

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

10-2023

Unsupervised anomaly detection in medical images with a memory-augmented multi-level cross-attentional masked autoencoder

Yu TIAN

Guansong PANG

Singapore Management University, gspang@smu.edu.sg

Yuyuan LIU

Chong WANG

Yuanhong CHEN

~~See next page for additional authors~~

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Artificial Intelligence and Robotics Commons](#), and the [Graphics and Human Computer Interfaces Commons](#)

Citation

TIAN, Yu; PANG, Guansong; LIU, Yuyuan; WANG, Chong; CHEN, Yuanhong; LIU, Fengbei; SINGH, Rajvinder; VERJANS, Johan W.; WANG, Mengyu; and CARNEIRO, Gustavo. Unsupervised anomaly detection in medical images with a memory-augmented multi-level cross-attentional masked autoencoder. (2023). *Machine Learning in Medical Imaging: 14th International Workshop, MLMI 2023, Vancouver, October 8: Proceedings*. 14349, 11-21.

Available at: https://ink.library.smu.edu.sg/sis_research/8329

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylids@smu.edu.sg.

Author

Yu TIAN, Guansong PANG, Yuyuan LIU, Chong WANG, Yuanhong CHEN, Fengbei LIU, Rajvinder SINGH, Johan W. VERJANS, Mengyu WANG, and Gustavo CARNEIRO

Unsupervised Anomaly Detection in Medical Images with a Memory-augmented Multi-level Cross-attentional Masked Autoencoder

Yu Tian¹, Guansong Pang⁵, Yuyuan Liu², Chong Wang², Yuanhong Chen², Fengbei Liu², Rajvinder Singh³, Johan W Verjans^{2,3,4}, Mengyu Wang¹, and Gustavo Carneiro⁶

¹ Harvard Ophthalmology AI Lab, Harvard University.

² Australian Institute for Machine Learning, University of Adelaide

³ Faculty of Health and Medical Sciences, University of Adelaide

⁴ South Australian Health and Medical Research Institute

⁵ Singapore Management University

⁶ Centre for Vision, Speech and Signal Processing, University of Surrey

Abstract. Unsupervised anomaly detection (UAD) aims to find anomalous images by optimising a detector using a training set that contains only normal images. UAD approaches can be based on reconstruction methods, self-supervised approaches, and Imagenet pre-trained models. Reconstruction methods, which detect anomalies from image reconstruction errors, are advantageous because they do not rely on the design of problem-specific pretext tasks needed by self-supervised approaches, and on the unreliable translation of models pre-trained from non-medical datasets. However, reconstruction methods may fail because they can have low reconstruction errors even for anomalous images. In this paper, we introduce a new reconstruction-based UAD approach that addresses this low-reconstruction error issue for anomalous images. Our UAD approach, the memory-augmented multi-level cross-attentional masked autoencoder (MemMC-MAE), is a transformer-based approach, consisting of a novel memory-augmented self-attention operator for the encoder and a new multi-level cross-attention operator for the decoder. MemMC-MAE masks large parts of the input image during its reconstruction, reducing the risk that it will produce low reconstruction errors because anomalies are likely to be masked and cannot be reconstructed. However, when the anomaly is not masked, then the normal patterns stored in the encoder’s memory combined with the decoder’s multi-level cross-attention will constrain the accurate reconstruction of the anomaly. We show that our method achieves SOTA anomaly detection and localisation on colonoscopy, pneumonia, and covid-19 chest x-ray datasets.

Keywords: Pneumonia · Covid-19 · Colonoscopy · Unsupervised Learning · Anomaly Detection · Anomaly Segmentation · Vision Transformer

1 Introduction

Detecting and localising anomalous findings in medical images (e.g., polyps, malignant tissues, etc.) are of vital importance [1, 4, 7, 12–15, 17–19, 27, 29, 30, 32, 34].

Systems that can tackle these tasks are often formulated with a classifier trained with large-scale datasets annotated by experts. Obtaining such annotation is often challenging in real-world clinical datasets because the amount of normal images from healthy patients tend to overwhelm the amount of anomalous images. Hence, to alleviate the challenges of collecting anomalous images and learning from class-imbalanced training sets, the field has developed unsupervised anomaly detection (UAD) models [3, 31] that are trained exclusively with normal images. Such UAD strategy benefits from the straightforward acquisition of training sets containing only normal images and the potential generalisability to unseen anomalies without collecting all possible anomalous sub-classes.

Current UAD methods learn a one-class classifier (OCC) using only normal/healthy training data, and detect anomalous/disease samples using the learned OCC [3, 8, 11, 16, 22, 25, 26, 33, 36]. UAD methods can be divided into: 1) reconstruction methods, 2) self-supervised approaches, and 3) Imagenet pre-trained models. Reconstruction methods [3, 8, 16, 25, 36] are trained to accurately reconstruct normal images, exploring the assumption that the lack of anomalous images in the training set will prevent a low error reconstruction of an test image that contains an anomaly. However, this assumption is not met in general because reconstruction methods are indeed able to successfully reconstruct anomalous images, particularly when the anomaly is subtle. Self-supervised approaches [28, 31, 34] train models using contrastive learning, where pretext tasks must be designed to emulate normal and anomalous image changes for each new anomaly detection problem. Imagenet pre-trained models [5, 24] produce features to be used by OCC, but the translation of these models into medical image problems is not straightforward. Reconstruction methods are able to circumvent the aforementioned challenges posed by self-supervised and Imagenet pre-trained UAD methods, and they can be trained with a relatively small amount of normal samples. However, their viability depends on an acceptable mitigation of the potentially low reconstruction error of anomalous test images.

In this paper, we introduce a new UAD reconstruction method, the Memory-augmented Multi-level Cross-attention Masked Autoencoder (MemMC-MAE), designed to address the low reconstruction error of anomalous test images. MemMC-MAE is a transformer-based approach based on masked autoencoder (MAE) [9] with of a novel memory-augmented self-attention encoder and a new multi-level cross-attention decoder. MemMC-MAE masks large parts of the input image during its reconstruction, and given that the likelihood of masking out an anomalous region is large, then it is unlikely that it will accurately reconstruct that anomalous region. However, there is still the risk that the anomaly is not masked out, so in this case, the normal patterns stored in the encoder’s memory combined with the correlation of multiple normal patterns in the image, utilised by the decoder’s multi-level cross-attention can explicitly constrain the accurate anomaly reconstruction to produce high reconstruction error (high anomaly score). The encoder’s memory is also designed to address the MAE’s long-range ‘forgetting’ issue [20], which can be harmful for UAD due to the poor reconstruction based on forgotten normality patterns and ‘unwanted’ generalisability to subtle anomalies during testing. Our contributions are summarised as:




Fig. 1: **Top:** overall MemMC-MAE framework. Yellow tokens indicate the unmasked visible patches, and blue tokens indicate the masked patches. Our memory-augmented transformer encoder only accepts the visible patches/tokens as input, and its output tokens are combined with dummy masked patches/tokens for the missing pixel reconstruction using our proposed multi-level cross-attentional transformer decoder. **Bottom-left:** proposed memory-augmented self-attention operator for the transformer encoder, and **bottom-right:** proposed multi-level cross-attention operator for the transformer decoder.

- To the best of our knowledge, this is the first memory-based UAD method that relies on MAE [9];
- A new memory-augmented self-attention operator for our MAE transformer encoder to explicitly encode and memorise the normality patterns; and
- A novel decoder architecture that uses the learned multi-level memory-augmented encoder information as prior features to a cross-attention operator.

Our method achieves better anomaly detection and localisation accuracy than most competing approaches on the UAD benchmarks using the public Hyper-Kvasir colonoscopy dataset [2], pneumonia [10] and Covid-X [37] Chest X-ray (CXR) dataset.

2 Method

2.1 Memory-augmented Multi-level Cross-attentional Masked Autoencoder (MemMC-MAE)

Our MemMC-MAE, depicted in Fig. 1, is based on the masked autoencoder (MAE) [9] that was recently developed for the pre-training of models to be used in downstream computer vision tasks. MAE has an asymmetric architecture, with a encoder that takes a small subset of the input image patches and a

smaller/lighter decoder that reconstructs the original image based on the input tokens from visible patches and dummy tokens from masked patches.

Our MemMC-MAE is trained with a normal image training set, denoted by $\mathcal{D} = \{\mathbf{x}_i\}_{i=1}^{|\mathcal{D}|}$, where $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^{H \times W \times R}$ (H : height, W : width, R : number of colour channels). Our method first divides the input image \mathbf{x} into non-overlapping patches $\mathcal{P} = \{\mathbf{p}_i\}_{i=1}^{|\mathcal{P}|}$, where $\mathbf{p} \in \mathbb{R}^{\hat{H} \times \hat{W} \times R}$, with $\hat{H} \ll H$ and $\hat{W} \ll W$. We then randomly mask out 75% of the $|\mathcal{P}|$ patches, and the remaining visible patches $\mathcal{P}^{(v)} = \{\mathbf{p}_v\}_{v=1}^{|\mathcal{P}^{(v)}|}$ (with $|\mathcal{P}^{(v)}| = 0.25 \times |\mathcal{P}|$) are used by the MemMC-MAE to encode the normality patterns of those patches, and all $|\mathcal{P}^{(v)}|$ encoded visible patches and $|\mathcal{P}| - |\mathcal{P}^{(v)}|$ dummy masked patches are used as the input of a new multi-level cross-attention decoder to reconstruct the image.

The training of MemMC-MAE is based on the minimisation of the mean squared error (MSE) loss between the input and reconstructed images at the pixels of the masked patches of the training images. The approach is evaluated on a testing set $\mathcal{T} = \{(\mathbf{x}, y, \mathbf{m})_i\}_{i=1}^{|\mathcal{T}|}$, where $y \in \mathcal{Y} = \{\text{normal}, \text{anomalous}\}$, and $\mathbf{m} \in \mathcal{M} \subset \{0, 1\}^{H \times W \times 1}$ denotes the segmentation mask of the lesion in the image \mathbf{x} . When testing, we also mask 75% of the image and the patch-wise reconstruction error indicates anomaly localisation, and the mean reconstruction error of all patches is used to detect image-wise anomaly. Below we provide details on the major contributions of MemMC-MAE, which are the memory-augmented transformer encoder that stores the long-term normality patterns of the training samples, and the new multi-level cross-attentional transformer decoder to leverage the correlation of features from the encoder to reconstruct the missing normal pixels.

Memory-augmented Transformer Encoder (Fig. 1 - bottom left) We modify the encoder from the transformer with our a novel memory-augmented self-attention, by extending the keys and values of the self-attention operation with learnable memory matrices that store normality patterns, which are updated via back-propagation. To this end, the proposed self-attention (SA) module for layer $l \in \{0, \dots, L - 1\}$ is defined as:

$$\mathbf{X}^{(l+1)} = f_{SA}(\mathbf{W}_Q^{(l)} \mathbf{X}^{(l)}, [\mathbf{W}_K^{(l)} \mathbf{X}^{(l)}, \mathbf{M}_K^{(l)}], [\mathbf{W}_V^{(l)} \mathbf{X}^{(l)}, \mathbf{M}_V^{(l)}]), \quad (1)$$

where $\mathbf{X}^{(0)}$ is the encoder input matrix containing $|\mathcal{P}^{(v)}|$ patch tokens formed from the visible image patches transformed through the linear projection $\mathbf{W}^{(0)}$, with $|\mathcal{P}^{(v)}|$ being the number of visible tokens/patches, $\mathbf{X}^{(l)}$, $\mathbf{X}^{(l+1)}$ are the input and output of layer l , $\mathbf{W}_Q^{(l)}$, $\mathbf{W}_K^{(l)}$, $\mathbf{W}_V^{(l)}$ are the linear projections of the encoder’s layer l for query, key and value of the self-attention operator, respectively, and $\mathbf{M}_K^{(l)}$, $\mathbf{M}_V^{(l)}$ are the layer l learnable memory matrices that are concatenated with $\mathbf{W}_K \mathbf{X}^{(l)}$ and $\mathbf{W}_V \mathbf{X}^{(l)}$ using the operator $[\cdot, \cdot]$. The self-attention operator $f_{SA}(\cdot)$ follows the standard ViT [6] and transformer [35], which computes a weighted sum of value vectors according to the cosine similarity distribution between query and key. Such memory-augmented self-attention aims to store normal patterns that are not encoded in the feature $\mathbf{X}^{(l)}$, forcing the decoder to reconstruct anomalous input patches into normal output patches during testing.

Multi-level Cross-Attention Transformer Decoder (Fig. 1 - bottom right). Our transformer decoder computes the cross-attention operation using the outputs from all encoder layers and the decoder layer output from the self-attention operator (see Fig. 1 - Bottom right). More formally, the layer $d \in \{0, \dots, D - 1\}$ of our decoder outputs

$$\mathbf{Y}^{(d+1)} = \sum_{l=1}^L \alpha^{(d,l)} \times f_{SA}(f_{SA}(\mathbf{Y}^{(d)}, \mathbf{Y}^{(d)}, \mathbf{Y}^{(d)}), \mathbf{W}_K^{(d)} \mathbf{X}^{(l)}, \mathbf{W}_V^{(d)} \mathbf{X}^{(l)}), \quad (2)$$

where $\mathbf{Y}^{(d)}$ and $\mathbf{Y}^{(d+1)}$ represent the input and output of the decoder layer d containing $|\mathcal{P}|$ tokens (i.e., $|\mathcal{P}^{(v)}|$ tokens from the visible patches of the encoder and $|\mathcal{P}| - |\mathcal{P}^{(v)}|$ dummy tokens from the masked patches), $\mathbf{X}^{(l)}$ denotes the output from encoder layer $l-1$, and $\mathbf{W}_K^{(d)}, \mathbf{W}_V^{(d)}$ are the linear projections of the layer d of the decoder for the key and value of the self-attention operator, respectively. Note that all $|\mathcal{P}|$ input tokens for the decoder are attached with positional embeddings. The multi-level cross-attention results in (2) are fused together with a weighted sum operation using the weight $\alpha^{(d,l)}$, which is computed based on a linear projection layer and sigmoid function to control the weight of different layers' cross-attention results, as in

$$\alpha^{(d,l)} = \sigma \left(\mathbf{W}_\alpha^{(d,l)} \left(\left[f_{SA}(\mathbf{Y}^{(d)}, \mathbf{Y}^{(d)}, \mathbf{Y}^{(d)}), \mathbf{Y}^{(d+1)} \right] \right) \right), \quad (3)$$

where $\sigma(\cdot)$ is the sigmoid function, and $\mathbf{W}_\alpha^{(d,l)}$ denotes a learnable weight matrix. Such fusion mechanism enforces the correlation of multiple normal patterns in the image present at different levels of encoding information to contribute at different decoding layers by adjusting their relative importance using the self-attention output from $f_{SA}(\cdot)$ and cross-attention output $\mathbf{Y}^{(d+1)}$.

2.2 Anomaly Detection and Segmentation

We compute the anomaly score [3] with multi-scale structural similarity (MS-SSIM) [38]. The anomaly scores are pooled from 10 different random seeds for masking image patches with a fixed 75% masking ratio, which enables a more robust anomaly detection and localisation. The anomaly localisation mask is obtained by computing the mean MS-SSIM scores for all patches, and the anomaly detection relies on the mean MS-SSIM scores from the patches [3].

3 Experiments and Results

Datasets and Evaluation Measures: Three disease screening datasets are used in our experiments. We test anomaly detection on the CXR images of the pneumonia chest X-ray dataset [10] and Covid-X dataset [37], and both anomaly detection and localisation on the colonoscopy images of the Hyper-Kvasir dataset [2]. The publicly available **pneumonia chest X-ray dataset** [10], consisting of normal and pneumonia-affected images, was obtained from a total

Methods	Publication	Pneumonia	Covid-X	Hyper-Kvasir
DAE [21]	ICANN'11	0.599	0.557	0.705
OCGAN [23]	CVPR'18	0.703	0.612	0.813
F-anoGAN [25]	IPMI'17	0.755	0.669	0.907
ADGAN [16]	ISBI'19	0.627	0.659	0.913
MS-SSIM [3]	AAAI'22	0.695	0.634	0.917
PANDA [24]	CVPR'21	0.657	0.629	0.937
PaDiM [5]	ICPR'21	0.663	0.614	0.923
IGD [3]	AAAI'22	0.734	0.699	0.939
CCD+IGD* [31]	MICCAI'21	0.775	0.746	0.972
Ours		0.879	0.917	0.972

Table 1: **Anomaly detection AUC** test results on Pneumonia and Covid-X Chest X-ray datasets and Hyper-Kvasir colonoscopy dataset. CCD+IGD* [31] requires at least 2×longer training time than other approaches in the table because of a two-stage self-supervised pre-training and fine-tuning.

of 6,480 patients. In accordance with [39], we structured the anomaly detection dataset such that the training set encompasses 1,349 normal images, and the testing set comprises 234 normal and 390 pneumonia images. Each chest X-ray image has been resized to the standardized dimensions of 256×256 pixels. **Covid-X** [37] has a training set with 1,670 Covid-19 positive and 13,794 Covid-19 negative CXR images, but we only use the 13,794 Covid-19 negative CXR images for training. The test set contains 400 CXR images, consisting of 200 positive and 200 negative images, each image with size 299×299 pixels. **Hyper-Kvasir** is a large-scale public gastrointestinal dataset. The images were collected from the gastroscopy and colonoscopy procedures from Baerum Hospital in Norway, and were annotated by experienced medical practitioners. The dataset contains 110,079 images from unhealthy and healthy patients, out of which, 10,662 are labelled. Following [31], 2,100 normal images are selected, from which we use 1,600 for training and 500 for testing. The testing set also contains 1,000 anomalous images with their segmentation masks. Detection is assessed with area under the ROC curve (AUC), and localisation is evaluated with intersection over union (IoU).

Implementation Details For the transformer, we follow ViT-B [6, 9] for designing the encoder and decoder, consisting of stacks of transformer blocks. Inspired by U-Net [40] for medical segmentation, we add residual connections to transfer information from earlier to later blocks for both the encoder and decoder. Each encoder block contains a memory-augmented self-attention block and an MLP block with LayerNorm (LN). Each decoder block contains a multi-level cross-attention block and an MLP block with LayerNorm (LN). We also adopt a linear projection layer after the encoder to match the different width between encoder and decoder [9]. We add positional embeddings (with the sine-cosine version) to both the encoder and decoder input tokens. RandomResizedCrop is used for data augmentation during training. Our method is trained for 2000 epochs in an end-to-end manner using the Adam optimiser with a weight decay of 0.05 and a batch size of 256. The learning rate is set to $1.5e-3$. In the beginning, we warm up the training process for 5 epochs. The method is implemented in PyTorch and runs on an NVIDIA 3090 GPU. The overall training time is around 22 hours, and the mean inference time takes 0.21s per image.




Fig. 2: Segmentation results of our proposed method on Hyper-Kvasir [2], with our predictions (Pred) and ground truth annotations (GT).




Fig. 3: Reconstruction of testing images from Covid-X (Top) and Hyper-Kvasir (Bottom). For each triplet, we show the masked image (left), our MemMC-MAE reconstruction (middle), and the ground-truth (right). Normal testing images are marked with green boxes, and anomalous ones are marked with red boxes.

Evaluation on Anomaly Detection on Chest X-ray and Colonoscopy

We compare our method with nine competing UAD approaches: DAE [21], OCGAN [23], f-anogan [25], ADGAN [16], MS-SSIM autoencoder [3], PANDA [24], PaDiM [5], CCD [31] and IGD [3]. We apply the same experimental setup (i.e., image pre-processing, training strategy, evaluation methods) to these methods above as the one for our approach for fair comparison. The quantitative comparison results for anomaly detection are shown in Table 1 for Pneumonia, Covid-X, and Hyper-Kvasir benchmarks. Our MemMC-MAE achieves the best AUC results on three datasets with 87.9%, 91.7% and 97.2%, respectively. On pneumonia chest x-ray dataset, our model surpasses the previous SOTA approaches by a minimum 10.4% AUC and a maximum 28% AUC. On Covid-X, our result outperforms all competing methods by a large margin with an improvement of 17.1% over the second best approach. For Hyper-Kvasir, our result is on par with the best result in the field produced by CCD+IGD [31], which has a training time $2\times$ longer than our approach.

Evaluation on Anomaly Localisation on Colonoscopy

We compare our anomaly localisation results on Table 3 with four recently proposed UAD baselines: IGD [3], PaDiM [5], CCD [31] and CAVGA- R_u [36]. The results of these methods on Table 3 are from [31]. Following [31], we randomly sample five groups of 100 anomalous images from the test set and compute the mean segmentation IoU. The proposed MemMC-MAE surpasses IGD, PaDiM, CAVGA- R_u and CCD

MAE Mem-Enc MC-Dec	AUC - Covid	AUC - Hyper	Methods	Localisation - IoU
✓	0.799	0.915	IGD [3]	0.276
✓ ✓	0.862	0.956	PaDiM [5]	0.341
✓ ✓ ✓	0.917	0.972	CAVGA- R_u [36]	0.349
			CCD + IGD [31]	0.372
			Ours	0.419

Table 2: **Ablation study** on Covid-X of the encoder’s memory-augmented operator (Mem-Enc) and the decoder’s multi-level cross-attention (MC-Dec).

Table 3: **Anomaly localisation:** Mean IoU test results on Hyper-Kvasir on 5 groups of 100 images.

by a minimum of 4.7% and a maximum of 14.3% IoU, illustrating the effectiveness of our model in localising anomalous tissues.

Visualisation of predicted segmentation. The visualisation of polyp segmentation results of MemMC-MAE on Hyper-Kvasir [2] is shown in Fig. 2. Notice that our model can accurately segment colon polyps of various sizes and shapes.

Visualisation of Reconstructed Images Figure 3 shows the reconstructions produced by MemMC-MAE on Covid-X (Top) and Hyper-Kvasir (Bottom) testing images. Notice that our method can effectively reconstruct the anomalous images with polyps/covid as normal images by automatically removing the polyps or blurring the anomalous regions, leading to larger reconstruction errors for those anomalies. The normal images are accurately reconstructed with smaller reconstruction errors than the anomalous images.

Ablation Study Tab. 2 shows the contribution of each component of our proposed method on Covid-X and Hyper-Kvasir testing set. The baseline MAE [9] achieves 79.9% and 91.5% AUC on the two datasets, respectively. Our method obtains a significant performance gain by adding the memory-augmented self-attention operator to the transformer encoder (Mem-Enc). Adding the proposed multi-level cross-attention operator into the decoder (MC-Dec) further boosts the performance on both datasets.

4 Conclusion

We proposed a new UAD reconstruction method, called MemMC-MAE, for anomaly detection and localisation in medical images, which to the best of our knowledge, is the first memory-based UAD method using MAE. MemMC-MAE introduced a novel memory-augmented self-attention operator for the MAE encoder and a new multi-level cross-attention for the MAE decoder to address the large reconstruction error of anomalous images that plague UAD reconstruction methods. The resulting anomaly detector showed SOTA anomaly detection and localisation accuracy on three public medical datasets. Despite the remarkable performance, the results can potentially improve if we use MemMC-MAE as a pre-training approach for other UAD methods, which we plan to explore in the future.

References

1. Baur, C., et al.: Scale-space autoencoders for unsupervised anomaly segmentation in brain mri. In: MICCAI. pp. 552–561. Springer (2020)
2. Borgli, H., et al.: Hyperkvasir, a comprehensive multi-class image and video dataset for gastrointestinal endoscopy. *Scientific Data* **7**(1), 1–14 (2020)
3. Chen, Y., Tian, Y., Pang, G., Carneiro, G.: Deep one-class classification via interpolated gaussian descriptor. arXiv preprint arXiv:2101.10043 (2021)
4. Chen, Y., et al.: Bomd: Bag of multi-label descriptors for noisy chest x-ray classification. arXiv preprint arXiv:2203.01937 (2022)
5. Defard, T., Setkov, A., Loesch, A., Audigier, R.: Padim: a patch distribution modeling framework for anomaly detection and localization. arXiv preprint arXiv:2011.08785 (2020)
6. Dosovitskiy, A., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
7. Fan, D.P., Ji, G.P., Zhou, T., Chen, G., Fu, H., Shen, J., Shao, L.: Pranet: Parallel reverse attention network for polyp segmentation. In: MICCAI. pp. 263–273. Springer (2020)
8. Gong, D., et al.: Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection. In: ICCV. pp. 1705–1714 (2019)
9. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. arXiv preprint arXiv:2111.06377 (2021)
10. Kermay, D.S., Goldbaum, M., Cai, W., Valentim, C.C., Liang, H., Baxter, S.L., McKeown, A., Yang, G., Wu, X., Yan, F., et al.: Identifying medical diagnoses and treatable diseases by image-based deep learning. *cell* **172**(5), 1122–1131 (2018)
11. Li, C.L., et al.: Cutpaste: Self-supervised learning for anomaly detection and localization. In: CVPR. pp. 9664–9674 (2021)
12. Litjens, G., et al.: A survey on deep learning in medical image analysis. *Medical image analysis* **42**, 60–88 (2017)
13. Liu, F., Tian, Y., Cordeiro, F.R., Belagiannis, V., Reid, I., Carneiro, G.: Noisy label learning for large-scale medical image classification. arXiv preprint arXiv:2103.04053 (2021)
14. Liu, F., Tian, Y., et al.: Self-supervised mean teacher for semi-supervised chest x-ray classification. arXiv preprint arXiv:2103.03629 (2021)
15. Liu, F., et al.: Acpl: Anti-curriculum pseudo-labelling for semi-supervised medical image classification. *CVPR* (2022)
16. Liu, Y., et al.: Photoshopping colonoscopy video frames. In: ISBI. pp. 1–5 (2020)
17. Liu, Y., Tian, Y., Wang, C., Chen, Y., Liu, F., Belagiannis, V., Carneiro, G.: Translation consistent semi-supervised segmentation for 3d medical images. arXiv preprint arXiv:2203.14523 (2022)
18. Luo, Y., et al.: Harvard glaucoma fairness: A retinal nerve disease dataset for fairness learning and fair identity normalization. arXiv preprint arXiv:2306.09264 (2023)
19. LZ, C.T.P., et al.: Computer-aided diagnosis for characterisation of colorectal lesions: a comprehensive software including serrated lesions. *Gastrointestinal Endoscopy* (2020)
20. Martins, P.H., Marinho, Z., Martins, A.F.: infinity-former: Infinite memory transformer. arXiv preprint arXiv:2109.00301 (2021)
21. Masci, J., et al.: Stacked convolutional auto-encoders for hierarchical feature extraction. In: ICANN. pp. 52–59. Springer (2011)

22. Pang, G., Shen, C., van den Hengel, A.: Deep anomaly detection with deviation networks. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. pp. 353–362 (2019)
23. Perera, P., Nallapati, R., Xiang, B.: Ocgan: One-class novelty detection using gans with constrained latent representations. In: CVPR. pp. 2898–2906 (2019)
24. Reiss, T., Cohen, N., Bergman, L., Hoshen, Y.: Panda: Adapting pretrained features for anomaly detection and segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2806–2814 (2021)
25. Schlegl, T., et al.: f-anogan: Fast unsupervised anomaly detection with generative adversarial networks. *Medical image analysis* **54**, 30–44 (2019)
26. Seeböck, P., et al.: Exploiting epistemic uncertainty of anatomy segmentation for anomaly detection in retinal oct. *IEEE transactions on medical imaging* **39**(1), 87–98 (2019)
27. Shi, M., et al.: Artifact-tolerant clustering-guided contrastive embedding learning for ophthalmic images in glaucoma. *IEEE Journal of Biomedical and Health Informatics* (2023)
28. Sohn, K., Li, C.L., Yoon, J., Jin, M., Pfister, T.: Learning and evaluating representations for deep one-class classification. arXiv preprint arXiv:2011.02578 (2020)
29. Tian, Y., Maicas, G., Pu, L.Z.C.T., Singh, R., Verjans, J.W., Carneiro, G.: Few-shot anomaly detection for polyp frames from colonoscopy. In: MICCAI. pp. 274–284. Springer (2020)
30. Tian, Y., Pang, G., Liu, F., Liu, Y., Wang, C., Chen, Y., Verjans, J., Carneiro, G.: Contrastive transformer-based multiple instance learning for weakly supervised polyp frame detection. In: MICCAI. pp. 88–98. Springer (2022)
31. Tian, Y., Pang, G., Liu, F., Shin, S.H., Verjans, J.W., Singh, R., Carneiro, G., et al.: Constrained contrastive distribution learning for unsupervised anomaly detection and localisation in medical images. *MICCAI 2021* (2021)
32. Tian, Y., et al.: One-stage five-class polyp detection and classification. In: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019). pp. 70–73. IEEE (2019)
33. Tian, Y., et al.: Pixel-wise energy-biased abstention learning for anomaly segmentation on complex urban driving scenes. arXiv preprint arXiv:2111.12264 (2021)
34. Tian, Y., et al.: Self-supervised multi-class pre-training for unsupervised anomaly detection and segmentation in medical images. arXiv preprint arXiv:2109.01303 (2021)
35. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
36. Venkataramanan, S., Peng, K.C., Singh, R.V., Mahalanobis, A.: Attention guided anomaly localization in images. In: ECCV. pp. 485–503. Springer (2020)
37. Wang, L., Lin, Z.Q., Wong, A.: Covid-net: A tailored deep convolutional neural network design for detection of covid-19 cases from chest x-ray images. *Scientific Reports* **10**(1), 1–12 (2020)
38. Wang, Z., et al.: Multiscale structural similarity for image quality assessment. In: The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003. vol. 2, pp. 1398–1402. Ieee (2003)
39. Zhao, H., Li, Y., He, N., Ma, K., Fang, L., Li, H., Zheng, Y.: Anomaly detection for medical images using self-supervised and translation-consistent features. *IEEE Transactions on Medical Imaging* **40**(12), 3641–3651 (2021)
40. Zhou, Z., et al.: Unet++: A nested u-net architecture for medical image segmentation. In: Deep learning in medical image analysis and multimodal learning for clinical decision support, pp. 3–11. Springer (2018)