

Unsupervised Detection of Cancerous Regions in Histology Imagery using Image-to-Image Translation

Dejan Štepec^{1,2}, Danijel Škočaj¹

¹University of Ljubljana, Faculty of Computer and Information Science
Večna pot 113, 1000 Ljubljana, Slovenia

²XLAB d.o.o.
Pot za Brdom 100, 1000, Ljubljana, Slovenia

dejan.stepec@xlab.si

Abstract

Detection of visual anomalies refers to the problem of finding patterns in different imaging data that do not conform to the expected visual appearance and is a widely studied problem in different domains. Due to the nature of anomaly occurrences and underlying generating processes, it is hard to characterize them and obtain labeled data. Obtaining labeled data is especially difficult in biomedical applications, where only trained domain experts can provide labels, which often come in large diversity and complexity. Recently presented approaches for unsupervised detection of visual anomalies approaches omit the need for labeled data and demonstrate promising results in domains, where anomalous samples significantly deviate from the normal appearance. Despite promising results, the performance of such approaches still lags behind supervised approaches and does not provide a one-fits-all solution. In this work, we present an image-to-image translation-based framework that significantly surpasses the performance of existing unsupervised methods and approaches the performance of supervised methods in a challenging domain of cancerous region detection in histology imagery.

1. Introduction

The capability to detect anomalies in different data modalities has important applications in different domains, including medical imaging [9]. Detecting visual anomalies is a particularly challenging problem that has recently seen a significant rise of interest, due to the prevalence of deep-learning-based methods. Nevertheless, large part of this success can be attributed to the availability of large-scale labeled data, which is hard to obtain, as anomalies generally occur rarely, in different shapes and forms, and are thus extremely hard or even impossible to label.

Supervised deep-learning-based anomaly detection approaches have seen great success in different industrial and medical application domains [12, 23]. The success of such methods is the most evident in the domains with well-known characterization (and possibly a finite set) of the anomalies and abundance of labeled data. Specific to the detection of visual anomalies, we usually also want to localize the actual anomalous region in the image. Obtaining such detailed labels to learn supervised models is a costly process and in many cases also impossible. There is an abundance of data available in the biomedical domain, but it is usually of much higher complexity and diversity. Usually only trained biomedical experts can annotate such data, preventing large-scale crowd annotation efforts.

Weakly-supervised approaches address such problems by requiring only image-level labels (e.g. disease present or not) and are able to detect and delineate anomalous regions solely from such weakly labeled data, without the need for detailed pixel or patch-level labels [8]. On the contrary, few-shot approaches reduce the number of required labeled samples to the least possible amount [24], which can be further boosted with active learning, where the aim is to come up with the most informative and effective subset of samples for labeling [27, 29].

In an unsupervised setting, only normal appearance samples are available (e.g. healthy, defect-free), which are usually available in larger quantities and are easier to obtain. Deep generative methods, in a form of autoencoders (AE) or generative adversarial networks (GAN), have been recently applied to the problem of unsupervised detection of visual anomalies and have shown promising results in different industrial and medical application domains [21, 6, 4, 7]. Current approaches require normal appearance samples for training, in order to detect and segment deviations from that normal appearance, without the need for labeled data. They usually model normal appearance with low-resolution AE

or GAN models and the overall performance still lags significantly behind supervised approaches.

In this work, we present a novel high-resolution image-to-image translation-based method for unsupervised detection of cancerous regions in histology imagery, that significantly surpasses the performance of existing GAN-based unsupervised approaches in that domain and also approaches the performance of the supervised counterparts.

2. Related Work

2.1. Unsupervised Detection of Visual Anomalies

In an unsupervised setting, we only have access to normal samples (e.g. healthy, defect-free), which are used to model normal visual appearance. This is achieved by learning deep generative models, which results in the capability to generate realistic-looking artificial normal samples [22, 20]. An anomalous (i.e. out-of-distribution) sample is detected by comparing the original input (i.e. query sample) with its reconstruction, by thresholding on a domain-specific similarity metric. This is possible due to the learned manifold of normal appearance and its inability to reconstruct anomalous samples, resulting in higher visual dissimilarity [22, 20].

Different approaches have been proposed for normal appearance modeling, as well as anomaly detection. Learning the normal visual appearance is based on autoencoders (AE) [6], generative adversarial networks (GAN) [20, 21], or combined hybrid models [1, 2] and was already investigated for the histology domain [22]. Most of the approaches learn the space of the normal sample distribution Z , from which latent vectors $z \in Z$ are sampled, that generate the closest normal appearance, to the presented query image. Different solutions have been proposed for latent vector optimization, that are usually independent of the used normal appearance modeling method (i.e. AE, GAN).

Autoencoders (AE) represent an approach towards modeling normal appearance and have a major advantage with their ability to reconstruct images with low reconstruction errors, due to a supervised training objective. Unfortunately, they suffer from memorization and usually produce images that are blurry and of a much lower resolution, quality, and diversity, in comparison with GAN-based approaches [5]. Variational implementation (VAE) turns AE into a generative model, which enables modeling of data distribution and they are usually also combined with a GAN discriminator, in order to produce sharper images [5, 1, 2].

In comparison with autoencoders, GANs do not automatically yield the inverse mapping from the image to latent space, which is needed for closest-looking normal sample reconstruction and consequently anomaly detection. In AnoGAN [20] an iterative optimization approach was proposed to optimize the latent vector z via backpropaga-

tion, using the residual and discrimination losses. In the f-AnoGAN method [21], an autoencoder replaces the iterative optimization procedure, using the trainable encoder and the pre-trained generator (normal appearance modeling), as the decoder. In comparison with AnoGAN [20], StyleGAN2 [17] enables high-resolution normal appearance modeling, while a similar iterative optimization procedure is used for latent space mapping.

2.2. Image-to-Image Translation

The goal of image-to-image translation is to learn a mapping between an input image and an output image of different domains. In the supervised setting, paired corresponding images from different domains are available (e.g. grayscale-color of the same image) and conditional GANs [16] are used to learn the mapping. In the unsupervised setting, only two independent sets of images are available from source and target domains, with no paired examples (e.g. grayscale-color of different images). Cycle-consistency loss, presented with CycleGANs [30], is a predominantly used constraint for such inherently ill-posed problems. CycleGANs [30] enforce original image reconstruction, when mapping the source image to target domain and back, thus capturing special characteristics of the target domain and figuring it how to transfer it to the source domain, while preserving source domain image characteristics.

It was recently discovered, that CycleGANs are masters of steganography [10], as it learns to "hide" information about the source image into images it generates, for almost perfect reconstruction. This intriguing property was exploited by SteGANomaly [4] for unsupervised anomaly detection. They used image-to-image translation to learn a mapping between the healthy brain MR images and a simulated distribution of healthy brain MR images with lower entropy. CycleGANs encode the source image information into the target image with a nearly imperceptible, high-frequency signal [10], thus enabling the reconstruction of unseen anomalous samples during the inference. In SteGANomaly [4] they alleviate this problem by removing high frequency, low amplitude signal during training using Gaussian filtering in the target domain, before performing a complete cycle. The choice of intermediate distribution with lower entropy is very important, for the Gaussian filtering to effectively remove the hidden information, that is not relevant to the target domain. Together with the size of the Gaussian kernel, which needs to be manually fine-tuned, this represents a major limiting factor for general applicability.

2.3. Anomaly Detection in Medical Imagery

In this work, we particularly focus on a challenging task of cancerous region detection from gigapixel histol-




Figure 1: Our proposed unsupervised anomaly detection method, based on the image-to-image translation. We disentangle a latent space into shared content and style spaces, implemented via domain-specific (blue and green colors) encoders E and decoders G . Anomaly detection is performed with an example-guided image translation.

ogy imagery, which has been already addressed in a supervised [12], as well as in a weakly-supervised setting [8]. There is a limited number of approaches in the literature that would approach that problem in an unsupervised fashion [19, 25]. Extremely large histology imagery (patch-based processing) and the highly variable appearance of the different tissue regions represent a unique challenge for existing unsupervised approaches. Unsupervised approaches for the detection of visual anomalies are also applied to other biomedical imaging data, especially brain MRI data analysis, where brain lesion segmentation is the predominantly addressed problem [11, 5, 4].

3. Methodology

We argue that image-to-image translation can be effective for unsupervised detection of visual anomalies and that this can be achieved with a direct mapping between unpaired sets of healthy cohorts, with an appropriate architecture, which successfully disentangles content, that needs to be preserved, apart from the style, which needs to change. CycleGAN [30] based methods perform this implicitly using the adversarial and cycle-consistency losses, which assumes a strong bijection between the two domains - resulting in steganography [10]. Inspired by the multimodal image-to-image translation methods [15, 18], we propose an example guided image translation method (Figure 1), which in comparison with SteGANomaly [4] enables anomaly detection without cycle-reconstruction during the inference, specially crafted intermediate domain distribution, and Gaussian filtering. Similar to MUNIT [15], we

assume that the latent space of images can be decomposed into content and style spaces. We also assume that images in both domains share a common content space C , as well as style space S . This differs from MUNIT [15], where style space is not shared, due to semantically different domains X and Y . Similar to MUNIT [15], our translation model consists out of encoder E_{ij} and decoder G_j networks for each space $i \in \{C, S\}$ and domains $j \in \{X, Y\}$. Those subnetworks are used for autoencoding, as well as cross-domain translation, by interchanging encoders and decoders from different domains. Style latent codes s_x and s_y are randomly drawn for cross-domain translation and used as Adaptive Instance Normalization (AdaIN) [14] parameters in residual blocks of decoders G_x and G_y , additionally transformed by a multilayer perceptron (MLP) network f . For autoencoding, E_{sx} and E_{sy} encoders are used directly, to extract style codes. This randomness during cross-domain translation in training prevents the effect of memorization, largely present in autoencoder-based anomaly detection approaches. Different losses between inputs and its reconstructions (e.g. x and \hat{x}) are used to train autoencoders (L1 loss), cross-domain translations (adversarial loss), as well as reconstructions (cycle-consistency loss). The architecture and training objectives closely follow the implementation of MUNIT [15], with the exception of the selection of X and Y domains.

During anomaly detection (Figure 1), an input image x is encoded with E_{cx} , to produce content vector c_x , which is then joined with the style code s_y , extracted from the original image x , with the style encoder E_{sy} of the target domain Y . This presents an input to decoder G_y , which gen-

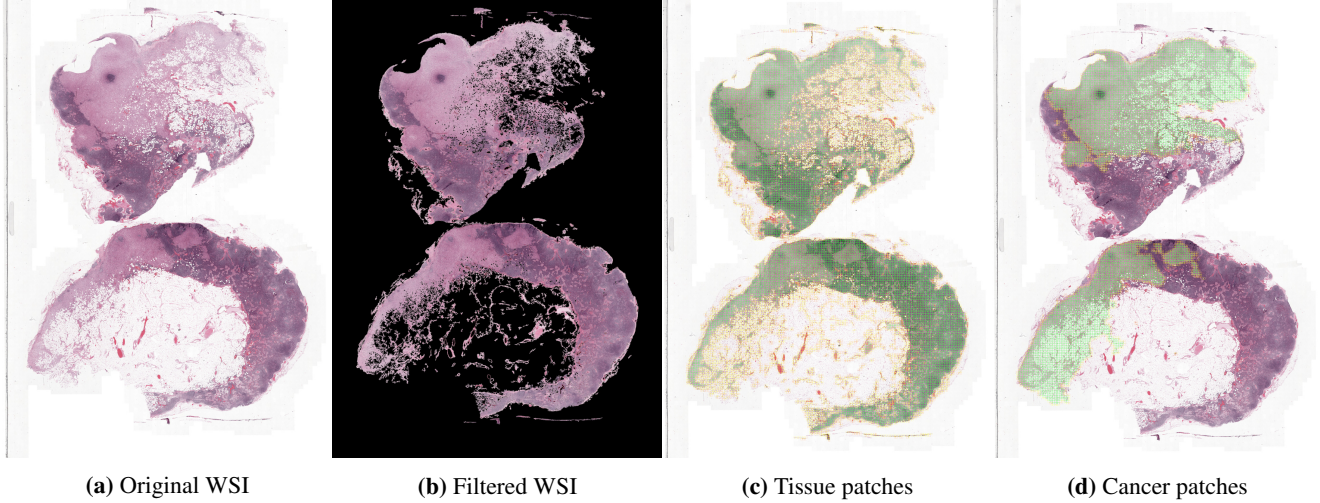


Figure 2: Preprocessing of the original WSI presented in a) consists of b) filtering tissue sections and c) extracting tissue patches, based on the tissue (green $\geq 90\%$, orange $\leq 10\%$ and yellow in-between) and d) cancerous region coverage (green $\geq 90\%$, orange $\leq 30\%$ and yellow in-between). Best viewed in a digital version with zoom.

erates y . This is basically an example guided image translation, used also in MUNIT [15] and DRIT++ [18] methods. Content-style space decomposition is especially well suited for histopathological analysis due to different staining procedures, which causes the samples to significantly deviate in their visual appearance. Style-guided translation ensures that the closest looking normal appearance is found, taking into account also the staining appearance. We then measure an anomaly score using distance metric d (e.g. perceptual LPIPS distance [28] or Structure Similarity Index (SSIM) [26]), between the original image x and its reconstruction \hat{x} .

4. Experiments and Results

4.1. Histology Imagery Dataset

We evaluate the proposed anomaly detection pipeline on whole-slide histology images (WSI), which are used for diagnostic assessment of the spread of the cancer. This particular problem was already addressed in a supervised setting [12], as a competition¹, with provided clinical histology imagery and ground truth data. A training dataset with (n=110) and without (n=160) cancerous regions is provided, as well as a test set of 129 images (49 with and 80 without anomalies). Raw histology imagery, presented in Figure 2a, is first preprocessed, in order to extract the tissue region (Figure 2b). We used the approach from IBM², which utilizes a combination of morphological and color space filtering operations and was also used in prior work [22] of synthesizing realistically looking histology

samples. Patches of 512 x 512 are then extracted from the filtered image and filtered according to the tissue (Figure 2c) and cancer (Figure 2d) coverage. We only use patches with tissue and cancerous region coverage over 90% (i.e. green patches).

We train the models on random 80,000 healthy tissue patches extracted from a training set of healthy and cancerous (cancer coverage=0%) WSIs (n=270). The baseline supervised approach is trained on randomly extracted healthy (n=25,000) and cancerous patches (n=25,000). The methods are evaluated on healthy (n=7673) and cancerous (n=16,538) patches extracted from a cancerous test set of WSIs (n=49). We mix healthy training patches of both cohorts (i.e. healthy patches from cancerous WSIs) in order to demonstrate the robustness of the proposed approach against a small percentage of possibly contaminated healthy appearance data (e.g. non-labeled isolated tumor cells).

4.2. Unsupervised Anomaly Detection

We compare the proposed method against GAN-based f-AnoGAN [21] and StyleGAN2 [17] methods. Both methods separately model normal appearance and perform latent space mapping for anomaly detection. The f-AnoGAN method models normal appearance using Wasserstein GANs (WGAN) [3], which is limited to a resolution of 64^2 and uses an encoder-based fast latent space mapping approach. The StyleGAN2 method enables high-resolution image synthesis (up to 1024^2) and also implements an iterative optimization procedure, based on Learned Perceptual Image Patch Similarity (LPIPS) [28] distance metric. We evaluate the performance of the proposed and StyleGAN2 methods on patches of 512^2 , while center-cropped

¹<https://camelyon16.grand-challenge.org/>

²<https://github.com/CODAIT/deep-histopath>




Figure 3: Distribution of anomaly scores on healthy and cancerous histology imagery patches (a) for the proposed method (SSIM metric), (b) proposed method (LPIPS metric), (c) f-AnoGAN (original metric), and (d) StyleGAN2 (LPIPS metric). Results for the proposed and StyleGAN2 methods are reported for 512^2 patches, while 64^2 patches are used for f-AnoGAN.

64^2 patches are used for the f-AnoGAN method. Additionally, we compare the performance against the supervised DenseNet-121 [13] baseline model, trained and evaluated on 512^2 patches. We evaluate the proposed method using Structural Similarity Index Measure (SSIM) [26] and LPIPS reconstruction error metrics as an anomaly score. We use the same metrics as also as an alternative to original f-AnoGAN anomaly score implementation, as well as to measure StyleGAN2 reconstruction errors.

We first evaluate the methods by inspecting the distribution of anomaly scores across healthy and cancerous patches, as presented in Figure 3. We compare our proposed approach (Figures 3a and 3b) against f-AnoGAN (Figure 3c) and StyleGAN2 (Figure 3d) methods and report significantly better distribution disentanglement.

Area under the ROC curve (AUC) scores are reported in Table 1 for all the methods and different anomaly scores. Corresponding ROC curves are presented in Figure 4. We also report F1 and classification accuracy measures, calculated at the Youden index of the ROC curve. We notice that the performance of the proposed method approaches the performance of the supervised baseline in terms of both reconstruction error metrics (i.e. LPIPS and SSIM). The performance of the f-AnoGAN significantly improves using SSIM and LPIPS metrics, in comparison with the originally proposed anomaly score. This shows the importance of the selection of the appropriate reconstruction error metric. The styleGAN2 method shows good distribution disentanglement using the LPIPS distance metric, while the SSIM metric fails to capture any significant differences between the two different classes (i.e. healthy and anomalous). The proposed method demonstrates consistent performance across both anomaly score metrics, as well as different evaluation measures.

In Figure 5 we present example reconstructions of healthy and cancer tissue samples for all the methods. The proposed method reconstructs healthy samples much more accurately in comparison with the StyleGAN2 method. Some level of *artificial healing* is visible on cancerous samples (i.e. visual appearance much more closely reassem-

Table 1: Performance statistics (F1, Classification Accuracy - CA) calculated at Youden index of Receiver Operating Characteristic (ROC) curve and the corresponding area under the ROC curve (AUC) and Average Precision (AP) scores summarizing ROC and Precision-Recall (PR) curves.

	AUC	AP	F1	CA
Supervised	0.954	0.974	0.925	0.901
Proposed (SSIM)	0.947	0.976	0.920	0.895
Proposed (LPIPS)	0.900	0.914	0.886	0.847
StyleGAN2 (LPIPS)	0.908	0.940	0.872	0.836
StyleGAN2 (SSIM)	0.580	0.711	0.674	0.588
f-AnoGAN (original)	0.650	0.443	0.502	0.637
f-AnoGAN (SSIM)	0.887	0.916	0.886	0.846
f-AnoGAN (LPIPS)	0.865	0.902	0.875	0.830




Figure 4: ROC curves for different methods and different reconstruction error metrics (i.e. anomaly scores).

bles healthy samples). The f-AnoGAN method is only able to operate on 64^2 resolution tissue samples, where similarly we notice better reconstructions of healthy ap-




Figure 5: Example reconstructions of ground truth healthy and cancerous tissue samples with the proposed and StyleGAN2 methods for 512^2 resolution and f-AnoGAN method for 64^2 . Best viewed in a digital version with zoom.

pearances. The visual patterns in the StyleGAN2 method reconstructions demonstrate significantly lower variability in the visual appearance, especially in comparison with its demonstrated capability to synthesize realistically looking, highly variable, high-resolution histology tissue samples [22]. In comparison, the proposed image-to-image translation-based enables high-resolution image synthesis, as well as example-based reconstruction, that can be effectively utilized for the detection of visual anomalies.

5. Conclusion

Detection of visual anomalies is an important process in many domains and recent advancements in deep generative-based methods have shown promising results towards applying them in an unsupervised fashion. This has sparked research in many domains, that did not benefit much from traditional supervised deep-learning-based approaches, due to difficulties in obtaining large quantities of labeled data. The medical image analysis domain is one such notable example, where the availability of imaging data in large quantities is usually not a problem, but the real challenge lies in the scarcity of human-expert annotations.

In this work, we presented an image-to-image translation-based unsupervised approach that significantly surpasses the performance of existing GAN-based unsupervised approaches for the detection of visual

anomalies in histology imagery and also approaches the performance of supervised methods. The method is capable of closely reconstructing presented healthy histology tissue samples, while unable to reconstruct cancerous ones and is thus able to detect such samples with an appropriate visual distance measure.

The image-to-image translation-based framework offers a promising multi-task platform for a wide range of problems in the medical domain and can be now further extended with the capabilities for anomaly detection. Additional research is needed to investigate effectiveness in other biomedical modalities, as well as to exploit the benefits of using such a framework in a multi-task learning setting.

Acknowledgment

This work was partially supported by the European Commission through the Horizon 2020 research and innovation program under grant 826121 (iPC) and by the Slovenian Research Agency (ARRS) project J2-9433 (DIVID).

References

- [1] Samet Akcay, Amir Atapour-Abarghouei, and Toby P Breckon. Ganomaly: Semi-supervised anomaly detection via adversarial training. In *ACCV*, pages 622–637. Springer, 2018.

- [2] Samet Akçay, Amir Atapour-Abarghouei, and Toby P Breckon. Skip-ganomaly: Skip connected and adversarially trained encoder-decoder anomaly detection. In *IJNN*, pages 1–8. IEEE, 2019.
- [3] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein Generative Adversarial Networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 214–223, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
- [4] Christoph Baur, Robert Graf, Benedikt Wiestler, Shadi Albarqouni, and Nassir Navab. Steganomaly: Inhibiting cy-clegan steganography for unsupervised anomaly detection in brain mri. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 718–727. Springer, 2020.
- [5] Christoph Baur, Benedikt Wiestler, Shadi Albarqouni, and Nassir Navab. Deep Autoencoding Models for Unsupervised Anomaly Segmentation in Brain MR Images. In Alessandro Crimi, Spyridon Bakas, Hugo Kuijf, Farahani Keyvan, Mauricio Reyes, and Theo van Walsum, editors, *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*, pages 161–169, Cham, 2019. Springer International Publishing.
- [6] Christoph Baur, Benedikt Wiestler, Shadi Albarqouni, and Nassir Navab. Scale-space autoencoders for unsupervised anomaly segmentation in brain mri. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 552–561. Springer, 2020.
- [7] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4183–4192, 2020.
- [8] Gabriele Campanella, Matthew G Hanna, Luke Geneslaw, Allen Miralflor, Vitor Werneck Krauss Silva, Klaus J Busam, Edi Brogi, Victor E Reuter, David S Klimstra, and Thomas J Fuchs. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature medicine*, 25(8):1301–1309, 2019.
- [9] Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly Detection: A Survey. *ACM Comput. Surv.*, 41(3):15:1–15:58, July 2009.
- [10] Casey Chu, Andrey Zhmoginov, and Mark Sandler. Cyclegan, a master of steganography. *arXiv preprint arXiv:1712.02950*, 2017.
- [11] Alessandro Crimi and Spyridon Bakas. *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries: 5th International Workshop, BrainLes 2019, Held in Conjunction with MICCAI 2019, Shenzhen, China, October 17, 2019, Revised Selected Papers, Part I*, volume 11992. Springer Nature, 2020.
- [12] Babak Ehteshami Bejnordi, Mitko Veta, Paul Johannes van Diest, Bram van Ginneken, Nico Karssemeijer, Geert Litjens, Jeroen A. W. M. van der Laak, , and the CAMELYON16 Consortium. Diagnostic Assessment of Deep Learning Algorithms for Detection of Lymph Node Metastases in Women With Breast Cancer. *JAMA*, 318(22):2199–2210, 12 2017.
- [13] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, pages 4700–4708, 2017.
- [14] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1501–1510, 2017.
- [15] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 172–189, 2018.
- [16] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [17] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, pages 8110–8119, 2020.
- [18] Hsin-Ying Lee, Hung-Yu Tseng, Qi Mao, Jia-Bin Huang, Yu-Ding Lu, Maneesh Singh, and Ming-Hsuan Yang. Dri++: Diverse image-to-image translation via disentangled representations. *International Journal of Computer Vision*, pages 1–16, 2020.
- [19] Louise Naud and Alexander Lavin. Manifolds for unsupervised visual anomaly detection. *arXiv preprint arXiv:2006.11364*, 2020.
- [20] 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 Marc Niethammer, Martin Styner, Stephen Aylward, Hongtu Zhu, Ipek Oguz, Pew-Thian Yap, and Dinggang Shen, editors, *Information Processing in Medical Imaging*, pages 146–157, Cham, 2017. Springer International Publishing.
- [21] 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.
- [22] Dejan Štepec and Danijel Škočaj. Image synthesis as a pretext for unsupervised histopathological diagnosis. In *International Workshop on Simulation and Synthesis in Medical Imaging*, pages 174–183. Springer, 2020.
- [23] Domen Tabernik, Samo Šela, Jure Skvarč, and Danijel Škočaj. Segmentation-Based Deep-Learning Approach for Surface-Defect Detection. *Journal of Intelligent Manufacturing*, May 2019.
- [24] Yu Tian, Gabriel Maicas, Leonardo Zorron Cheng Tao Pu, Rajvinder Singh, Johan W Verjans, and Gustavo Carneiro. Few-shot anomaly detection for polyp frames from colonoscopy. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 274–284. Springer, 2020.

- [25] Nina Tuluptceva, Bart Bakker, Irina Fedulova, Heinrich Schulz, and Dmitry V Dylov. Anomaly detection with deep perceptual autoencoders. *arXiv preprint arXiv:2006.13265*, 2020.
- [26] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [27] Lin Yang, Yizhe Zhang, Jianxu Chen, Siyuan Zhang, and Danny Z Chen. Suggestive annotation: A deep active learning framework for biomedical image segmentation. In *International conference on medical image computing and computer-assisted intervention*, pages 399–407. Springer, 2017.
- [28] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595, 2018.
- [29] Hao Zheng, Lin Yang, Jianxu Chen, Jun Han, Yizhe Zhang, Peixian Liang, Zhuo Zhao, Chaoli Wang, and Danny Z Chen. Biomedical image segmentation via representative annotation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5901–5908, 2019.
- [30] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.