



Unsupervised Liver Tumor Segmentation with Pseudo Anomaly Synthesis

Zhaoxiang Zhang^(✉) , Hanqiu Deng^(ID) , and Xingyu Li^(ID) 

University of Alberta, Edmonton, AB T6G 2R3, Canada
{zhaoxia2,hanqiu1,xingyu}@ualberta.ca

Abstract. Liver lesion segmentation is a challenging task. Liver lesions often appear as regional heterogeneity in various shapes and intensities, while collecting a comprehensive dataset for supervised learning is costly. To address this issue, this study formulates unsupervised liver tumor segmentation as an anomaly segmentation problem and presents a pseudo-supervised anomaly segmentation solution with synthetic anomalies. In this regard, we investigate two fundamental, yet under-explored questions: (1) how to generate anomalies? and (2) how to address a covariant shift between synthesis data and real tumor samples in model training? To the first question, instead of fabricating anomalies approximating the known abnormal patterns, we propose to generate anomalies spreading over a broader spectrum to encourage a model to learn the cluster boundary of normal samples. Our rationale toward the second question suggests light training on synthesis data for model generalizability. Based on these insights, this study incorporates a random-shaped anomaly synthesis module and two-stage training strategy into the DRAEM architecture for unsupervised liver tumor segmentation. Experiments on the public benchmark show that the proposed method trained on various synthetic anomalies has good generalizability on real tumor and achieves a comparable performance to prior arts. Our code is available at: <https://github.com/nono-zz/LiTs-Segmentation>.

Keywords: Lesion Segmentation · Unsupervised Learning · Anomaly Synthesis · Two-stage Training · CT

1 Introduction

Liver tumors are one of the leading causes of cancer-related deaths, and accurately segmenting them in medical images such as computed tomography (CT) is crucial for early detection and diagnosis. While supervised tumor segmentation methods show promising results, their performance is heavily dependent on high-quality annotated data, which can be expensive to obtain. Furthermore,

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/978-3-031-44689-4_9.

due to the high heterogeneity of tumors, the generalizability of supervised models may be limited in identifying rare lesions or anomalies. Recently, there is an increased interest in treating tumors as anomalies in medical images and exploring unsupervised learning approaches, i.e. anomaly segmentation, to address the aforementioned challenges. In the context of unsupervised anomaly segmentation, a model is expected to identify and segment potential abnormalities by learning from a healthy cohort of patients during model training.

Prior arts in unsupervised anomaly segmentation can be categorized into two paradigms. The first paradigm stems from anomaly detection, where a model predicts if a query is normal or not and image regions that contribute most to this results are taken as the anomaly segmentation results. Among various anomaly detection methods, generative models, such as Variational Auto-Encoders (VAEs) [13], Generative Adversarial Networks (GANs) [8], Denoising Diffusion Probabilistic Models (DDPM) [9] have been extensively exploited [1, 24, 27, 28, 30, 34]. Relying on the assumption that abnormal regions would be poorly reconstructed, the residue between the input and the generative model’s output is used to detect abnormalities. Feature modelling in the embedding space is another approach. [5, 23] employ the teacher-student architecture to extract features for normal and abnormal sample discrimination. [4, 22] propose to detect the anomalies by out-of-distribution feature embedding. Despite their success in industrial anomaly detection [2], their effectiveness may be limited when applied to medical domain.

Alternative, anomaly synthesis has emerged as a prominent approach that incorporates pseudo-positive samples to enhance anomaly segmentation. By overlaying color, texture, and semantic outliers on normal samples, a model is trained to segment the synthetic anomalous regions [6, 10, 11, 16, 25, 26, 31, 33]. Despite yielding promising results, there exists significant variation in methods for generating pseudo anomalies. For instance, [15, 25, 26] generate anomalies by utilizing in-distribution image patches, while [10, 16, 31, 33] focus on producing lesions that closely resemble real anomalies. Additionally, prior arts usually focus on model design [16, 32]. However, there is little study explicitly tackling the following two fundamental questions behind this paradigm.

- *Should pseudo-anomalies approximate the queries in the test phase?*
- *How should the segmentation model be trained on the synthesis data?*

In this study, we explore these questions and introduce two principles for pseudo-supervised anomaly segmentation with synthetic anomalies. We apply these principles to unsupervised liver tumor segmentation through adapted discriminative joint reconstruction anomaly embedding (DRAEM) [32]. Our approach introduces a varied anomaly synthesis pipeline and a balanced two-stage training strategy for DRAEM, resulting in outstanding performance on the liver tumor segmentation dataset (LiTs) [3].

2 Preliminaries

This section tackles two fundamental, yet under-explored questions in pseudo-supervised anomaly segmentation with synthetic abnormalities. The reasoning offers insights for designing the proposed solution.

Q1: About pseudo anomaly generation: Should pseudo-anomalies approximate the common queries in the test set?

Pseudo anomaly is introduced to establish the boundary that distinguishes abnormality, transforming the unsupervised problem into pseudo-supervision, which helps the model learn normal patterns by providing negative samples. Since there is no clear definition of what constitutes an anomaly, there shouldn't be any bound or limit on pseudo anomaly synthesis. Instead of focusing on creating pseudo anomalies that match known abnormal patterns in queries, we advocate generating a diversity of anomalies to facilitate a model to learn the comprehensive normal spectrum. In particular, when dealing with unsupervised tumor segmentation, we believe that generating a large diversity of pseudo anomalies in terms of intensity, shape, and textures facilitates addressing the high heterogeneity in tumors. This motivates the design of the proposed pseudo anomaly generation module.

Q2: About model training: Should the model training follow the exact supervised training principles on synthetic anomalies?

The success of supervised learning relies on the IID assumption that both the training and test data follow an identical distribution. Under this assumption, a model is usually well-trained on the training set with multiple iterations. However, we argue that one shouldn't follow the same philosophy to train a model on pseudo anomalies in anomaly detection and segmentation. According to the reasoning in **Q1**, a covariate shift is likely to exist between the synthesized and query anomalies. We visualize this covariate shift by 2-D TSNE in Fig. 1(C), where both tumor samples and normal images are from the LiTs dataset [3]. Consequently, due to the potential covariate shift between the synthesized and the common anomalies in pseudo-supervised segmentation, training a model on the pseudo anomalies may cause a bias and harm its performance on real queries. In another words, a good-fit model on the pseudo-anomaly data may fail on real testing data. Our ablation experiment shown in Fig. 3 validates this hypothesis. Therefore, unlike conventional supervised learning that requires a relatively long training time, we argue that model optimization on anomaly synthesis for pseudo-supervised segmentation should stop early to preserve the model's generalizability on queries. Our answer to **Q2** inspires us to design the two-phase training strategy in this study.

3 Method

Toward unsupervised liver tumor segmentation, we incorporate our reasoning to **Q1** and **Q2** into the DRAEM-similar [32] architecture. As depicted in

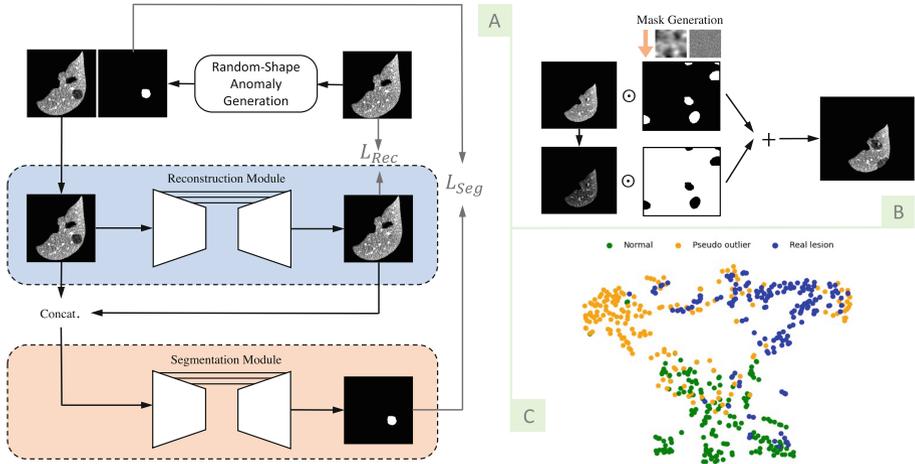


Fig. 1. (A) Systematic diagram of the proposed unsupervised liver tumor segmentation scheme. During training, synthetic abnormalities are fed to a restoration net followed by a segmentation net. To avoid model overfitting on synthesis, the two models are trained in two phases represented by blue and orange, respectively. In inference, a query is directly passed to the two networks for segmentation. (B) Proposed synthesis pipeline based on Gaussian noise stretching. (C) Liver image embedding by 2-D TSNE. (Color figure online)

Fig. 1(A), the framework comprises random-shape anomaly generation, a restoration network, and a segmentation network. Unlike DRAEM training both networks jointly, we propose a two-phase learning to avoid segmentation model over-fitting on synthetic abnormalities. In inference, only the reconstructive network and segmentation network are deployed on queries. Compared to DRAEM, our experiments show that both the proposed anomaly generation module and the two-phase learning strategy boost the liver tumor segmentation performance in terms of segmentation accuracy and model stability.

3.1 Pseudo Anomaly Generation

The anomalous training samples are simulated by the anomaly synthesis module, which generates masks of random shapes and sizes through Gaussian noise and morphological transformations. Initially, Gaussian noise is generated with the same resolution as a normal image and then blurred with a Gaussian kernel. The noise is then stretched and thresholded to produce a binarized mask. Subsequently, closing and opening operations with the elliptical kernel are applied to the binarized mask to obtain an anomaly segmentation mask. The detailed algorithm is shown in Algorithm 1 in Supplementary.

Using the generated anomaly mask M_s , we proceed to synthesize the abnormal sample I_s . In CT slides, unhealthy patterns in liver regions are demonstrated by abnormal Hounsfield-Unit (HU) values. Therefore, we propose to randomly

shift the intensity of the slice and overlay the new intensity values on the original image I within the mask regions (as shown in Fig. 1(B)). We demonstrate the proposed abnormality synthesis in Fig. 1(B) and formulate it as

$$I_s = (1 - M_s) \odot (I + C) + M_s \odot I, \quad |C| \in (\minRange, \maxRange), \quad (1)$$

where I_s represents synthesized anomalies, \odot is element-wise multiplication, and C is a random value drawn from a Gaussian distribution within a defined range.

It is noteworthy that unlike [10, 33] that aims to fabricate pseudo anomalies to approximate the common patterns of liver tumors, we follow our principle to **Q1**, leverage the stochastic nature in the proposed synthesis process to generate a wide spectrum of anomalies deviating from normal patterns (as shown in Fig. 1(C)). Our experiment shows that our method outperforms [33] by 12% in Dice.

3.2 Model Architecture and Training Functions

The reconstruction network is trained to restore anomalous regions while preserving the normal regions. The segmentation network takes the concatenation of the restoration and pseudo-anomalous image as input and targets to estimate an accurate segmentation map for the anomaly. For the reconstruction network, we use U-Net [21] with 3 encoder and decoder blocks as backbones. The specific encoder block in the restoration network adopts the architecture proposed in [11], where it consists of 2 weight-standardized convolutions [19] followed by swish activation [20] and group normalization [29].

To address diverse levels of model optimization complexity, we train the two networks consequently in two phases. The reconstruction model is first trained to restore the anomalous region in synthetic abnormal images with L_1 loss:

$$L_{rec}(I_s, \tilde{I}_s) = |I_s - \tilde{I}_s|, \quad (2)$$

where I_s, \tilde{I}_s are the pseudo outlier augmented sample and the reconstruction image. After freezing the well-trained generative module, we slightly train the segmentation model to avoid bias introduced by the covariance shift. To accommodate potential small tumors, Focal Loss [17] is adopted:

$$L_{seg}(M_s, \tilde{M}_s) = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^C \alpha_j (1 - \tilde{m}_{s,ij})^\gamma \log(\tilde{m}_{s,ij}) \quad (3)$$

where $\tilde{m}_{s,ij}$ is the predicted probability of class j at pixel i and α_j is the weight for class j , and M_s, \tilde{M}_s are the ground truth and the estimated anomaly masks.

4 Experiments

4.1 Experimental Setting

Dataset Preparation: We evaluate the proposed method on the Liver Tumor Segmentation (LiTs) dataset [3] from MICCAI 2017 challenge. LiTs dataset

consists of 131 abdominal computed tomography (CT) scans with the paired liver and liver tumour ground truth labels. Unlike with previous works [7, 16], which perform the cross-fold validation on the LiTs dataset, we argue that training on the retrieved partial samples from an unhealthy CT scan is not ideal for the model to learn the complete liver feature distribution. Therefore, we train our model on an anomaly-free dataset BTCV [14], which provides 40 healthy CT abdomen scans and the corresponding organ masks.

Table 1. Liver tumor segmentation on LiTs [3]. Our method exhibits the best Dice with a standard deviation 1.78. Results with * are directly copied from original papers.

Methods	Supervision	Dice
Zhang et al. [33]	✓	61.91*
DRAEM [32]	X	14.75
Zhang et al. [33]	X	40.78*
ASC-Net [7]	X	32.24*
ASC-Net + postprocessing [7]	X	50.23*
Hu et al. [10]	X	59.77*
Ours	X	53.03

For all CT volumes in training and test, HU values are transformed into grayscale and the liver Region of Interest (ROI) is extracted according to the organ annotations. Then 2D slices are obtained along the Axial plane, resized to 256×256 , and normalized independently by histogram equalization.

Implementation Details: We run the experiments on dual Nvidia RTX-3090 GPUs. The threshold for pseudo mask generation is set to be 200, and the intensity range of the random intensity shift is $[-100, 100]$. The focal loss parameters are defined as $\alpha = 1$ and $\gamma = 2$. We use PyTorch [18] to implement the proposed method. The model is trained for 200 epochs for the first stage and just 1 epoch for the second stage to avoid bias introduced by pseudo anomalies. The learning rate is set to 0.0001, with a batch size of 8 using Adam [12] optimizer. We follow previous studies and use the Dice score as our evaluation metric.

4.2 Results and Discussion

Comparison to SOTA: We quantitatively compare the proposed method with state-of-the-art unsupervised liver tumor segmentation methods including Zhang et al. [33], Hu et al. [10] ASC-Net [7] both with and without manually-designed post-processing and report the results in Table 1. The fully supervised method is taken as performance upper bound. As shown in Table 1, our approach significantly outperforms the other methods, with the exception of [10] and shrinks the gap between unsupervised method and fully-supervision. Notably, [10] leverages

extensive clinical prior knowledge to synthesize pseudo anomalies resembling real tumors. Furthermore, our approach achieves a substantial reduction in runtime at 0.018 s/slice , compared to 0.476 s/slice in [10] which operates on 3D volume, incurring higher memory usage and slower inference time. In Fig. 2, we show our segmentation results on real tumor data in the LiTs dataset.

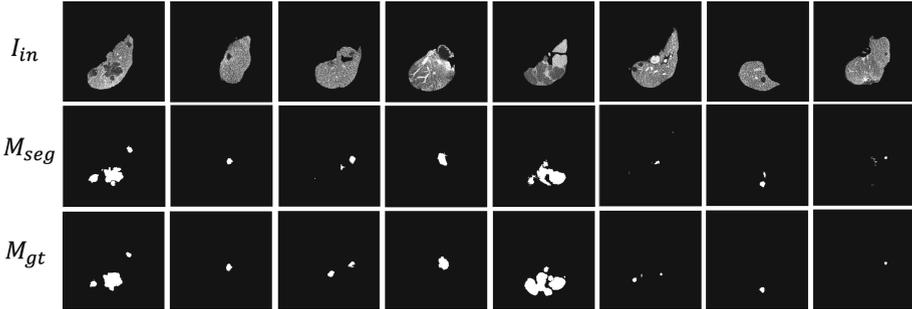


Fig. 2. Tumor segmentation on real liver tumor data, from easy (left) to difficult (right). I_{in} : Input, M_{seg} : segmentation mask, and M_{gt} : Ground-Truth.

Table 2. Ablation study of two-phase training (TP), pseudo anomaly (PA), and reconstructive network. The baseline is DRAEM model [32]. Asterisks indicate statistical significance (*: $p \leq 0.05$, **: $p \leq 0.001$) when using a paired Student’s t -test compared to the baseline.

Method	+TP	+PA	+U-Net	Dice
Baseline				14.75 ± 14.28
Baseline	✓			21.31 ± 12.54
Baseline	✓	✓		$30.17 \pm 5.50^*$
Baseline		✓	✓	$40.06 \pm 6.85^*$
Baseline	✓	✓	✓	$53.03 \pm 1.78^{**}$

Ablation on Model Components: The proposed method and DRAEM differ in three aspects: pseudo anomaly generation (corresponding to **Q1**), two-phase training (corresponding to **Q2**), and U-Net backbone in the restoration net. In this ablation study, we take the DRAEM as baseline, decouple these factors, and evaluate their impact in terms of tumor detection (by AUROC) and segmentation (by DICE) on LiTs. We run this ablation 3 times, and the performance is reported in Table 2. Due to page limitation, additional results such as AUROC, anomaly size mask threshold selection, and a wide spectrum of pseudo anomalous samples can be found in supplementary.

The two-phase training strategy improves Dice by 6.5% on the baseline. When combined with U-Net and PA, there’s a significant 13% performance

boost compared to using only U-Net and PA, validating our hypothesis that light segmentation training on pseudo anomalies helps address the covariant shift between synthetic anomalies and real tumors. We further extended the training of the segmentation net to 200 epochs and captured tumor detection performance (by AUROC) and segmentation quality (by Dice) every 5 epochs. It’s worth emphasizing that these experiments incorporate the synergistic application of TA, U-Net and PA, as this combination has proven to demonstrate optimal outcomes with TA. Therefore, Fig. 3 results diverge from the Dice score in Table 2, where TA solely influences the baseline, yielding a comparatively less pronounced impact on reducing training perturbation. As shown in Fig. 3, the mean AUROC keeps decreasing, and the standard deviation keeps increasing. The perturbation also occurs in Dice after 50 epochs. We attribute this to model overfitting on the pseudo data, which hurts model’s generalizability on queries.

Additionally, the proposed anomaly synthesis module and U-Net designed in our restoration net significantly boost the segmentation performance. Figure 4 presents a visualization comparison of reconstructions generated by autoencoder (AE) and U-Net. Compared to the AE-based network, the skip-connection in U-Net helps preserve the texture details in liver reconstruction images, which facilitates the downstream segmentation task.

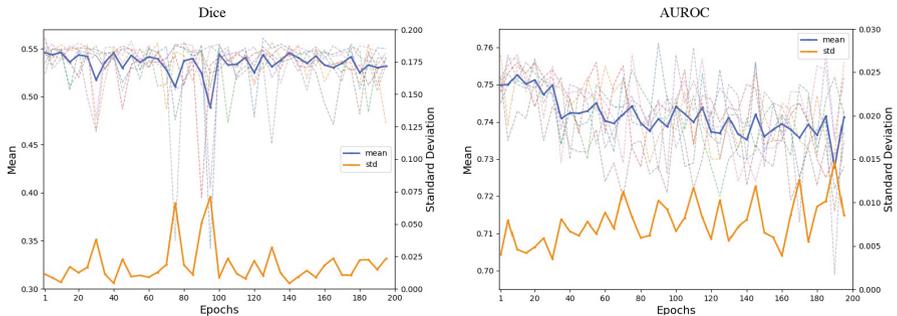


Fig. 3. An illustrative depiction of the evaluation performance of the segmentation network reveals a tendency to overfit shortly after a short period of training. Throughout this process, the reconstruction network maintains a frozen state.

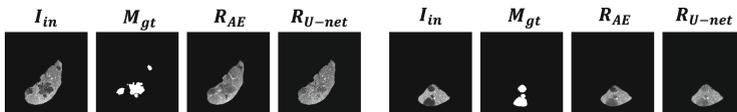


Fig. 4. Visualization of image reconstruction by AE and U-Net.

5 Conclusion

In this study, we tackled the challenging problem of unsupervised liver tumor segmentation and proposed a two-stage pseudo-supervision solution with synthetic anomalies. By generating anomalies spreading over a large spectrum, the synthesis data facilitated the model in finding normal sample boundary in embedding space. The two-stage training strategy mitigated the impact of covariant shift between synthesis data and actual tumor data on model optimization, and thus avoid segmentation model's overfitting on synthetic anomalies. Experimentation suggested that the proposed method performs comparably to SOTA methods. Looking ahead, we aspire to extend our exploration of model performance to encompass various other diseases and data modalities and investigate the integration of both real and synthetic tumor within the model training pipeline.

References

1. Akçay, S., Atapour-Abarghouei, A., Breckon, T.P.: Skip-ganomaly: skip connected and adversarially trained encoder-decoder anomaly detection. In: 2019 International Joint Conference on Neural Networks (IJCNN), pp. 1–8 (2019)
2. 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, pp. 9592–9600 (2019)
3. Bilic, P., et al.: The liver tumor segmentation benchmark (LITS). [arXiv:1901.04056](https://arxiv.org/abs/1901.04056) (2019)
4. Defard, T., Setkov, A., Loesch, A., Audigier, R.: PaDiM: a patch distribution modeling framework for anomaly detection and localization. In: Del Bimbo, A., et al. (eds.) ICPR 2021. LNCS, vol. 12664, pp. 475–489. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-68799-1_35
5. Deng, H., Li, X.: Anomaly detection via reverse distillation from one-class embedding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 9737–9746 (2022)
6. Deng, H., Li, X.: Self-supervised anomaly detection with random-shape pseudo-outliers. In: 2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), pp. 4768–4772 (2022)
7. Dey, R., Hong, Y.: ASC-Net: adversarial-based selective network for unsupervised anomaly segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 236–247 (2021)
8. Goodfellow, I.: Generative adversarial networks. *Commun. ACM* **63**(11), 139–144 (2020)
9. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: Advances in Neural Information Processing Systems, vol. 33, pp. 6840–6851 (2020)
10. Hu, Q., et al.: Label-free liver tumor segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 7422–7432 (2023)
11. Kascenas, A., Pugeault, N., O’Neil, A.Q.: Denoising autoencoders for unsupervised anomaly detection in brain MRI. In: Medical Imaging with Deep Learning (2021)

12. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
13. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. [arXiv:1312.6114](https://arxiv.org/abs/1312.6114) (2013)
14. Landman, B., Xu, Z., Igelsias, J., Styner, M., Langerak, T., Klein, A.: MICCAI multi-atlas labeling beyond the cranial vault-workshop and challenge. In: Proceedings of the MICCAI Multi-Atlas Labeling Beyond Cranial Vault-Workshop Challenge, vol. 5, p. 12 (2015)
15. Li, C.L., Sohn, K., Yoon, J., Pfister, T.: CutPaste: self-supervised learning for anomaly detection and localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 9664–9674 (2021)
16. Li, H., Iwamoto, Y., Han, X., Lin, L., Hu, H., Chen, Y.W.: An accurate unsupervised liver lesion detection method using pseudo-lesions. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) MICCAI 2022. LNCS, vol. 13438, pp. 214–223. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-16452-1_21
17. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2980–2988 (2017)
18. Paszke, A., et al.: PyTorch: an imperative style, high-performance deep learning library. In: Advances in Neural Information Processing Systems, vol. 32, pp. 8024–8035 (2019)
19. Qiao, S., Wang, H., Liu, C., Shen, W., Yuille, A.: Micro-batch training with batch-channel normalization and weight standardization. [arXiv:1903.10520](https://arxiv.org/abs/1903.10520) (2019)
20. Ramachandran, P., Zoph, B., Le, Q.V.: Searching for activation functions. [arXiv:1710.05941](https://arxiv.org/abs/1710.05941) (2017)
21. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
22. Roth, K., Pemula, L., Zepeda, J., Schölkopf, B., Brox, T., Gehler, P.: Towards total recall in industrial anomaly detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 14318–14328 (2022)
23. Salehi, M., Sadjadi, N., Baselizadeh, S., Rohban, M.H., Rabiee, H.R.: Multiresolution knowledge distillation for anomaly detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 14902–14912 (2021)
24. Schlegl, T., Seeböck, P., Waldstein, S.M., Schmidt-Erfurth, U., Langs, G.: Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In: International Conference on Information Processing in Medical Imaging, pp. 146–157 (2017)
25. Tan, J., Hou, B., Batten, J., Qiu, H., Kainz, B.: Detecting outliers with foreign patch interpolation. [arXiv:2011.04197](https://arxiv.org/abs/2011.04197) (2020)
26. Tan, J., Hou, B., Day, T., Simpson, J., Rueckert, D., Kainz, B.: Detecting outliers with poisson image interpolation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 581–591 (2021)
27. Wang, M., et al.: Unsupervised anomaly detection with local-sensitive VQVAE and global-sensitive transformers. [arXiv:2303.17505](https://arxiv.org/abs/2303.17505) (2023)
28. Wolleb, J., Bieder, F., Sandkühler, R., Cattin, P.C.: Diffusion models for medical anomaly detection. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 35–45 (2022)

29. Wu, Y., He, K.: Group normalization. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11217, pp. 3–19. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01261-8_1
30. Wyatt, J., Leach, A., Schmon, S.M., Willcocks, C.G.: AnoDDPM: anomaly detection with denoising diffusion probabilistic models using simplex noise. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 650–656 (2022)
31. Yao, Q., Xiao, L., Liu, P., Zhou, S.K.: Label-free segmentation of COVID-19 lesions in lung CT. *IEEE Trans. Med. Imaging* **40**(10), 2808–2819 (2021)
32. 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, p. 8330–8339 (2021)
33. Zhang, X., Xie, W., Huang, C., Zhang, Y., Wang, Y.: Self-supervised tumor segmentation through layer decomposition. [arXiv:2109.03230](https://arxiv.org/abs/2109.03230) (2021)
34. Zimmerer, D., Isensee, F., Petersen, J., Kohl, S., Maier-Hein, K.: Unsupervised anomaly localization using variational auto-encoders. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11767, pp. 289–297. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32251-9_32