorithm must be generalized for robust performance when the data are of varying density.

2.2. Segmentation

Image segmentation aims to simplify and represent an image by dividing it into multiple regions denoted as segments. Each segment represents pixels that share the same features. This process helps to analyze the image data more meaningfully in a spatial sense. Unlike clustering, segmentation also considers boundaries and structures. Various segmentation algorithms use specific conditions to segment the image. Felzenszwalb and Huttenlocher [5] introduced a fast algorithm that generates segments based on a boundary that separates the regions. By applying Felzenszwalb's segmentation with $scale = 25$, $sigma = 1$, and $min_size = 500$ to an image of braid from DTD Figure 1a, the algorithm can find segments as shown in Figure 1b.

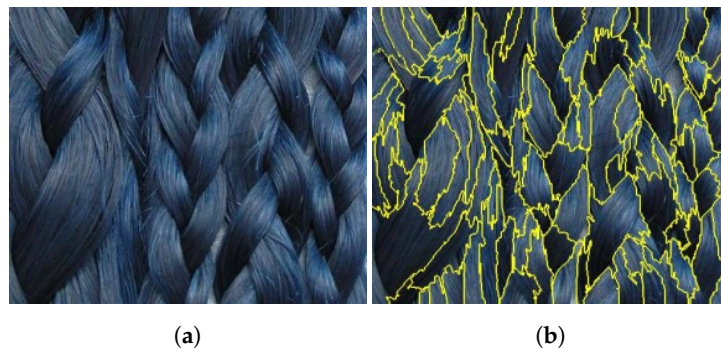


Figure 1. DTD (a) image of braid, and (b) its Felzenszwalb segmented boundaries.

2.3. SSIM

Structural similarity (SSIM), introduced by Wang et al. [6], is a commonly used algorithm to determine the similarity or difference between two images. The algorithm is designed specifically for grayscale images, considering their inherent properties such as luminance, contrast, and structure. It adopts mechanisms of human vision for effectively identifying structural information. The SSIM score ranges from zero to one, while unity indicates perfect similarity and zero indicates complete dissimilarity. For example, the SSIM score between the original image Figure 2a and its blurred version Figure 2b with a kernel of (99, 99) is 0.9646.

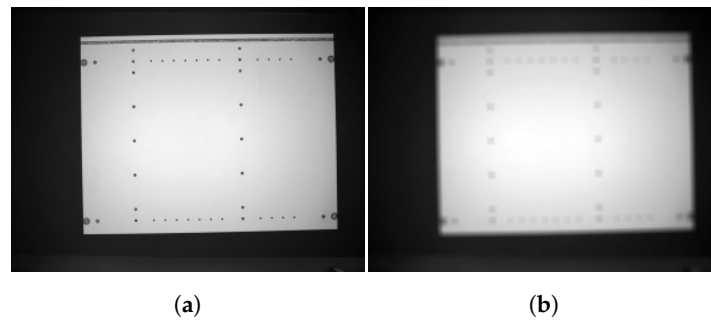


Figure 2. Comparing melamine-faced board sample image: (a) actual vs. (b) blurred image.

2.4. Autoencoders

The generative model autoencoder is a class of unsupervised learning algorithms in which the output shape is the same as the input shape [7]. Such a model allows the network to learn basic representations in a better way when compared to raw, unprocessed data, thereby learning features and patterns while ignoring noise. The network has an encoder that maps the information to a latent representation. Following it is decoding the latent space to reconstruct the original data. The model optimizes by minimizing the MSE loss between the target and the reconstructed image. Usually, the convolutional autoencoders (CAEs) are not very promising for localizing anomalies; using a denoising autoencoder (DAE) with an SSIM-based loss typically increases the performance of CAE.

2.5. Generative Adversarial Networks

Generative adversarial networks (GANs), as described by Goodfellow et al. [8], are generative models based on game theory, where the players (the generator and the discriminator) each try to beat the other through strategic adjustments. In the context of image data, the model works competitively, where a generator creates images from random noise that resemble real images, and the discriminator distinguishes between the real and the generated images based on a probability score. In short, GANs learn the probability distribution of the data to generate synthetic data. GAN-based methods fail where the discriminator gets stuck in a local minimum, and the generators produce a particular output repeatedly, making it difficult to reliably reconstruct anomaly-free images, especially textures [9].

2.6. U-Net

U-net, as introduced by Ronneberger et al. [10], was initially designed for segmentation in the field of medical imaging. It depends on data augmentation and can be trained even with a couple of images. The main difference between autoencoder and u-net architecture is the implementation of skip connections. Skip connections help the network maintain high-resolution features lost during downsampling.

2.7. Feature Extraction with Pre-Trained CNN

A deep neural network can extract feature descriptors for a completely different dataset. This approach is known as feature extraction. Figure 3 depicts the architecture of the popular VGG16 [11] deep neural network architecture. The first part of the fully connected layer of the model provides the features of the input image. Simple clustering or classification utilizes these feature vectors from multiple dataset images as postprocessing.

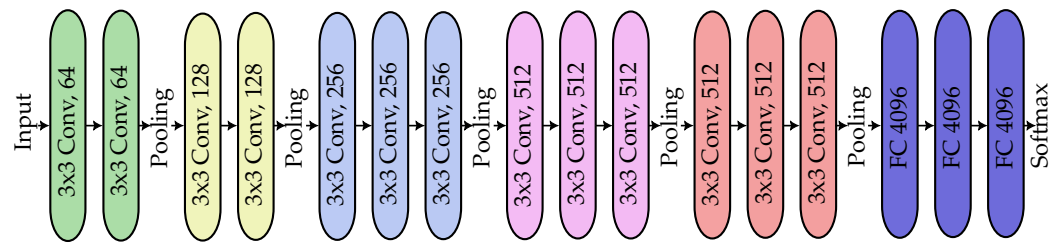


Figure 3. The VGG16 model architecture.

3. Related Work

Anomaly detection tackled with AI methodologies is broadly categorized into supervised, semisupervised, and unsupervised models. The lack of anomaly ground-truth values in an image dataset is the foremost cause for researchers to employ unsupervised anomaly localization models. Due to the lack of data on anomalous classes, modern research and development in unsupervised anomaly localization relies on reconstructive or generative deep neural networks. A popular method that follows Goodfellow et al. [12] is to train the model with anomaly-free classes only, and the difference between the reconstruction and the input data would localize the anomalies. The methodology suffers from insufficient information on generating or reconstructing an image of a non-anomalous class.

Based on the criteria of the type of methodology incorporated to try to localize anomalies, we have two subsections that provide a better understanding of the existing research. The studies follow two approaches: one where the input data are encoded in some representation and then fed for training, and the other where the focus is mainly on identifying the difference with the reconstruction.

3.1. Encoding-Based Anomaly Localization

The presence of knots in wood plays a crucial role in assessing the eventual quality/strength of the end product. Kamal et al. [13] introduces a unique technique, employing feedforward backpropagation neural networks with Laws texture energy measures (LTEM) [14] as input parameters. This innovative approach aims to predict knot defects in wood through a supervised classification model. While the model performs well for multiclass classification, it struggles with generalization when dealing with unfamiliar defects. Nakanishi et al. [15] proposes an alternative approach using autoencoders and a weighted frequency domain loss to identify various wood anomalies effectively. The study reveals that the weighted frequency domain loss significantly improves the autoencoder's ability to detect anomalies by emphasizing specific frequency components. However, we need further investigation to assess its effectiveness on real-time datasets and, as mentioned by the authors, on high-frequency components.

3.2. Reconstruction-Based Anomaly Localization

The widely recognized MVTec AD dataset, established by Bergmann et al. [16], serves as the standard benchmark for unsupervised anomaly detection and localization. Many researchers utilize the autoencoder architecture for feature learning on this dataset and employ image inpainting techniques to enhance the robustness of generalized predictions. DRAEM, as introduced by Zavrtnik et al. [17], is a notable model that is a reconstructive-discriminative subnetwork trained on the MVTec AD dataset specifically for visual surface anomaly detection. DRAEM incorporates a Perlin noise generator [18] to produce random shapes of artificial anomalies, while the shape content derives from randomly augmented DTD [19] images. DRAEM outperforms several other methods, achieving impressive mean detection/localization, e.g., an AUROC score of 0.98 and 0.973, respectively, could be achieved. It improves the accuracy of anomaly localization and achieves nearly fully supervised results on surface defect datasets.

Another novel model, named CS-ResNet, proposed by Zhang et al. [20], is introduced for PCB cosmetic defect detection using convolutional neural networks. It addresses issues

related to unbalanced class distribution and misclassification cost by incorporating a cost-sensitive adjustment layer in the standard ResNet [21]. This modification results in higher accuracy and lower misclassification cost compared to Auto-VRS [22].

DeRA, as introduced by Hida et al. [9] and akin to DRAEM, is an unsupervised anomaly detection method tested on the MVTec AD dataset. It leverages Felzenszwalb's graph-based segmentation method [5] on segments of the DTD [19] dataset. These segments are superimposed on non-anomalous images with random transparency, creating a more diverse and complex anomalous dataset. The DeRA neural network combines U-net [10] and autoencoder architectures. The performance of DeRA is notably enhanced by incorporating the neural style transfer (NST) loss [23] using a pre-trained VGG19 [11] network as a discriminator. This loss function, trained on the ImageNet dataset [24], aids in improving the anomaly detection results. DeRA achieves a pixel-wise mean AUROC of 0.97, surpassing methodologies described by Bergmann et al. [16]. However, this method is limited to grayscale images and performs poorly on transparent artifacts.

In contrast to DeRA, Schlüter et al. [25] propose a self-supervision task called natural synthetic anomalies (NSA) using Poisson image editing [26] to generate a wide range of realistic synthetic anomalies for anomaly detection and localization. The NSA architecture, a ResNet-based encoder–decoder, achieves mean image-level and pixel-level AUROC scores of 0.972 and 0.963 on the MVTec AD dataset. While outperforming particular methods that do not use additional data, NSA lacks robustness when dealing with minute anomalies.

To compare different autoencoder models for real-time anomaly detection, Mujkic et al. [27] evaluated the following three models: the denoising autoencoder (DAE) [28], semisupervised autoencoder (SSAE), and variational autoencoder (VQ-VAE) [29], against the baseline YOLOv5 [30]. Although YOLOv5 slightly outperformed SSAE in the AUROC score (0.945 vs. 0.8849), SSAE demonstrated better performance in critical cases.

Lastly, Huang et al. [31] introduced a self-supervised masking (SSM) method for anomaly localization. They use random masking to augment each image, creating a diverse set of training triplets [32], enabling the autoencoder to reconstruct masks of various shapes and sizes during training. For inference, a progressive mask refinement method gradually reveals non-anomalous regions and eventually localizes anomalies. On the MVTec AD dataset, the SSM achieves a mean AUROC score of 0.92.

Our research paper leverages the autoencoder architecture, incorporating the autoencoder architecture, an artificial anomaly overlay, and a loss function akin to DeRA. Unlike conventional approaches, our unique pipeline accommodates diverse features within our training dataset. Given the real-world limitations of training high-resolution images directly, we adopt the sliding window technique to capture and comprehend image features effectively. Implementing the sliding window approach, we address memory constraints and augment the richness of our training data. Such an implementation additionally allows our model to be versatile for all dimensions of the boards without an additional camera system. Different from CS-ResNet, we integrate k-means clustering to identify and prioritize non-frequent image features that could be overshadowed by dominant classes, refining the focus of our model and rectifying class imbalances. The resultant balanced training data propels the overall efficacy of our approach. Similar to the methodology in DRAEM, we generate artificial anomalies akin to DeRA using Felzenszwalb's method. Our workflow stands for an innovative and pragmatic solution for advancing anomaly detection within complex image datasets.

4. Methodology

This paper aims to provide a novel solution for localizing artifacts on a melamine-faced board surface subject on the basis of a highly imbalanced dataset. The methodology deals with the frequently arising problem of insufficient data featuring anomalies and requiring a high-resolution image for anomaly identification. The first step in the approach is to capture high-resolution images of the melamine-faced boards. Moreover, we employ the sliding window technique to slide a small window to slice the high-resolution images

into small crops, as detailed in Sections 4.1 and 4.2. The features of the sliced images are extracted with the help of a pre-trained VGG16 and clustered into groups with the k-means algorithm, as described in Section 4.3. As a separate process, Section 4.4 explains the artificial anomaly generation by extracting segments from the images of DTD. Finally, the artificial anomalies are dynamically and randomly overlaid on the fly for training. Figure 4 represents the pipeline of the method, while the following subsections detail the elements.

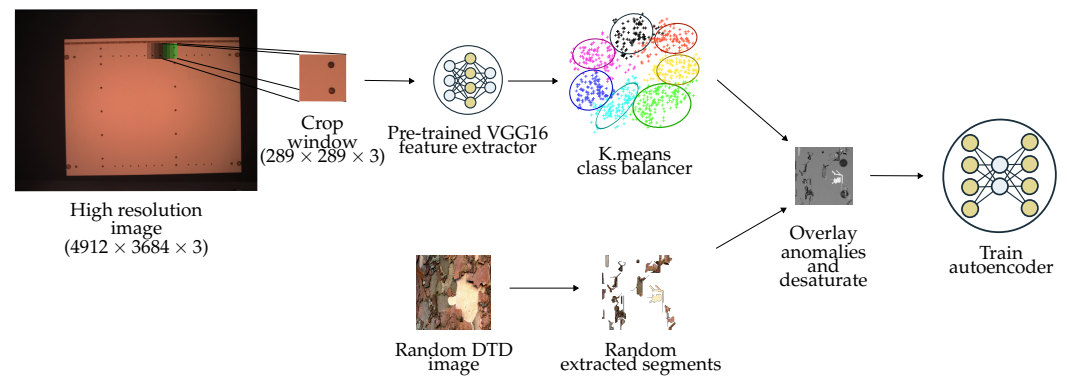


Figure 4. Unsupervised defect localizing training pipeline.

4.1. Imaging System and Dataset Gathering

The demonstrator is operated at controlled conditions, featuring an enclosed lightproof setup utilizing a precise artificial light source, camera position, and focus. A graphical user interface (GUI) running on a Raspberry Pi 4 [33] empowers the system to initiate image capturing when a product sample takes its position beneath the anticipatory camera lens. Our data collection venture encompassed a multitude of captures, each revealing the melamine-faced board's distinctive facets through diverse locations and orientations within the enclosure. The resulting image dataset showcases a lateral viewpoint (Figure 5a) alongside a transverse (Figure 5b). The image dataset presents invaluable insights, culminating in 348 images, for our analysis and model development pursuits.

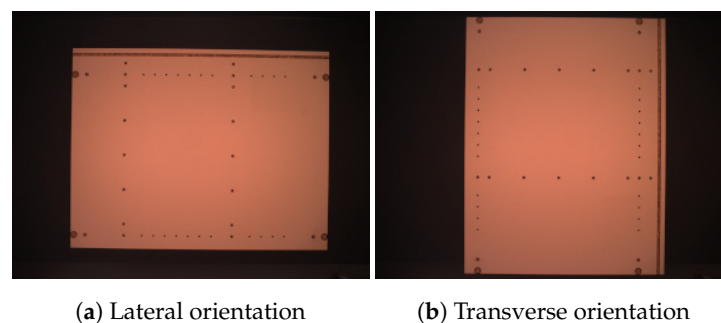


Figure 5. Prototype setup training image.

4.2. Image Dataset Analysis

The demonstrator captures images with a resolution of (4912, 3684). Each melamine-faced board has a barcode label glued on its surface. However, the label feature space resembles the melamine-faced board's surface feature space, presenting a critical issue. Hence, the barcode labels require manual removal beforehand. Seven distinct classes meticulously considered in the images are as follows: (a) background, (b) surface, (c) drilling holes, (d) edges, (e) corners, (f) slots, and (g) combinations of edges, holes, and slots. The board carefully selected for this endeavor is Figure 5 because it included the most diverse classes, with each board exhibiting unique features.

The challenge arose in training the network directly on such high-resolution images. Hence, the crop windows of the image, Figure 6, each of resolution (289, 289), are generated with a precise stride of (67, 97) and zero padding. This approach resulted in 2'520 cropped

tiles derived from a single high-resolution image and 876'960 tiles extracted from all images (348 images in total).

The image dataset demonstrated exceptional consistency in the feature space, featuring only minute variance. This consistency instrumentally overcame the inherent challenges of lens distortion, where features further away from the camera are captured with a bulge, elongation, or stretch (barrel distortion) without camera calibration.

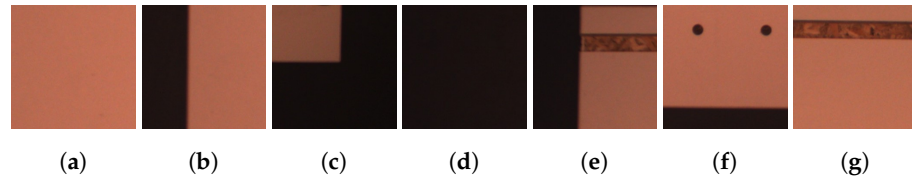


Figure 6. Sample windows of furniture finishing: (a) surface, (b) edge, (c) corner, (d) background, (e) edge and groove, (f) hole and edge, (g) groove.

4.3. Unsupervised Class Selection

The prototype image dataset features noticeable class imbalances as surface and background tiles dominate the remaining classes. To mitigate this issue, we adopt an approach of two steps: Firstly, we extract feature vectors for each cropped tile using the pre-trained VGG16 model with ImageNet1k_v1 weights. This step significantly contributes to resolving the class imbalance by enhancing the representation of each class. Next, we utilize the unsupervised k-means algorithm to cluster the obtained feature vectors into seven distinct clusters. In doing so, we identify inherent patterns and group samples regarding their similarity. Moreover, we deliberately drop the clusters with the highest frequency, corresponding to the surface class, and the second-highest frequency, associated with the background class. By discarding these clusters, we ensure that only the most relevant and informative ones are retained, thus facilitating a more balanced and meaningful representation of the dataset.

4.4. Image Dataset Augmentation

Data augmentation in computer vision is the process employed to increase the diversity of a training dataset by applying diverse transformations such as rotation, flipping, scaling, and cropping to the original images. This technique helps to enhance the ability of the model to generalize by exposing it to different variations of the data, thereby improving performance and reducing overfitting. We achieve data augmentation by adopting the extraction of segments from the Describable Textures Dataset (DTD) [19]. The DTD is a collection of natural patterns and textures, serving as a foundation for developing better methods to recognize and understand texture attributes in images. Notably, the DTD contains image resolutions higher than the cropped-out tiles resulting from the sliding window technique. We employ Felzenswalb's segmentation algorithm to extract segments from the DTD with specific values for the algorithm parameters, as detailed in Table 1. Next, the resulting image segments are randomly selected and superimposed onto the tile, as demonstrated in Figure 7. This approach enables the generation of diverse artificial anomaly variations, significantly contributing to the overall performance of the anomaly detection model and increasing its robustness to detect anomalous patterns in various real-world scenarios. This augmentation strategy is instrumental in improving the model's generalization capabilities and practical applicability in real-world anomaly detection tasks. We additionally employ data augmentation techniques, as detailed in Table 2. During the training process, these augmentations are randomly applied to each image, enriching the diversity of the data and enhancing the model's ability to generalize effectively.

Table 1. Felzenswalb’s parameter values.

Parameter	Value
scale	2
sigma	5
min_size	100

Table 2. Augmentation values.

Augmentation	Value
horizontal flip	probability = 0.5
vertical flip	
brightness range	[0.98, 1.5]
contrast range	[1, 1.2]



Figure 7. DTD (left) and artificial anomaly overlaid on crop window of prototype dataset (right).

4.5. Network Architecture

Inspired by the work of Hida et al. [9], we propose the network architecture, as depicted in Figure 8, with the principal objective of encoding an input image of resolution 289×289 into a compact latent representation of size $512 \times 1 \times 1$ using the encoder. After the encoding, the decoder reconstructs the original input image. As indicated in Figure 8, the autoencoder features skip connections that bypass the middle part of the network. In general, skip connections enable the network to more easily retain finer details from the input, thus facilitating the learning of a more accurate reconstruction of the original data. By bypassing specific layers, skip connections can also alleviate the vanishing gradient problem, which is particularly beneficial when training deeper autoencoders.

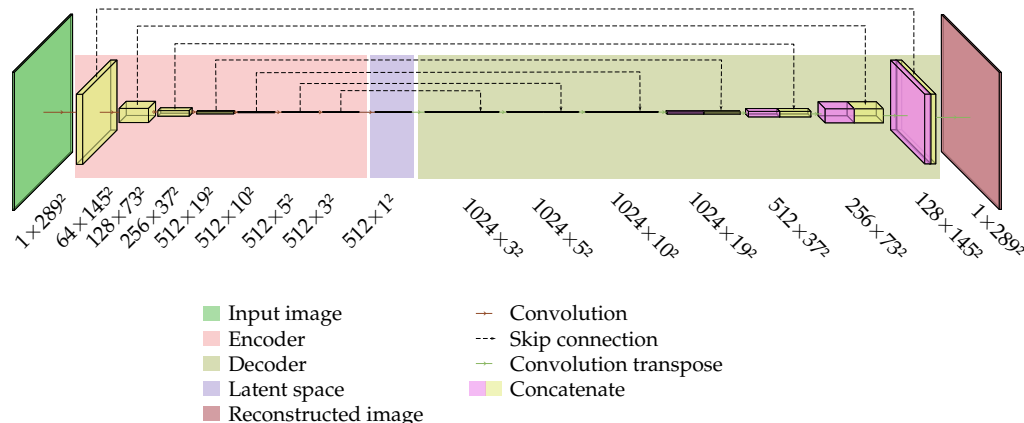


Figure 8. Autoencoder architecture.

4.6. Loss Function

At the outset, we implemented the mean squared error (MSE) loss function alone, which yielded unsatisfactory results. However, we observed a significant improvement after introducing the structural similarity index (SSIM) metric as an additional term in the

loss function. This addition notably enhanced the preservation of intricate image details, resulting in greater accuracy. Given that SSIM is for grayscale images, we adjusted the dataset accordingly, leading us to utilize grayscale images for training. Further increasing overall performance, we integrate an extra MSE loss component tailored for overlay regions. This addition proved crucial in emphasizing precise reconstruction within these critical areas, contributing to enhanced outcomes. The primary responsibility for emphasizing accurate reconstruction in these vital areas, where anomalies exist, lies with the MSE at overlay loss term. Its influence leads to improved results, particularly in overlay regions. We culminate these insights and develop the final version of the loss function (see Equation (1)), which combines MSE (Equation (2)), SSIM (Equation (3)), and MSE at overlay (Equation (4)) for training. The individual loss weights in Equation (1), λ_{MSE} , λ_{SSIM} , and $\lambda_{MSE_artificial_anomaly}$ use a value of one. In the last step, capitalizing on the strengths of these three distinct loss metrics, we train the model to achieve excellent accuracy and robustness in addressing artificial anomalies and irregularities such as unexpected variations, inconsistencies, abnormal patterns, noise, errors, or outliers in the data. This comprehensive approach ensured our model's proficiency in handling complex challenges within the dataset.

$$Loss = \lambda_{MSE} L_{MSE}(Y, \hat{Y}) + \lambda_{SSIM}(1 - L_{SSIM}(Y, \hat{Y})) + \lambda_{MSE_artificial_anomaly} L_{MSE_artificial_anomaly}(Y, \hat{Y}) \quad (1)$$

$$L_{MSE}(Y, \hat{Y}) = \frac{1}{p} \sum_{i=1}^p (Y_i - \hat{Y}_i) \quad (2)$$

$$L_{SSIM}(Y, \hat{Y}) = \frac{1}{q} \sum_{j=1}^q \frac{(2\mu_{Y_j} \mu_{\hat{Y}_j} + c_1)(2\sigma_{Y_j \hat{Y}_j} + c_2)}{(\mu_{Y_j}^2 + \mu_{\hat{Y}_j}^2 + c_1)(\sigma_{Y_j}^2 + \sigma_{\hat{Y}_j}^2 + c_2)}, \quad (3)$$

where $c_1 = 0.01$, and $c_2 = 0.03$

$$L_{MSE_artificial_anomaly}(Y, \hat{Y}) = \frac{1}{r} \sum_{k=1}^r (Y_k - \hat{Y}_k), \quad (4)$$

where $k \in \{\text{artificial anomaly pixels}\}$

4.7. Training

Our neural network implementation utilizes the PyTorch framework [34]. We ensure reproducibility by setting a random seed of 42 and using modules such as NumPy and PyTorch. For optimization, we employ the Adam optimizer with a learning rate of 2×10^{-4} , betas set at (0.9, 0.999), an eps value of 1×10^{-8} , and with the amsgrad option disabled. We incorporate a learning rate scheduler and an early stopping strategy to optimize training performance. The "reduce on plateau" learning rate scheduler has a configuration with mode set to minimum, a factor of 0.7, patience of three epochs, a threshold of 1×10^{-4} , an eps value of 1×10^{-8} , verbose mode disabled, and cooldown and minimum learning rate set to zero. We assign a patience of 40 epochs and a minimum change in validation loss of 1×10^{-6} for early stopping. If the validation loss does not show improvement (change less than 1×10^{-6}) for 40 consecutive epochs, the training process halts. We save the model's weights when the validation loss during training is lower than the previous lowest value, thus capturing the best performance achieved during training. These optimization strategies enhance performance and reproducibility in network training, surpassing non-optimized models. This results in more consistent and reliable outcomes. Figure 9 shows the training performance through these applied optimizations, achieving completion in two days and four hours with 285 epochs.

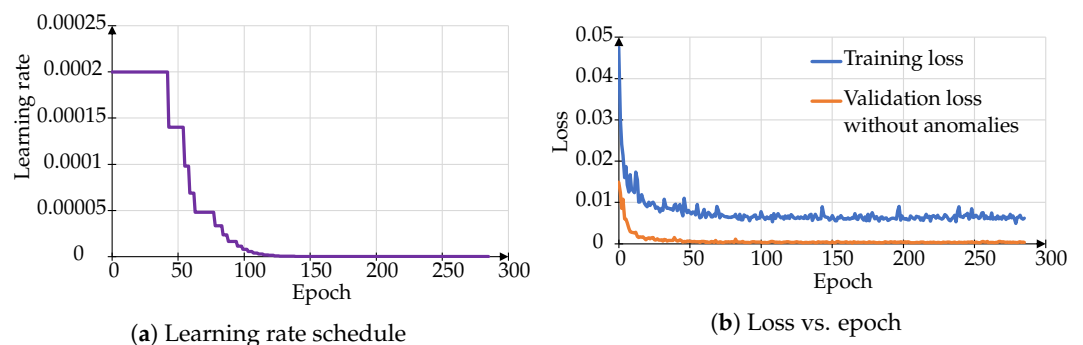


Figure 9. Training performance.

4.8. Hardware Setup

The demonstrator is a lightproof box with an inner dimension of $1200 \times 800 \times 2033 \text{ mm}^3$. It incorporates a clearance height of 402.00 mm at the bottom for an AGV or movable table to pass through. The camera mounting is 1088 mm above the clearance, ensuring optimal image capturing conditions. A diffuse artificial light source is positioned behind the camera, providing a constant and uniform light distribution throughout the imaging process. A Raspberry Pi 4 computer with a touch display is on the side of the prototype setup. The model trains on an AMD Ryzen 9 7950x 16-Core Processor with 64 GB RAM and Nvidia GeForce RTX 3070 Lite Hash Rate with 8GB VRAM.

5. Results & Discussion

In the context of classification analysis, the false positive rate (FPR) and true positive rate (TPR) are pivotal metrics employed to evaluate the efficacy of a classifier in alignment with actual ground-truth labels. In the context of an unsupervised model akin to ours, synthetically generated anomalies influence the assessment of the model's performance. It is noteworthy to consider that incorporating bona fide ground-truth data could enhance the model's interpretative capacity.

In our investigation, a suite of seven discrete melamine-faced board images assumes the role of the litmus test for gauging the model's capabilities, portrayed in Figure 10's ROC plot. The subsequent narrative endeavors to unfold the nuanced interpretations latent within this ROC plot. The zone where TPR attains unity while FPR dwells at zero demonstrates an impeccable proficiency in demarcating anomalies from their non-anomalous counterparts. However, a discernible pattern emerges in the investigation of the shown ROC plot. Initially, at exceedingly low thresholds, the TPR experiences a precipitous ascent until 0.4 while the FPR remains close to zero. This dynamic underscores the model's capacity to effectively localize anomalies, albeit with a trade-off that precision remains high while accuracy is fair. Proceeding along this ROC plot, a shift in the equilibrium is palpable. The velocity at which FPR escalates surpasses that of TPR as thresholds ascend. This divergence indicates the model's tendency to misclassify non-anomalous pixels as anomalies. Though accuracy registers an uptick, a proportionate decline in precision is also observed. Such a juxtaposition warrants a meticulous inquiry to study the underlying catalyst—whether inherent noise in the data or a manifestation of model overfitting to specific patterns—contributing to the abrupt spike in FPR. As thresholds scale higher, a unique dynamic manifests. The model traverses the image where all anomalous pixels are successfully localized; however, this feat coexists with a counterpoint: the erroneous labeling of non-anomalous pixels. This paradox underscores a control of heightened accuracy counterbalanced by a degrading precision.

In sum, this investigation of the ROC plot within our context shines a spotlight on the intricate interplay between the true positive rate and false positive rate, unraveling insights into the model's discrimination prowess, precision, and potential pitfalls attributed to varying thresholds. We select the average threshold value 0.04, corresponding to the TPR value 0.4, to obtain accurate predictions while maintaining low misclassification.

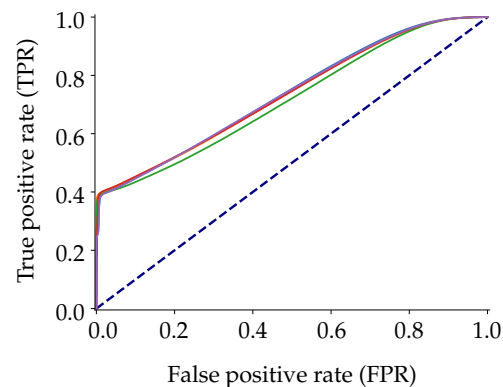


Figure 10. ROC plot of the model predictions on seven different melamine-faced board images, each plotted for a range of threshold values.

As previously stated, the uncropped images are of high resolution, and most of the anomalies are relatively tiny compared to the size of the other features. To better visualize the defects, the evaluation focuses on identifying specific crops of the melamine-faced board to showcase the performance, encompassing corners, edges, grooves, holes, and surfaces with their actual sharpness. Magnifying Figure 11 can provide a better visualization of the anomalies. While correct predictions generally characterize the non-anomalous area, accuracy fluctuates amidst anomalous regions. Given an unsupervised defect localizing model, achieving precise pixel-level performance is challenging due to the lack of definitive ground truth for the model to learn. Addressing this is done by generating heatmaps from the difference between the original and reconstructed images. These heatmaps are then overlaid with the corresponding actual anomalous crop areas to gauge the model behavior and the quality of its localization (refer to Figure 11). Each tile features a 289×289 resolution crop, organized in sets of three (a , b , and c) in each of the four columns. The first tile, a , represents the actual anomalous area crop, followed by the second tile, b , displaying the heatmap resulting from the difference, and the third tile, c , presenting an overlay of the heatmap on the actual anomalous crop area, with opacities set at 75% and 50%, respectively.

The model demonstrates decent capabilities regarding localization of tiny artifacts such as smudges, dirt, and deformities on the corners of the melamine-faced boards in Figure 11A1,A3,A4. The model accurately identifies these imperfections. However, the model's performance falls short when predicting the presence of large imperfections on the melamine-faced boards, as evidenced in Figure 11A2. It appears that the model struggles to detect and localize large-scale defects.

In the context of anomaly localization close to the edges of the boards, as shown in Figure 11B,C, the model's performance in predicting significant defects of various shapes and sizes is decent. However, challenges arise when dealing with relatively small-scale or (and) blurred anomalies, such as the one located on the left edge in Figure 11B2. In such cases, the model's confidence in its prediction decreases, leading to less accurate results. Figure 11C1 showcases the model's capability to accurately localize artifacts even when positioned slightly away from the edge, demonstrating its potential for robust anomaly detection in various scenarios.

The model exhibits mixed confidence when dealing with artifacts around the groove, as evidenced in Figure 11D1–D4. This behavior can be attributed to the artifact closely resembling the surrounding environment, making it challenging for the model to distinguish it as an anomaly. However, the model's performance receives a significant boost when dealing with instances of discontinuity. Discontinuities in the data are more straightforward for the model to identify and classify as anomalies, resulting in higher confidence predictions. Moreover, the defects found at the end of the groove are also localized by the model with moderate confidence, as seen in Figure 11E1, suggesting that the model

can detect these defects to some extent, but there may still be room for improvement in accuracy and precision.

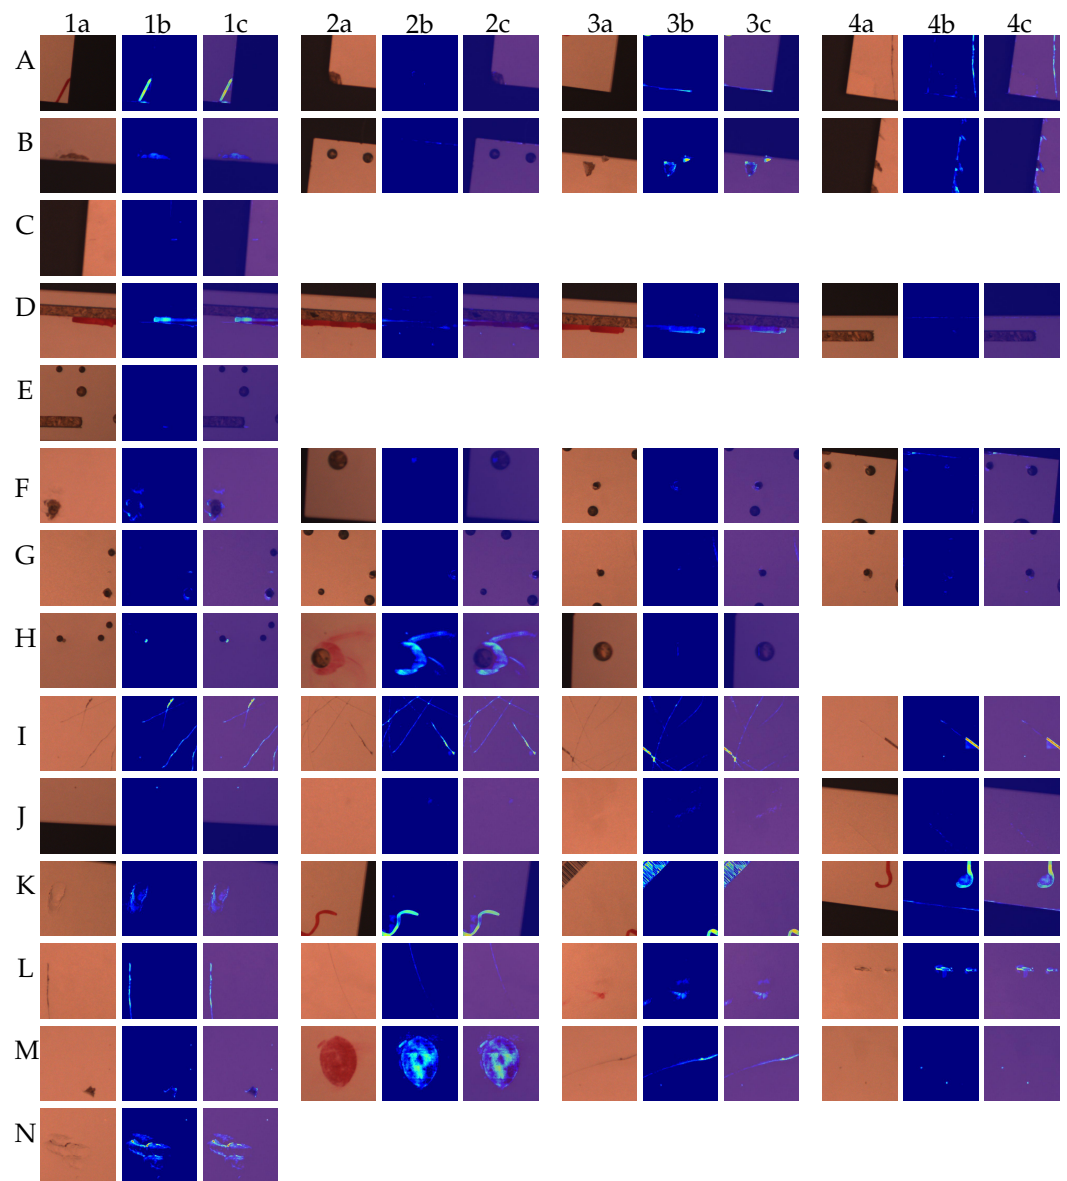


Figure 11. The defect localization on the melamine-faced board involves four columns with the following three grouping sets: (a) the actual anomaly, (b) heatmap of prediction, and (c) the overlay of the heatmap on the original defect, for evaluation. The rows are categorized as follows: corners -> (A), edges -> (B,C), grooves -> (D,E), holes -> (F-H), and plain surface -> (I-N).

The model's performance in localizing anomalies around holes is observable in Figure 11F–H. It successfully identifies anomalies within hole textures. Moreover, the model exhibits its proficiency in accurately pinpointing localizing different types of defects, such as defects along edges and holes (Figure 11F4) as well as defects on surfaces and holes (evident in Figure 11G1,G3,H1). The versatile capability to address diverse anomaly types underscores the model's effectiveness and potential.

Figure 11I–N evidences the model's localization performance for plain surface defects. It successfully detects prominent anomalies and localizes even the most minute plain surface defects, as evident in Figure 11J1–J4. Moreover, the model demonstrates its ability to recognize defects resembling closely to holes in shape and texture, as observed in Figure 11M1.

Overall, the results in Figure 11 highlight the model's robustness, limitations, and accuracy in tackling complex defect localization tasks by identifying and localizing a wide range of surface defects occurring on melamine-faced boards.

6. Conclusions

Our paper presents a hybrid approach for detecting surface defects on melamine-faced boards. Unlike traditional methods that use images with a specific region of interest, our methodology utilizes a dataset with images of high resolution captured with a fixed field-of-view camera, enabling the inspection of boards of varying sizes. This model combines techniques, including slicing high-resolution images, addressing imbalanced datasets, performing feature extraction and k-means clustering due to feature frequency variations, and utilizing an autoencoder model for anomaly prediction.

The unsupervised defect localizing model evaluation reveals strong performance in identifying anomalies within the melamine-faced board. The model recognizes and localizes minor artifacts such as smudges, dirt, and deformities, even in corners and edges. It also accurately predicts significant defects around edges and holes, showcasing its robust anomaly detection capabilities. However, challenges arise when dealing with relatively small and blurred anomalies, particularly around the edges. Moreover, the model struggles to identify and localize larger missing sections on melamine-faced boards. The model's behavior around grooves is a mixed bag; it identifies instances of discontinuity but faces difficulties distinguishing artifacts that blend with the surrounding environment. When it comes to plain surface defects, the model performs effectively. It successfully detects large and small surface defects, highlighting its versatility in identifying diverse anomaly types, including those resembling holes in both shape and texture.

In conclusion, the model exhibits promising potential for improving quality control and inspection processes on surfaces of melamine-faced board. This potential is especially evident in its ability to identify defects on plain surfaces, corners, edges, and holes. However, it is necessary to acknowledge the need for continuous refinement to enhance its performance in identifying specific defect types and managing anomalies with less distinct attributes, as discussed in the following subsection.

Future Scope

This paper introduces an approach to localize anomalies on melamine-faced boards, addressing class imbalance challenges while maintaining quality standards. However, let us look at avenues for enhancing our model and its performance in anomaly localization on melamine-faced boards. One potential improvement involves using a higher-resolution camera with a larger aperture to aid in capturing anomalies more effectively, particularly significant ones. Moreover, such a camera could facilitate the imaging of larger boards, contributing to better feature learning with higher memory use. Another approach is the implementation of a camera array, requiring proper calibration and considerations for overlap regions, rectifying barrel distortion to minimize variations in class features, effectively reducing the number of training images. Addressing class imbalance warrants further investigation, especially in identifying damaged corners. One possible strategy is to train separate models for each feature following an initial classifier, involving weighing trade-offs related to the overall model size. Another is to assign weights to each class, which requires an excellent classifier at the initial stage. Optimizing the hyperparameters for Felzenszwalb's segmentation could lead to an artificial dataset representing natural anomalies. This optimization can enhance the model's robustness and improve performance metrics. Exploring alternative architectures, such as UNet++ introduced by Zhou et al. [35], may offer improved power and generalization capabilities, even though it could lead to additional training parameters. Incorporating techniques like the weighted frequency domain loss, as explored by Nakanishi et al. [15], and NST loss, as investigated by Hida et al. [9], could enhance the model's ability to generalize on the edges and textures of board classes.

By pursuing these avenues of improvement, our model could achieve higher accuracy, greater robustness, and improved generalization for localizing anomalies on melamine-faced boards, contributing to elevating quality control standards within the manufacturing industry.

Author Contributions: Conceptualization, D.M.; methodology, D.M.; software, D.M.; formal analysis, D.M.; investigation, D.M.; data curation, D.M.; writing—original draft preparation, D.M.; writing—review and editing, N.K.; supervision, N.K.; project administration, N.K. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: For interested parties, the image dataset utilized in this study is available upon request. To obtain the dataset, please contact the corresponding author, Noah Klarmann, via email at noah.klarmann@th-rosenheim.de. Requests will be processed and the dataset will be transferred accordingly.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

Adam	Adaptive moment optimization
AUC	Area under curve
AUROC	Area under the receiver operator characteristic
ADM	Artificial defect module
AI	Artificial intelligence
AGV	Automated guided vehicle
CV	Computer vision
DL	Deep learning
DTD	Describable textures dataset
DAE	Denoising autoencoder
FPR	False positive rate
GAN	Generative adversarial network
MSE	Mean square error
NST	Neural style transfer
OpenCV	Open source computer vision library
ROC	Receiver operator characteristics
SSAE	Semisupervised autoencoder
SSIM	Structural similarity
TPR	True positive rate
VQ-VAE	Vector-quantized variational autoencoder

References

- Li, R.; Zhao, S.; Yang, B. Research on the application status of machine vision technology in furniture manufacturing process. *Appl. Sci.* **2023**, *13*, 2434. [CrossRef]
- O'Mahony, N.; Campbell, S.; Carvalho, A.; Harapanahalli, S.; Hernandez, G.V.; Krpalkova, L.; Riordan, D.; Walsh, J. Deep Learning vs. Traditional Computer Vision. In *Proceedings of the Advances in Computer Vision*; Arai, K., Kapoor, S., Eds.; Springer: Cham, Switzerland, 2020; pp. 128–144. [CrossRef]
- Weiss, K.; Khoshgoftaar, T.M.; Wang, D. A survey of transfer learning. *J. Big Data* **2016**, *3*, 1–40. [CrossRef]
- Hartigan, J.A.; Wong, M.A. A K-Means Clustering Algorithm. *J. R. Stat. Soc. Ser. C (Appl. Stat.)* **1979**, *28*, 100–108. [CrossRef]
- Felzenszwalb, P.F.; Huttenlocher, D.P. Efficient graph-based image segmentation. *Int. J. Comput. Vis.* **2004**, *59*, 167–181. [CrossRef]
- Wang, Z.; Bovik, A.; Sheikh, H.; Simoncelli, E. Image quality assessment: From error visibility to structural similarity. *IEEE Trans. Image Process.* **2004**, *13*, 600–612. [CrossRef] [PubMed]
- Vasilev, I.; Slater, D.; Spacagna, G.; Roelants, P.; Zocca, V. *Python Deep Learning: Exploring Deep Learning Techniques and Neural Network Architectures with Pytorch, Keras, and TensorFlow*; Packt Publishing Ltd.: Birmingham, UK, 2019. Available online: <https://www.packtpub.com/product/python-deep-learning-second-edition/9781789348460> (accessed on 1 December 2023).

8. Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative adversarial nets. In *Proceedings of the Advances in Neural Information Processing Systems*; Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N., Weinberger, K., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2014; Volume 27. Available online: https://proceedings.neurips.cc/paper_files/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf (accessed on 1 December 2023).
9. Hida, Y.; Makariou, S.; Kobayashi, S. Smart Image Inspection using Defect-Removing Autoencoder. *Procedia CIRP* **2021**, *104*, 559–564. [[CrossRef](#)]
10. Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Proceedings of the Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015*, Munich, Germany, 5–9 October 2015; Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F., Eds.; Springer: Cham, Switzerland, 2015; pp. 234–241. [[CrossRef](#)]
11. Simonyan, K.; Zisserman, A. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv* **2014**, arXiv:1409.1556. [[CrossRef](#)]
12. Goodfellow, I.; Bengio, Y.; Courville, A. *Deep Learning*; MIT Press: Cambridge, MA, USA, 2016. Available online: <http://www.deeplearningbook.org> (accessed on 1 December 2023).
13. Kamal, K.; Qayyum, R.; Mathavan, S.; Zafar, T. Wood defects classification using laws texture energy measures and supervised learning approach. *Adv. Eng. Inform.* **2017**, *34*, 125–135. [[CrossRef](#)]
14. Laws, K.I. Texture energy measures. In *Proceedings of the Image Understanding Workshop*, Palo Alto, CA, USA, 24–25 April 1979; pp. 47–51. Available online: <https://apps.dtic.mil/sti/pdfs/ADA077568.pdf#page=57> (accessed on 1 December 2023).
15. Nakanishi, M.; Sato, K.; Terada, H. Anomaly Detection By Autoencoder Based On Weighted Frequency Domain Loss. *arXiv* **2021**, arXiv:2105.10214. [[CrossRef](#)]
16. Bergmann, P.; Fauser, M.; Sattlegger, D.; Steger, C. MVTEC AD—A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, USA, 15–20 June 2019; pp. 9592–9600. Available online: https://openaccess.thecvf.com/content_CVPR_2019/papers/Bergmann_MVTEC_AD_-_A_Comprehensive_Real-World_Dataset_for_Unsupervised_Anomaly_CVPR_2019_paper.pdf (accessed on 1 December 2023).
17. Zavrtnik, V.; Kristan, M.; Skočaj, D. DRAEM—A Discriminatively Trained Reconstruction Embedding for Surface Anomaly Detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, Montreal, BC, Canada, 11–17 October 2021; pp. 8330–8339. Available online: https://openaccess.thecvf.com/content_ICCV2021/papers/Zavrtnik_DRAEM_-_A_Discriminatively_Trained_Reconstruction_Embedding_for_Surface_Anomaly_ICCV_2021_paper.pdf (accessed on 1 December 2023).
18. Perlin, K. An Image Synthesizer. *SIGGRAPH Comput. Graph.* **1985**, *19*, 287–296. [[CrossRef](#)]
19. Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; Vedaldi, A. Describing Textures in the Wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Columbus, OH, USA, 23–28 June 2014; pp. 3606–3613. Available online: https://openaccess.thecvf.com/content_cvpr_2014/papers/Cimpoi_Describing_Textures_in_2014_CVPR_paper.pdf (accessed on 1 December 2023).
20. Zhang, H.; Jiang, L.; Li, C. CS-ResNet: Cost-sensitive residual convolutional neural network for PCB cosmetic defect detection. *Expert Syst. Appl.* **2021**, *185*, 115673. [[CrossRef](#)]
21. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep Residual Learning for Image Recognition. In *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, USA, 27–30 June 2016; pp. 770–778. [[CrossRef](#)]
22. Deng, Y.S.; Luo, A.C.; Dai, M.J. Building an automatic defect verification system using deep neural network for pcb defect classification. In *Proceedings of the 2018 4th International Conference on Frontiers of Signal Processing (ICFSP)*, Poitiers, France, 24–27 September 2018; pp. 145–149. [[CrossRef](#)]
23. Jing, Y.; Yang, Y.; Feng, Z.; Ye, J.; Yu, Y.; Song, M. Neural Style Transfer: A Review. *IEEE Trans. Vis. Comput. Graph.* **2020**, *26*, 3365–3385. [[CrossRef](#)] [[PubMed](#)]
24. Deng, J.; Dong, W.; Socher, R.; Li, L.J.; Li, K.; Fei-Fei, L. ImageNet: A large-scale hierarchical image database. In *Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition*, Miami, FL, USA, 20–25 June 2009; pp. 248–255. [[CrossRef](#)]
25. Schlüter, H.M.; Tan, J.; Hou, B.; Kainz, B. Natural Synthetic Anomalies for Self-supervised Anomaly Detection and Localization. In *Proceedings of the Computer Vision—ECCV 2022*, Tel Aviv, Israel, 23–27 October 2022; Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T., Eds.; Springer: Cham, Switzerland, 2022; pp. 474–489. [[CrossRef](#)]
26. Pérez, P.; Gangnet, M.; Blake, A. Poisson Image Editing. *ACM Trans. Graph.* **2003**, *22*, 313–318. [[CrossRef](#)]
27. Mujkic, E.; Philipsen, M.P.; Moeslund, T.B.; Christiansen, M.P.; Ravn, O. Anomaly Detection for Agricultural Vehicles Using Autoencoders. *Sensors* **2022**, *22*, 3608. [[CrossRef](#)] [[PubMed](#)]
28. Vincent, P.; Larochelle, H.; Bengio, Y.; Manzagol, P.A. Extracting and Composing Robust Features with Denoising Autoencoders. In *Proceedings of the 25th International Conference on Machine Learning*, New York, NY, USA, 5–9 July 2008; ICML '08, pp. 1096–1103. [[CrossRef](#)]
29. Van Den Oord, A.; Vinyals, O.; Kavukcuoglu, K. Neural Discrete Representation Learning. In *Proceedings of the Advances in Neural Information Processing Systems*; Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2017; Volume 30. Available online: https://proceedings.neurips.cc/paper_files/paper/2017/file/7a98af17e63a0ac09ce2e96d03992fbc-Paper.pdf (accessed on 1 December 2023).

30. Jocher, G.; Chaurasia, A.; Stoken, A.; Borovec, J.; Kwon, Y.; Michael, K.; Fang, J.; Yifu, Z.; Wong, C.; Montes, D.; et al. *ultralytics/yolov5: v7.0 - YOLOv5 SOTA Realtime Instance Segmentation*, Version v7.0; Zenodo: Geneva, Switzerland, 2022. [CrossRef]
31. Huang, C.; Xu, Q.; Wang, Y.; Wang, Y.; Zhang, Y. Self-Supervised Masking for Unsupervised Anomaly Detection and Localization. *IEEE Trans. Multimed.* **2023**, *25*, 4426–4438. [CrossRef]
32. Hoffer, E.; Ailon, N. Deep Metric Learning Using Triplet Network. In *Proceedings of the Similarity-Based Pattern Recognition*; Feragen, A., Pelillo, M., Loog, M., Eds.; Springer: Cham, Switzerland, 2015; pp. 84–92. [CrossRef]
33. Raspberry Pi Documentation. Raspberry Pi Documentation—Raspberry Pi hardware. Available online: <https://www.raspberrypi.com/documentation/computers/raspberry-pi.html> (accessed on 14 August 2023).
34. Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. *arXiv* **2019**, arXiv:1912.01703. [CrossRef].
35. Zhou, Z.; Siddiquee, M.M.R.; Tajbakhsh, N.; Liang, J. UNet++: Redesigning Skip Connections to Exploit Multiscale Features in Image Segmentation. *IEEE Trans. Med. Imaging* **2020**, *39*, 1856–1867. [CrossRef] [PubMed]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.