

Can I Trust My Anomaly Detection System? A Case Study Based on Explainable AI

Muhammad Rashid¹, Elvio Amparore^{1(⊠)}, Enrico Ferrari², and Damiano Verda²

¹ Computer Science Department, University of Torino, C.so Svizzera 185, 10149 Torino, Italy

{muhammad.rashid,elviogilberto.amparore}@unito.it ² Rulex Innovation Labs, Via Felice Romani 9, 16122 Genova, Italy {enrico.ferrari,damiano.verda}@rulex.ai

Abstract. Generative models based on variational autoencoders are a popular technique for detecting anomalies in images in a semi-supervised context. A common approach employs the anomaly score to detect the presence of anomalies, and it is known to reach high level of accuracy on benchmark datasets. However, since anomaly scores are computed from reconstruction disparities, they often obscure the detection of various spurious features, raising concerns regarding their actual efficacy.

This case study explores the robustness of an anomaly detection system based on variational autoencoder generative models through the use of eXplainable AI methods. The goal is to get a different perspective on the real performances of anomaly detectors that use reconstruction differences. In our case study we discovered that, in many cases, samples are detected as anomalous for the wrong or misleading factors.

Keywords: anomaly detection \cdot variational autoencoder \cdot eXplainable AI

1 Introduction

The popularity of machine learning methods in difficult tasks, like the detection of anomalies in industrial quality-control processes, has witnessed a significant surge over the past decade. Variational AutoEncoders paired with a Generative Adversarial Networks, commonly referred as VAE-GAN [1] models, are particularly prominent in this regard, due to their high potential in representation learning. Anomaly Detection (AD) on image data with Deep Generative Models (DGM) [2] operates on the premise that a model can be trained to learn a representation of the normal features of a sample, while deliberately excluding the capacity to represent and generate any anomalies. An *anomaly score* can then be defined on the difference between the original image and its reconstruction, thus quantifying the representational gap for the sample abnormalities. While successful results have been reported using this strategy [3], significant challenges remain. An important issue with this approach is that reconstruction differences may actually be either real anomalies, or could be caused by the inability of the generative model to faithfully reproduce the input image. Additionally, VAE-GAN models often produce images that lack sharpness and details, amplifying differences, particularly at the borders. Even VAE model with vector quantization exhibit limited improvement in the reconstruction task [4].

This paper presents a small case study of the performances of a VAE-GAN AD system applied on the popular MVTec dataset [5]. We review the general framework for anomaly detection using autoencoders by Ravi & al. [3], which was outlined qualitatively but lacked quantitative evaluation. Our study reproduces that framework, augmenting it with additional insights for the explanation part. The work of [3] leveraged eXplainable AI (XAI) techniques like LIME and SHAP specifically adapted for anomaly detection (AD). However, their focus was on using XAI for visual explanation to improve anomaly localization compared to basic residual maps, rather than ensuring that the explained anomalies themselves were valid. Additionally, they did not quantify their findings.

In this paper we:

- Review an explainable AD system architecture that combines VAE-GAN models with the LIME and SHAP explanation methods;
- Quantify the AD system efficacy using anomaly scores;
- Use XAI methods to determine if anomalies are indeed detected for the right reason by comparing them with a ground truth, improving the framework of [3]. Our results reveal instances where samples were classified as anomalous but for incorrect reasons. To identify such samples, we employ a methodology based on the optimal Jaccard score.

2 Literature Review

AD is a well developed field, that has received a lot of attention due to its critical role in numerous practical applications. Creating effective detection systems is challenging due to several factors, like the difficulty of precisely define what an abnormality is within specific contexts, or the the lack of anomalous samples.

For these reasons, explaining the behaviour of an AD system remains a complex task. While general purpose interpretability techniques such as GradCAM [6], LIME [7] or SHAP [8,9] are available, some scholars regard them as imprecise and unreliable [10]. Moreover, their application in the realm of anomaly detection is inherently challenging, due to the lack of a probabilistic black-box function to explain. Nonetheless, these methodologies can be adapted to offer invaluable insights into understanding the rationale behind the behavior of AD systems. In this study we focus on LIME and SHAP systems, due to their (partially) comparable characteristics and their capability in localizing activation areas in anomaly maps. A broader recent review on AD systems is [11].

An XAI method for VAE-based systems is VAE-LIME [12], which is based on generating random samples in the latent space of the VAE model. However, it is unclear how this approach can be used in an anomaly detection setting, as it is not obvious how perturbed latent dimensions maps back to the original image segments. A methodology for explaining anomalies detected by VAE models using SHAP has been developed in [13], and our study considers this approach.

A general anomaly detection framework using autoencoders for images is discussed qualitatively in [3], with our study focusing on reproducing and refining it, particularly in the explanation aspect. In that framework, AD relies on anomaly scores, requiring threshold calibration. The challenges of perturbationbased methods, such as difficulty in setting appropriate thresholds, are addressed in [14]. Alternatives like residual explainers for AD have been explored in [15].

In XAI for anomaly detection, *anomaly localization* [16] is crucial. It improves interpretability by transitioning from pixel-based scores to region localization, especially challenging given the small size of anomalies in real-world datasets.

3 Preliminaries

We describe the relevant preliminaries following the workflow depicted in Fig. 1. The approach shares many similarity with [3]. Consider the problem for the



Fig. 1. AD system using a VAE-GAN model with LIME explanations.

domain of $h \times w$ images $\mathcal{I} \in [0 - 255]^{h \times w \times 3}$, where a sample $\xi \in \mathcal{I}$ may be normal or anomalous. We consider images from the high-quality open industrial dataset MVTec [5], namely the categories *hazelnut* and *screw*. From a training set (Fig. 1/A) containing only normal data (i.e. without anomalies) a VAE-GAN model is trained (Fig. 1/B).

3.1 VAE-GAN Models

A Variational Autoencoder Generative Adversarial Network (VAE-GAN) combines [1,17] the strengths of both variational autoencoders (VAEs) and generative adversarial networks (GANs) [18]. A VAE-GAN consists of an encoder e, a decoder d and a discriminator s. The encoder function $e : \mathcal{I} \to \mathcal{Z}$ maps input data, such as images \mathcal{I} , to a lower-dimensional latent space $\mathcal{Z} \in \mathbb{R}^{z}$, where each point in \mathcal{Z} represents a potential data sample. The decoder function $d : \mathcal{Z} \to \mathcal{I}$ estimates a potential input from a latent space representation, i.e. d approximates e^{-1} . Therefore, encoding and decoding an input image ξ results in its reconstruction (Fig. 1/C) through the latent representation z, given by

$$z = e(\xi), \qquad \xi' = d(z)$$

The distribution of the latent space is learnt using a probabilistic approach, and adopts both a regularization of the latent distribution (usually Gaussian) and a GAN approach for adversarial (joint) training of both d and e using the discriminator function s (a classifier trained to distinguish between real and generated data). When encoded, each data point is described by a Gaussian distribution, with mean μ and (log)-variance σ , from which new samples z can be drawn.

3.2 Semi-supervised Anomaly Detection Using Variational Models

While the task of identifying anomalies, particularly in image-based data, holds significant interest across various application domains [5,19], creating effective anomaly detectors remains a challenge. Imbalanced datasets are common, with anomalous data being significantly underrepresented (due to the infrequency of anomalous events). Furthermore, the definition of what constitutes an anomaly is often ambiguous, making supervised learning approaches impractical. Therefore, a relevant approach is based on the use of *semi-supervised* learning, where models are trained to detect anomalies from "normal" data only. Several approaches are possible to perform anomaly detection in a semi-supervised way [20], and in this study we consider a VAE-GAN-based approach [21].

A VAE-GAN model (e, d, s) for AD is trained exclusively on "normal" data, ensuring that only normal data has a proper representation in the latent space. Consider an input image ξ , and let $\xi' = d(e(\xi))$ be its encoding-decoding through the VAE-GAN model. If the sample is normal and lies in-distribution with the model, it should be reconstructed accurately, with minimal reconstruction errors. Conversely, if ξ has anomalous regions, its reconstruction ξ' is likely to resemble that of a normal sample, thereby allowing anomalies to be detected by difference.

Following [3], an anomaly reconstruction error map $m \in \mathbb{R}^{h \times w}$ assigns to each pixel of an image ξ its likelihood of being anomalous (Fig. 1/D), using

$$m = |gs(\xi) - gs(\xi')|, \qquad \alpha = \max(m)$$

where $gs: \mathbb{R}^{h \times w \times 3} \to \mathbb{R}^{h \times w}$ performs a per-pixel maximization of the three color channel values, α is the maximum anomaly value found, denoted as the *anomaly score* (Fig. 1/E). Alternative definitions of anomaly scores have also been explored [13]. While the anomaly map m can be used to visually inspect the reconstruction error, it suffers from limitations:

- it does not distinctly identify the anomaly per se, being at the pixel level;
- it provides only superficial insights into why a sample may be deemed anomalous.

An anomaly detection threshold τ is used to decide if a sample is classified as anomalous, i.e. when $\alpha \geq \tau$. An optimal threshold τ^* for the whole dataset can be determined using a calibration set (in this study, the test set) as

$$\tau^* = \operatorname*{argmax}_{\tau} \sqrt{\mathrm{TPR}(\tau) \times (1 - \mathrm{FPR}(\tau))}$$

where TPR and FPR denote the true positive rate and false positive rate, respectively, for the anomaly detection on the calibration set. Note that this threshold calibration is a critical and fragile part of this class of AD systems, as it is hard to generalize across different domains or datasets.

3.3 Explaining Anomaly Maps Using Model-Agnostic XAI Methods

While anomaly maps reveal the reconstruction errors, they only provide a superficial indication of potential anomaly areas within the input image, lacking precise localization of anomalies. To address this limitation, XAI methods have been adopted to help in localizing these areas for anomalous samples. We focus on model-agnostic methods based on perturbations of input data. Although many XAI methods rely on classifier predictions, reconstruction-based AD does not inherently provide such probability scores, and a special setup is needed [3, 4.1]. We consider two XAI methods, LIME and SHAP, adapted as described.

LIME. The Local Interpretable Model-Agnostic Explanations [7], is a method for explainable AI that works by creating a simpler, interpretable model that approximates the behavior of a more complex model in a synthetic neighborhood of a particular instance being explained. Let $f : \mathcal{I} \to \mathbb{R}$ be a prediction regression function that assigns probability scores to input images $\xi \in \mathcal{I}$. LIME produces an *high-level explanation* consisting of feature attributions (i.e. real-valued scores) assigned not at the pixel-level, but at the level of $k \ll (w \cdot h)$ superpixels. These superpixels represent pre-determined regions of the input image ξ characterized by a combination of color and spatial continuity. A common algorithm used to identify superpixels is *Quickshift* [22].

The k superpixels are used for masking, which is the step that generates the synthetic neighborhood $\mathcal{N}(\xi)$ made of perturbed images (Fig. 1/F). A mask $x \in \{0,1\}^k$ is a binary vector representing whether each of the k superpixels should be kept (value 1) or replaced (value 0). In standard LIME, masking vectors are sampled from an unbiased Bernoulli distribution B having probability 0.5, but more advanced sampling strategies have been proposed [23].

Let ξ_x be the perturbation of image ξ according to the masking vector x. The synthetic neighborhood $\mathcal{N}(\xi) = \{\xi_x \mid x \in X\}$ is then generated from a set X of n masking vectors, resulting in the corresponding dependent variables $Y = \{f(\xi_x) \mid \xi_x \in \mathcal{N}(\xi)\}.$

As previously mentioned, LIME is designed to explain a prediction function f, and it is not directly applicable to AD, since there is no function f producing probability scores. Nonetheless, it can be used to explain the reconstruction error as follows. A perturbed image ξ_x for mask x is defined, for every pixel p, as

$$\xi_x[p] = \begin{cases} \xi'[p] & \text{if pixel } p \text{ belongs to a masked superpixel in } x \\ \xi[p] & \text{otherwise} \end{cases}$$

where the reconstruction error of ξ_x is measured as the mean squared error w.r.t. the original input ξ , as

$$f(\xi_x) = \text{MSE}(\xi - \xi_x)$$

An explanation in LIME (Fig. 1/G) is obtained by fitting a simple linear model: $Y = X \cdot b + \epsilon$, where the vector *b* represents the weighted least squares estimator of the regression coefficients of *Y* on *X*, weighted by an appropriate distance function. A linear function g(x) with coefficients *b* acts as a local approximation of the square loss function *f*, and the real coefficients b[i] for each superpixel $1 \leq i \leq k$ are interpreted as *feature attribution* scores. An image-level *feature attribution explanation* β_L assigns feature attribution scores to individual pixel, such that each pixel of the *k* superpixels receive the corresponding coefficient in *b*.

SHAP. The SHapley Additive exPlanation method [8,24] provides a gametheoretical approach to assign feature importance scores to an input classified by a black-box model. Similarly to LIME, it is based on the concept of generating perturbations of the original input (with features masked using one or more "background" values). In the *KernelSHAP* method, perturbations are drawn from the Shapley distribution function. However, unlike LIME, explanation scores are computed from the marginal contribution that each input feature brings to the explained function f. The *SHAP partition explainer* [8] is a specialized image method that employs a recursive cut approach to localize relevant features within an input image. An explanation β_S generated by the SHAP partition explainer assigns feature attribution scores directly to pixels (Fig. 1/H). The granularity of these scores depend on a budget of n perturbed images that the XAI method can produce to explain an input sample ξ .

The application of SHAP to explain the anomalies revealed by an autoencoder has been developed in [13] and, similarly to LIME, is based on a reconstruction error function $f(\xi)$ but without relying on any predetermined superpixels.

3.4 Comparing Explained Anomalies Against a Ground Truth

A pixel-level feature attribution explanation β generated by an XAI method is a real matrix of feature attribution scores assigned to the pixels of the image. To assess the method's capability of localizing the anomalous regions in an input image, we adopt the following methodology. A Boolean ground truth $\gamma \in \{0,1\}^{h \times w}$ is a matrix that assigns, to each pixel of the input image ξ , a value whether the pixel belongs to the anomaly being localized or not (Fig. 1/I).

We assume that γ is available for the anomalous samples of the test set. Since the explanation β is a real-valued matrix, it is not directly comparable with γ . An effective way to perform such comparison is to define an *explanation* threshold θ , and define a boolean explanation γ' , derived from β , that marks as anomalous those pixels of ξ for which the feature attribution score in β is greater than θ . A comparison between γ and γ' can then be performed using standard metrics like the Jaccard coefficient (a.k.a. Intersection over Union - IoU)

$$J(\gamma, \gamma') = \frac{\gamma \land \gamma'}{\gamma \lor \gamma'}$$

However, determining an optimal threshold θ is not straightforward. Hence, we select, for each explained sample, a corresponding optimal threshold $\theta^*(\xi)$ for which $J(\gamma, \gamma')$ is maximal (Fig. 1/J). The mismatch between γ and γ' can then be inspected and visualized¹ (Fig. 1/K). Note that this coefficient can only be computed when γ is available and it is not empty (otherwise it would be meaningless). Thus it can be used only to explain anomalies for "abnormal" samples, but it cannot be used on "good" samples.

4 Experimental Evaluation

We present the results on a set of experiments made on the MVtec dataset [5] and considering two categories, *hazelnut* and *screw*, each comprising images of these objects with and without defects. The tests use a VAE-GAN model implemented in Keras [25], where the encoder model e uses 4 nested convolutional layers (3×3 kernel, stride 2), with each layer using ReLU activation and followed by a batch normalization, and using a final Dense decision layer. The discriminator s is similar to e, but using three convolutional layers with larger kernels (8×8, 5×5 and 4×4, respectively) and followed also by max pooling. The decoder d

¹ We adopt a threshold-maximization approach instead of a threshold-independent metric like AU-IoU, because the former has a more intuitive visualization.



Fig. 2. Maximum IoU vs the anomaly scores in the two test datasets.

mirrors the structure of e, but in reverse order and using transposed convolutions. Input images are scaled to 128×128 . Training is performed on $30\,000$ epochs on batches of 64 images, incorporating mild augmentation techniques (rotation, width/height shift, brightness adjustment, zoom) to mitigate overfitting and make the model more robust to variations in background light and shadows.

Due to the dependency of LIME on the quality of the segmentation in superpixels, we consider three evaluation setups:

- **S1**: LIME explanations with segmentation performed on the input image, without prior knowledge of the anomalies (fair setup). Potential misbehaviors may arise from either LIME's failure to localize anomalies or inaccuracies in the segmentation method in identifying anomaly boundaries. All explanations are computed using k = 100 segments, $n = 5\,000$ samples.
- **S2**: LIME explanations with segmentation performed knowing both the image and the ground truth. In this setup, we remove the segmentation method as a potential cause of LIME misbehaviors (anomalies fall into distinct segments). However, this setup is unrealistic since it exposes the ground truth. As before, we use k = 100 segments and $n = 5\,000$ samples.
- **S3**: SHAP explanations using partition explainer, with $n = 5\,000$ samples.



Fig. 3. Explanations for a few anomalous samples of the hazelnut dataset.

Explaining using $n=5\,000$ samples takes about 20 seconds on a M1 laptop. The plots in Fig. 2 illustrate the performance of the AD system (X axis) and its explainability in terms of maximal $J(\gamma, \gamma')$ scores (Y axis) on the test sets of the two considered datasets (left and right columns) in the three setups (rows). We denote LIME and SHAP explanations with J_L and J_S , respectively. Anomaly scores remain consistent within each column, with only the maximal J_L (resp. J_S) scores varying. The hazelnut dataset comprises 40 good (37 correctly classified, 3 misclassified) and 70 anomalous (62 correctly classified, 8 misclassified) samples, reaching 90% accuracy using the optimal threshold. The screw dataset includes 41 good (31 correctly classified, 10 misclassified) and 119 anomalous (97 correctly classified, 22 misclassified) samples, achieving 80% accuracy using the optimal threshold.

While it is expected that the maximal IoU should not be perfect, the scores obtained from the XAI methods already reveal that some samples exhibit very poor localization of the anomalies. Given that both LIME and SHAP compute explanations based on residual reconstruction errors, it is plausible that some samples are classified as good or anomalous for incorrect reasons. To evaluate this, we conduct manual inspection of the samples.

Hazelnut Dataset. Figure 3 illustrates a few selected anomalous samples from the hazelnut dataset². Each row shows, from left to right, the sample ξ and its reconstruction ξ' , the anomaly reconstruction error map m, the explanations β_L and β_S generated from LIME and SHAP, resp., the visualization of the maximal

 $^{^{2}}$ All test sample explanations are provided separately (link at the end of the paper).

 J_L and J_S for both explanation methods (the *J* value is reported in the upper-left corner), and the boundary of the ground truth region γ . All LIME explanations β_L come from the S1 setup, unless explicitly labeled as S2. SHAP explanations β_S are computed using the S3 setup.

Sample (A) from Fig. 3 shows a case of a hazelnut with a small surface crack that is properly localized and detected (with some negligible mistakes).

Sample (B) looks similar, but it is misclassified as good, having the anomaly score α below the threshold τ^* . However, the XAI methods would still localize the anomalous region.

In (C1), β_L shows significant confusion, attributing large values to the border instead of the small hole at the center. The primary issue lies in the segmentation: employing a segmentation that accurately encloses the anomaly (as in C2 with the S2 setup) results in better localization(even if some confusion remains). This underscores how LIME can be greatly influenced by inadequate segmentation.

Sample (D) shows an example where both LIME and SHAP fail to identify the anomaly accurately: since the reconstruction ξ' is not entirely faithful, both XAI methods mislocate the anomalous region to the top of the image, overlooking the actual one (a cut on the hazelnut shell).



Fig. 4. Explanations for a few anomalous samples of the screw dataset.

Screw Dataset. Detecting anomalies in this dataset presents greater difficulty as they typically occupy small portions of the image. While many samples are correctly classified and explained, accurately localizing the anomalous area proves challenging for others. In the four samples in Fig. 4, all correctly classified as anomalous, the feature attribution scores are maximal in areas that do not contain any anomaly (as evidenced by the large false-positive areas in the $\max(J)$ plots). Sample (C) is particularly critical, as both LIME and SHAP assign low scores to the region containing the anomaly (the right thread side). This suggests that the sample may have been classified as anomalous for the wrong reason, and this could only be detected through the use of XAI methods.

5 Conclusions

In this case study we replicated the framework of [3], enhancing it by quantifying both AD and XAI performances. Our aim was to highlight the relevance of XAI methods in finding the true drivers behind anomaly detection, particularly when utilizing reconstruction error maps generated from VAE-GAN models.

The results show that relying solely on the anomaly score is insufficient for comprehending the classification process. A sample may be detected as anomalous for the wrong reasons, yet this misbehaviour may not be detectable from the information provided by the anomaly map alone. We used two model-agnostic XAI methods to obtain explanations from the anomalous samples, to inspect if the anomalies were correctly localized. Region localization through a XAI method with Jaccard score maximization allows the user to inspect the AD system, identifying potential misbehaviors in the detection and providing a better understanding of the system.

Both tested XAI methods successfully localizes activation regions, with some discrepancies. Specifically, LIME exhibited a slightly inferior performance compared to SHAP, attributable to its reliance on a pre-determined segmentation that is not aware of the ML process and does not get any feedback from it. This fragility can be seen by the variations between the S1 and S2 test setups (like in Fig. 3/C1-C2).

Code availability. All code needed to replicate the experiments, including all the explanations for all test samples, are available at: https://github.com/rashidrao-pk/anomaly_detection_trust_case_study.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

- Larsen, A., Sønderby, S.K., Larochelle, H., Winther, O.: Autoencoding beyond pixels using a learned similarity metric. In: International Conference on ML, vol. 48, 1558–1566. PMLR (2016)
- Zhou, C., Paffenroth, R.C.: Anomaly detection with robust deep autoencoders. In Proceedings of ACM SIGKDD, pp. 665–674 (2017)
- Ravi, A., Yu, X., Santelices, I., Karray, F., Fidan, B.: General frameworks for anomaly detection explainability: comparative study. In: 2021 IEEE International Conference on Autonomous Systems (ICAS), pp. 1–5 (2021)
- Van Den Oord, A., Vinyals, O., et al.: Neural discrete representation learning. In: NeurIPS, vol. 30 (2017)
- Bergmann, P., Batzner, K., Fauser, M., Sattlegger, D., Steger, C.: The MVTec anomaly detection dataset: a comprehensive real-world dataset for unsupervised anomaly detection. Int. J. Comput. Vision 129, 1038–1059 (2021)
- Selvaraju, R.R., et al.: Grad-cam: visual explanations from deep networks via gradient-based localization. In: ICCV, pp. 618–626 (2017)

- Ribeiro, M.T., Singh, S., Guestrin, C.: Why should i trust you? explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD International Conference, pp. 1135–1144 (2016)
- Lundberg, S.M., Lee, S.-I.: A unified approach to interpreting model predictions. In: NeurIPS, vol. 30 (2017)
- 9. Rozemberczki, B., et al.: The Shapley value in machine learning. IJCAI (2022). arXiv:2202.05594
- Kascenas, A.: Anomaly detection in brain imaging. PhD thesis, University of Glasgow (2023)
- Liu, J., et al.: Deep industrial image anomaly detection: a survey. Mach. Intell. Res. 21, 104–135 (2024). issn: 2731-5398
- Schockaert, C., Macher, V., Schmitz, A., et al.: VAE-LIME: deep generative model based approach for local data-driven model interpretability applied to the ironmaking industry. arXiv:2007.10256 (2020)
- Antwarg, L., Miller, R.M., Shapira, B., Rokach, L.: Explaining anomalies detected by autoencoders using SHAP. arXiv:1903.02407 (2019)
- Tritscher, J., Krause, A., Hotho, A., et al.: Feature relevance XAI in anomaly detection: reviewing approaches and challenges. Front. AI 6, 1099521 (2023)
- Oliveira, D.F., et al.: A new interpretable unsupervised anomaly detection method based on residual explanation. IEEE Access 10, 1401–1409 (2021)
- Venkataramanan, S., Peng, K.-C., Singh, R.V., Mahalanobis, A.: Attention guided anomaly localization in images. In: Proceedings of ECCV, pp. 485–503 (2020)
- Kingma, D.P., Welling, M.: Auto-Encoding Variational Bayes. STAT 1050, 1 (2014)
- 18. Goodfellow, I., et al.: Generative adversarial nets. In: NeurIPS, vol. 27 (2014)
- Chow, J., et al.: Anomaly detection of defects on concrete structures with the convolutional autoencoder. Adv. Eng. Inf. 45, 101105 (2020). issn: 1474-0346
- Pang, G., Shen, C., Cao, L., Hengel, A.V.D.: Deep learning for anomaly detection: a review. ACM Comput. Surv. 54 (2021). issn: 0360-0300
- An, J., Cho, S.: Variational autoencoder based anomaly detection using reconstruction probability. Spec. Lect. IE 2, 1–18 (2015)
- Vedaldi, A., Soatto, S.: Quick shift and kernel methods for mode seeking. In: Proceedings of ECCV, pp. 705–718 (2008)
- Rashid, M., Amparore, E., Ferrari, E., Verda, D.: Using stratified sampling to improve LIME image explanations. In: AAAI-2024 (2024)
- Fumagalli, F., Muschalik, M., Kolpaczki, P., Hüllermeier, E., Hammer, B.: SHAP-IQ: unified approximation of any-order shapley interactions. In: NeurIPS, vol. 36 (2024)
- 25. Generative Deep Learning (2024). https://keras.io/examples/generative/