

Unsupervised Anomaly Detection In Digital Pathology Using GANs

Milda Pocevičiūtė, Gabriel Eilertsen and Claes Lundström

The self-archived postprint version of this conference paper is available at Linköping University Institutional Repository (DiVA):

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-178631>

N.B.: When citing this work, cite the original publication.

Pocevičiūtė, M., Eilertsen, G., Lundström, C., (2021), Unsupervised Anomaly Detection In Digital Pathology Using GANs, 2021. *IEEE 18th International Symposium On Biomedical Imaging (ISBI)*, pp. 1878-1882. <https://doi.org/10.1109/ISBI48211.2021.9434141>

Original publication available at:

<https://doi.org/10.1109/ISBI48211.2021.9434141>

©2021 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

<http://www.ieee.org/>

IEEE Copyright Notice

© 20XX IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

UNSUPERVISED ANOMALY DETECTION IN DIGITAL PATHOLOGY USING GANS

Milda Pocevičiūtė^{*†}, Gabriel Eilertsen^{*†}, Claes Lundström^{*†§}

^{*}Department of Science and Technology, Linköping University, Sweden

[†]Center for Medical Image Science and Visualization, Linköping University

[§]Sectra AB, Linköping, Sweden

ABSTRACT

Machine learning (ML) algorithms are optimized for the distribution represented by the training data. For outlier data, they often deliver predictions with equal confidence, even though these should not be trusted. In order to deploy ML-based digital pathology solutions in clinical practice, effective methods for detecting anomalous data are crucial to avoid incorrect decisions in the outlier scenario. We propose a new unsupervised learning approach for anomaly detection in histopathology data based on generative adversarial networks (GANs). Compared to the existing GAN-based methods that have been used in medical imaging, the proposed approach improves significantly on performance for pathology data. Our results indicate that histopathology imagery is substantially more complex than the data targeted by the previous methods. This complexity requires not only a more advanced GAN architecture but also an appropriate anomaly metric to capture the quality of the reconstructed images.

Index Terms— digital pathology, anomaly detection, GAN, unsupervised learning

1. INTRODUCTION

Machine learning, and in particular deep learning, is showing great potential in pathology applications, with state-of-the-art methods achieving pathologist-level performance [1]. To guarantee reliable use of these approaches in practice, we have to detect the situations when the algorithm encounters outlier data and cannot be trusted. Hence, anomaly detection is an important topic in digital pathology, making it possible to flag potential mispredictions to the pathologist.

A fundamental challenge in anomaly detection is that the anomaly is unknown or unavailable at training time, requiring an unsupervised training strategy. For this purpose, frameworks based on generative adversarial networks (GANs) [2] have shown great promise [3, 4, 5]. As GANs learn to capture data distribution characteristics, they can be used to estimate if a sample fits the distribution or should be considered an outlier/anomaly. With the fast development in GAN architectures and training strategies, e.g. including Wasserstein GAN (WGAN) [6], Progressive Growing GAN (PG-GAN) [7] and




Fig. 1. Scheme of GAN-based unsupervised anomaly detection. The original input image x is fed through projector E which outputs a latent vector z . This latent vector is used to generate a reconstruction image via generator G . An anomaly score is a measure of some difference between the input image x and its reconstruction \hat{x} .

StyleGAN2 [8], the potential for GAN-based anomaly detection is constantly growing. For an in-depth introduction to the field of anomaly detection, we refer to a recent survey by Bulusu et al [9].

In this work, we introduce s2-AnoGAN, a GAN-based unsupervised anomaly detection method. Our approach stems from two important observations for the histopathology case. First, a powerful GAN architecture is needed to capture the high data complexity. Second, anomaly scores relying on pixel-wise comparisons are not suitable for this application domain. Instead, we propose a new anomaly score that uses edge information. We show that in combination these two components allow for substantial improvements over existing methods when evaluated on pathology data.

2. METHODS AND MATERIALS

We evaluate the performance on pathology data of our proposed s2-AnoGAN method, and compare it to two previous methods: f-AnoGAN [4] and a framework introduced by Berg et al. [5] which we denote pg-AnoGAN.

Fig. 1 illustrates the general idea behind unsupervised anomaly detection based on GANs. Given a trained generator G , and a projector E (estimated inverse of G), the anomaly


score is a distance function $\mathcal{A}(x, \hat{x})$ that measures the difference between input image x and its reconstructed projection $\hat{x} = G(E(x))$. The anomaly detection frameworks differ in a choice of GAN, projection method, and anomaly measure \mathcal{A} .

Our proposed s2-AnoGAN anomaly detector is based on the StyleGAN2 algorithm (configuration F) and its projector [8]. The projector is running a gradient-based optimisation to minimise the LPIPS distance [10] between the input image and its reconstruction. The f-AnoGAN framework is based on Wasserstein GAN with gradient penalty [6] and the projector is trained separately using *image-to-image* mapping approach. The pg-AnoGAN framework uses PG-GAN [7] and designs the encoder inspired by [11]. For more details, please refer to the aforementioned papers [4], [5] and [8]. The training of each algorithm was stopped once the perceived quality of the generated images was not changing for several epochs, resulting in 50 training epochs for f-AnoGAN, and 80 epochs for the pg-AnoGAN and s2-AnoGAN frameworks. Following the original implementation, f-AnoGAN encoder was trained separately for 50,000 iterations.


Another differing aspect of the three frameworks is the anomaly score used. Our proposed framework includes a new anomaly score defined by $\mathcal{A}_{Canny} = N_{x_{Canny}} - N_{\hat{x}_{Canny}}$, where x_{Canny} and \hat{x}_{Canny} are the output images by the Canny edge detector [12] applied to the original images and their reconstructions, respectively. N_y is the count of non-zero pixels in a binary image y . In our experiments, the parameters were set to kernel 5, sigma 3 as well as 100 and 200 for the lower and the upper thresholds, respectively. The intuition behind the \mathcal{A}_{Canny} metric is that even if the exact position of the features in the reconstructed images may vary, the overall edge characteristics should be similar in the original patch and its reconstruction given that the projection was successful.

The anomaly score of f-AnoGAN is given by $\mathcal{A}_{f-AnoGAN} = \mathcal{A}_{MSE} + \kappa \cdot \mathcal{A}_D$, where \mathcal{A}_{MSE} is a mean square error (MSE) between the original image and its reconstruction, \mathcal{A}_D is the MSE between discriminator features of original image and its reconstruction, and κ is a weight factor set to 1. The anomaly score for pg-AnoGAN is defined as $\mathcal{A}_{pg-AnoGAN} = \alpha \cdot \mathcal{A}_{res} + (1 - \alpha) \cdot \mathcal{A}_0(z)$, where \mathcal{A}_{res} is the l_2 -norm between original image and its reconstruction after min-max normalisation, $\mathcal{A}_0(z)$ is the distance to the origin of the latent vector, and α is a weight factor set to 0.05.


There are many possible types of anomalies in pathology applications. For our experiments, we adopted a setup where images of healthy tissue constituted normal data, whereas tumour images were considered as anomalies. While we do not propose this setup to represent a feasible anomaly detection scenario in the clinic as such, it was chosen as a suitable proxy task to evaluate anomaly detection methods. The difference between healthy and tumorous tissue stems from complex histology characteristics, and this is a main challenge for any anomaly detection scenario in our domain.



(a) ROC curves of anomaly detection



(b) Histogram of \mathcal{A}_{MSE} (f-AnoGAN)



(c) Histogram of \mathcal{A}_{Canny} (s2-AnoGAN)

Fig. 2. (a) shows ROC curves with the corresponding AUC achieved by the tested frameworks, and the baseline (for which only the number of edges is used as anomaly score). (b) is the histogram of \mathcal{A}_{MSE} scores computed on images and their projections done by f-AnoGAN. (c) is the histogram of anomaly scores computed by our s2-AnoGAN framework.

	f-AnoGAN	pg-AnoGAN	s2-AnoGAN
$\mathcal{A}_{f-AnoGAN}$	0.21	0.32	0.37
$\mathcal{A}_{pg-AnoGAN}$	0.74	0.70	0.71
\mathcal{A}_{Canny}	0.75	0.57	0.82

Table 1. AUC values of anomaly detection when combined with the anomaly measure of each framework, respectively.

The data used for the experiments consists of two types of patches, healthy and tumorous tissue, extracted from the CAMELYON17 data set [13], which contains whole-slide images (WSI) of hematoxylin and eosin (H&E) stained lymph node sections. Patches are 64×64 pixels, matching the image size used in the f-AnoGAN paper [4]. We worked with grayscale images, in order to make fair comparisons to previous results, and to reduce the effect of colour variation in the images due to being sourced from different hospitals and scanners. The training set consisted of 87 521 healthy patches, while the held-out test data contained 9 729 healthy and 9 729 tumour patches. For the implementation of f-AnoGAN, we also extracted 12 503 healthy patches for training the encoder.




Fig. 3. Examples of test data and their projections by f-AnoGAN, pg-AnoGAN and s2-AnoGAN frameworks.

3. RESULTS AND DISCUSSION

3.1. Anomaly detection results

We tested the three anomaly frameworks as well as other combinations of GAN+projector and anomaly measure. The performance is measured in terms of AUC (Area Under the Receiver Operating Characteristics Curve). In Table 1, we see that the performance of $\mathcal{A}_{f-AnoGAN}$ and $\mathcal{A}_{pg-AnoGAN}$ on the pathology data is considerably worse than reported in [4] and [5] for all GAN frameworks. The highest AUC score of 0.82 was achieved by our framework: StyleGAN2 with gradient descent projector and the anomaly metric \mathcal{A}_{Canny} .

The ROC curves for the three frameworks are presented in Sub-figure (a) of Fig. 2, further illustrating the top result from s2-AnoGAN. As a sanity check, we also include a baseline measure not involving any GAN training: using $N_{x_{Canny}}$ as an anomaly score (number of edge pixels in the input image). We can see that s2-AnoGAN and pg-AnoGAN outperformed the baseline by a large margin but f-AnoGAN actually achieved a lower AUC. To gain a better understanding of the results, in the next sections we will discuss the two main factors affecting the anomaly detection performance: GAN expressivity and the chosen anomaly metric.

3.2. Quality of learned underlying distribution

GANs and their projectors were evaluated based on qualitative as well as quantitative inspection. Here we mainly focus

	f-AnoGAN	pg-AnoGAN	s2-AnoGAN
$LPIPS$	0.36	0.36	0.64
\mathcal{A}_D	0.24	0.35	0.36
\mathcal{A}_{MSE}	0.21	0.32	0.37
\mathcal{A}_{res}	0.19	0.33	0.27
$\mathcal{A}_0(z)$	0.75	0.70	0.71

Table 2. AUC values achieved by different anomaly measures. The latter four are components of the measures used by f-AnoGAN and pg-AnoGAN.

on the healthy data as it was the target of the training. The overall qualitative impression is that s2-AnoGAN was most successful in reproducing realistic images, as illustrated in Fig. 3. From healthy example patches 1-3, we can see how s2-AnoGAN reconstructs typical healthy samples, with several visible cell nuclei, better than the other two frameworks. The amount of cell nuclei and their positions shows a better agreement with the input in these cases.

This observation is also supported by the peak signal-to-noise ratio (PSNR) between the original healthy image and its reconstruction. Average PSNR score on the healthy data for s2-AnoGAN is 28.7 (standard deviation: 1.7) while for f-AnoGAN and pg-AnoGAN frameworks it is 16.6 (standard deviation: 3.3) and 13.4 (standard deviation: 2.6), respectively. The combination of the highest PSNR mean value and the smallest standard deviation achieved by s2-AnoGAN indicates that the framework will be better suited for a distance based anomaly score.

However, less common types of healthy patches, such as example 4 in Fig. 3, are poorly reconstructed by all three frameworks. For reference, Fig. 3 also include two tumour images, showing the expected effect of poor reconstruction.

3.3. Problem with pixel-wise anomaly measures

An interesting result is that the pixel-wise difference measures work poorly for anomaly detection in pathology data. From Fig. 2(b), we can see that the anomaly score using \mathcal{A}_{MSE} actually correlated with healthy tissue, the opposite of what was expected. In theory, this indicates that the projections of tumour patches are better than the healthy. The same effect is seen for the related \mathcal{A}_{res} score (see Table 2). The explanation is two-fold. The MSE values are based on the locational comparison of the pixels, hence, it can be high even when the result is realistic. For example, when the exact positions of cell nuclei turn up to be slightly moved in the reconstructed images (see Fig. 3). Conversely, the reconstructions of tumour tissue can have the characteristics of a blurred image, giving relatively low MSE even though the realism is poor.

To overcome the problems with the direct comparisons between input and reconstructed pathology images, we introduced the \mathcal{A}_{Canny} anomaly score. We also investigated other

distance functions not affected by the exact locations of pixels, including LPIPS [10]. LPIPS is a recent learning-based measure that has shown good performance in image comparison problems. LPIPS achieved an AUC of 0.64 (see Table 2) when combined with StyleGAN2 and its projector versus AUC of 0.82 when using \mathcal{A}_{Canny} (see Table 1). This is likely due to LPIPS being based on features learned by a neural network trained on natural images, which seems to be a poor representation of the important features in pathology data.

4. CONCLUSIONS AND FUTURE WORK

In this work we explored unsupervised anomaly detection with GANs and how it works on histopathology data. Due to the more complex structures in the images than used in previous work, the current methods did not perform well and were outperformed by our introduced s2-AnoGAN framework with anomaly score based on canny edges. One limitation of the StyleGAN2 projector is that it is based on gradient descent. Thus, it is much slower than the projection methods used by Schlegl et al. [4] and Berg et al. [5]. For future work, we will explore different encoder-based projectors and tailor them to work well with the StyleGAN2 architecture. Furthermore, we will explore if the results could be further improved by deriving other metrics for comparing original pathology images with their reconstructions.

5. ACKNOWLEDGMENTS

This work was supported by the Swedish e-Science Research Center and VINNOVA, grant 2017-02447. Claes Lundström is an employee and shareholder of Sectra AB.

6. REFERENCES

- [1] Wouter Bulten, Maschenka Balkenhol, Jean-Joël Awoumou Belinga, Américo Brilhante, Aslı Çakır, Xavier Farré, Katerina Geronatsiou, Vincent Molinié, Guilherme Pereira, Paromita Roy, Günter Saile, Paulo Salles, Ewout Schaafsma, Joëlle Tschui, Anne-Marie Vos, Hester van Boven, Robert Vink, Jeroen van der Laak, Christina Hulsbergen-van de Kaa, and Geert Litjens, “Artificial intelligence assistance significantly improves gleason grading of prostate biopsies by pathologists.” 2020.
- [2] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, “Generative adversarial nets,” in *Proceedings of the 27th International Conference on Neural Information Processing Systems (NIPS 2014)*, 2014, p. 2672–2680.
- [3] 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 *Proceedings of the 25th International Conference on Information Processing in Medical Imaging (IPMI 2017)*, 2017, p. 146.
- [4] Thomas Schlegl, Philipp Seeböck, Sebastian Waldstein, Georg Langs, and Ursula Schmidt-Erfurth, “f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks,” *Medical Image Analysis*, vol. 54, 01 2019.
- [5] Amanda Berg, Jörgen Ahlberg, and Michael Felsberg, “Unsupervised adversarial learning of anomaly detection in the wild,” in *24th European Conference on Artificial Intelligence (ECAI 2020)*, 2020.
- [6] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville, “Improved training of wasserstein gans,” in *Advances in neural information processing systems (NeurIPS 2017)*, 2017, pp. 5767–5777.
- [7] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen, “Progressive growing of GANs for improved quality, stability, and variation,” in *International Conference on Learning Representations (ICLR 2018)*, 2018.
- [8] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila, “Analyzing and improving the image quality of stylegan,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2020)*, 2020, pp. 8110–8119.
- [9] Saikiran Bulusu, Bhavya Kaikhura, Bo Li, Pramod K Varshney, and Dawn Song, “Anomalous example detection in deep learning: A survey,” *IEEE Access*, vol. 8, pp. 132330–132347, 2020.
- [10] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang, “The unreasonable effectiveness of deep features as a perceptual metric,” in *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR 2018)*, 2018, pp. 586–595.
- [11] Sudipto Mukherjee, Himanshu Asnani, Eugene Lin, and Sreeram Kannan, “Clustergan : Latent space clustering in generative adversarial networks.” *CoRR*, vol. abs/1809.03627, 2018.
- [12] John Canny, “A computational approach to edge detection,” *IEEE Transactions on pattern analysis and machine intelligence*, , no. 6, pp. 679–698, 1986.

- [13] Geert Litjens, Peter Bandi, Babak Ehteshami Bejnordi, Oscar Geessink, Maschenka Balkenhol, Peter Bult, Altuna Halilović, Meyke Hermsen, Rob Loo, Rob Vogels, Quirine Manson, Nikolas Stathonikos, Alexi Baidoshvili, Paul Diest, Carla Wauters, Marcory van Dijk, and Jeroen van der Laak, “1399 H&E-stained sentinel lymph node sections of breast cancer patients: The camelyon dataset,” *GigaScience*, vol. 7, 05 2018.

7. COMPLIANCE WITH ETHICAL STANDARDS

This research study was conducted retrospectively using human subject data made available in open access by Litjens et al [13]. Ethical approval was not required as confirmed by the license attached with the open access data.