LoRIS - Weakly-supervised Anomaly Detection for Ultrasound Images

Marco Colussi[1], Dragan Ahmetovic[1], Gabriele Civitarese[1], Claudio Bettini[1], Aiman Solyman[1], Roberta Gualtierottixx[2,3], Flora Peyvandi[2,3] and Sergio Mascetti[1]

[1] Università degli Studi di Milano, Department of Computer Science, Via Celoria, 18, 20133, Milan, Italy
[2] Università degli Studi di Milano, Department of Pathophysiology and Transplantation, Via Pace, 9, 20122, Milan, Italy
[3] Fondazione IRCCS Ca' Granda Ospedale Maggiore Policlinico, Angelo Bianchi Bonomi Hemophilia and Thrombosis Center, Via Pace, 9, 20122, Milan, Italy
marco.colussi@unimi.it

Abstract. This paper presents **LoRIS** (Localized Reconstruction-by-Inpainting with Single-mask), a novel weakly-supervised anomaly detection technique designed to identify knee joint recess distension in musculoskeletal ultrasound images, which are noisy and unbalanced (as distended cases are rarer). In this context, supervised techniques require a high number of annotated images of both classes (distended and non-distended). On the other hand, we show that existing unsupervised anomaly detection techniques, which can be trained with images from a single class, are ineffective and often unable to correctly localize the anomaly. To overcome these issues, **LoRIS** is trained with nondistended images only and uses the recess bounding box as location prior to guide the reconstruction. Experimental results show that **LoRIS** outperforms state-of-the-art unsupervised anomaly detection techniques. When compared to a state-of-the-art fully supervised solution, **LoRIS** presents similar performance but has two key advantages: during training it requires images from a single class only, and it also outputs the recess segmentation, without the need for segmentation annotations.

Keywords: Anomaly detection· Weak-supervision· Ultrasound imaging

1 Introduction

For patients with hemophilia, joint bleeding is a common complication. If not treated promptly, it can lead to synovial hyperplasia, osteochondral damage, and hemophilic arthropathy [9]. Ultrasound (US) imaging is emerging as a practical approach for detecting bleeding in the joint recess [22]. In order to support medical practitioners in the diagnosis process, Computer-Aided Diagnosis (CAD) systems based on US images have been extensively researched [10]. Specifically, to support the diagnosis of joint bleeding, prior works have proposed techniques to detect joint recess distention caused by joint bleeding. The proposed solutions

are based on binary classification [27] and multi-task learning (combining classification and detection) [3]. In these works, the task of distinguishing between distended and non-distended recesses is addressed with supervised classification. In addition to classification, segmentation is also of utmost importance in medical imaging, as it facilitates the identification of structures or regions of interest, hence enabling visual guidance for professionals [2]. A major problem of these solutions is the reliance on labeled images, which are scarce, imbalanced between the two classes (distended cases are rarer than non-distended ones), and have a high annotation cost.

In the literature, a common approach to tackle these types of problems is unsupervised anomaly detection [26], in which the model is trained only using normal data samples and is used to identify anomalous samples deviating from the learned distribution. Several solutions have been proposed for unsupervised anomaly detection: reconstruction by inpainting [30], conditional GANs [1], patch-based memory banks [24], synthesizing anomaly samples [16, 29], and normalizing flows [8]. Other works propose specific solutions tackling the problems of medical imaging: normalizing reconstruction error with uncertainty [18] or through patch-interpolation [25]. These techniques provide both an overall anomaly score and, often, a pixel-level anomaly map that can be used for anomaly segmentation. However, as we show in this paper, these techniques are ineffective in the specific domain considered in this paper.

To address the ineffectiveness of unsupervised anomaly detection techniques, we propose a solution inspired by weakly supervised segmentation approaches that have been extensively researched in the segmentation domain, where acquiring the segmentation masks is not always feasible [5]. These approaches rely on *weak* labels that contain partial information compared to the labels used in the supervised approach. In particular, previous works suggest that the use of a *location prior*, in the form of the bounding box of the element of interest, can effectively mitigate the cost of annotation while still providing high accuracy in semantic segmentation [14], referring image segmentation [6], and in medical image segmentation [12, 17]. However, to the best of our knowledge, these approaches have never been applied in the field of anomaly detection.

In this paper, we present **LoRIS** (Localized Reconstruction-by-Inpainting with a Single mask), a weakly supervised anomaly detection approach that uses the joint recess bounding box as prior knowledge during the inpainting. We also propose *Directional Distance* (DD), a new image similarity deviation metric that yields better anomaly segmentation results than existing metrics, like Multi-Scale Gradient Magnitude Similarity Deviation (MSGMSD) [28]. Experimental results, conducted on a dataset of 483 images, show that **LoRIS** is more accurate in detecting recess distention when using MSGMSD (image-level AUROC 0.78), outperforming state-of-the-art unsupervised techniques and providing similar results as a previous approach specifically designed for this problem [3]. Instead, considering the segmentation problem, **LoRIS** provides better results when adopting DD (Dice score of 0.35), outperforming the existing unsupervised techniques. To summarize, our main contributions are:

- We propose **LoRIS**, the first weakly supervised anomaly detection and segmentation technique for ultrasound images.
- We demonstrate, through a comprehensive evaluation and an ablation study, that a) state-of-the-art unsupervised anomaly detection approaches are ineffective in this domain, and that b) **LoRIS** is effective and also provides similar performance as a supervised solution. The advantages over the supervised solution are that **LoRIS** can be trained with images of a single class and also produces recess segmentation.
- We show an effective solution to automatically compute the location prior, hence achieving a fully automated detection pipeline at inference time.

2 Methodology

After defining the problem (Section 2.1), we describe the two main steps of **LoRIS**: *localized reconstruction* (Section 2.2) and *anomaly detection* (Section 2.3). Finally, Section 2.4 describes how to automatically compute the location prior.

2.1 Problem Formulation

Hemophilia is a rare disease and its management has improved dramatically in the last decade for two reasons. First, the use of ultrasound imaging emerged as a practical solution for the detection of recess joint distention, caused by joint bleeding [19]. A second factor is the increased availability of replacement treatments (coagulation factor VIII and factor IX) and non-replacement treatments [20]. This has led to a reduction in the number of acute bleeding episodes, including intra-articular bleeding [7], which is otherwise a common cause of recess distention. Since in this paper we consider a cohort of patients treated with these drugs, images of distended recesses are rarer than non-distended ones. For this reason, we propose to frame the problem as an anomaly detection task in which a distended recess represents the anomalous case.

Specifically, we address the problem of detecting the distension of the subquadricipital recess (SQR), which is the main recess of the knee joint. Our approach uses images of the longitudinal US scan of the knee joint, which are commonly used to diagnose SQR distention by medical practitioners [19].

2.2 Localized reconstruction

The *localized reconstruction module* takes in input an image of the longitudinal US scan of the knee joint and the recess bounding box location prior (see Fig. 1). The module first inpaints the area in the image defined by the location prior with a black rectangle. Then, it reconstructs the inpainted area using a specifically trained network. One advantage of reconstructing the detected recess area only is that this solution avoids reconstructing areas that are of no interest for the given problem and that, due to noise and high inter-patient variability, can be reconstructed imprecisely also for physiological (non-pathological) images.

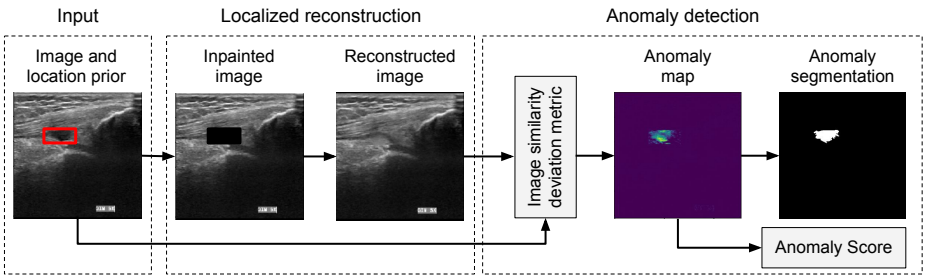

Fig. 1: **LoRIS** procedure schema

The network used is a U-Net [23], trained on a single class (images with non-distended recess) to reconstruct the inpainted image while focusing solely on the masked region. This is achieved through skip connections, directly propagating the information from low-level layers to the higher ones, facilitating the reconstruction process by preserving fine details, and maintaining contextual information from the original input. Consistently with previous works [30], we trained the network with the sum of three different losses:

$$\mathrm{L}_{\mathrm{tot}} = \mathrm{L}_{\mathrm{MSGMS}}(I, I_r) + \mathrm{L}_{\mathrm{SSIM}}(I, I_r) + \mathrm{L}_2(I, I_r)$$

where $\mathrm{L}_{\mathrm{MSGMS}}$ is the Multi-Scale Gradient Magnitude Similarity loss, $\mathrm{L}_{\mathrm{SSIM}}$ indicates the structural similarity index loss and the pixel-wise loss L_2 between the original image I and the reconstructed one I_r

At inference time, image reconstruction is achieved in a single iteration that reconstructs the entire masked area. This is in contrast with the approach of using multiple masks, adopted by existing reconstruction-by-inpainting techniques, that iteratively mask and reconstruct portions of the image, finally joining all the reconstructed areas to obtain the entire reconstructed image [30]. The problem with the multiple-masks approach is that, during its iterations, only a portion of the recess could be masked at a time, hence resulting in the image being precisely reconstructed even when the recess is distended. Instead, by using a single mask, the entire recess is inpainted, so it is more likely that it will be reconstructed as non-distended also in distended images, hence revealing the anomaly.

2.3 Anomaly detection

At inference time, **LoRIS** runs the localized reconstruction module to obtain the reconstructed image. Then, an *anomaly map* is computed, indicating an anomaly score for each pixel of the original image, using an image similarity deviation metric (see Fig 1). An overall anomaly score is computed at image level by average pooling the pixel-wise anomaly scores of the anomaly map. The anomaly is segmented by first selecting the set of pixels in the anomaly map whose value is above a threshold that maximizes the dice score (as in [16]) and

then by applying a post-processing step using morphological closing, followed by opening with kernel $3x3$.

In this paper, we propose a novel image similarity deviation metric called **directional difference** (DD) that is based on the following observation: a distended recess appears in a US image as a thick dark area, whereas a non-distended recess appears as a thin dark line on a lighter background. If an image containing a distended recess is provided in input, we expect the reconstruction to produce an image that resembles a non-distended recess, with the recess bounding box containing lighter pixels, on average, than the original image. The DD metric measures the increase of light intensity for the pixels in the reconstructed image with respect to the original one, ignoring the pixels where the light intensity actually decreases. Formally:

$$\mathrm{DD}(p, p_r) = \max((p_r - p), 0)$$

where p_r is the intensity of a pixel in the reconstructed image and p is the intensity of the corresponding pixel in the original image.

We experimented **LoRIS** also considering alternative image similarity deviation metrics. Some of them are derived from the existing literature on similarity deviation between images, including Gradient Magnitude Similarity Deviation (GMSD), and Multi-Scale Gradient Magnitude Similarity Deviation (MSGMSD) [31]. We also experimented with similarity scoring functions between images by computing their dual, such as the Structural Similarity Index (SSIM). Among all these image similarity deviation metrics, **LoRIS** obtained the best results with MSGMSD in terms of image-wise and pixel-wise accuracy, while best Dice score was obtained using **DD**.

2.4 Automatic detection of the recess bounding box

LoRIS requires the recess bounding box as location prior both at training and inference time. The use of (manually annotated) bounding box priors limits the real-world applicability of the proposed approach. To address this issue, we further propose the use of object detection for automatically annotating the bounding box location priors, thus achieving a fully automated pipeline (from image acquisition to anomaly prediction). Note that, also in this case, the object detection has to be trained on non-distended images only to maintain the applicability of the approach in the anomaly detection setting.

3 Experimental evaluation

This section describes the experimental methodology (Section 3.1), the experimental results in terms of anomaly detection and segmentation performance (Section 3.2), and the impact of automatic location prior detection (Section 3.3).

3.1 Experimental methodology

We used the same dataset as in [3], containing 483 US images of the knee recess, 123 of which are distended, according to the annotation of a physician who is an expert US reader in this specific field. The same physician also annotated the images with the recess bounding box (the location prior) and the recess segmentation, which is used to compute pixel-wise segmentation accuracy. The images are divided into 5 folds using patient-based splits, thus ensuring that no images of the same patient are simultaneously in the training and test folds. Due to this, the exact number of images in each fold can vary. Approximately, each fold contains 308, 78, and 97 images for the training, validation, and test sets, respectively. Note that the images of distended recesses in the training and validation sets are ignored for the training of the proposed anomaly detection technique. Therefore, for each fold, we use approximately 226, 63, and 97 images in the training, validation, and test sets, respectively.

For what concerns the model training, for each fold, the U-net model was trained for 1000 epochs with an early-stopping criterion of 50 epochs on the validation loss, a learning rate of 0.0001 with Adam optimizer [13], and a batch size of 4. Parameters were selected empirically. We ran the experiments on a Ubuntu Server with a partitioned NVIDIA A100 GPU, 42Gb of RAM, and an AMD EPYC 8-core CPU. The code is publicly available[4].

To assess the accuracy of the anomaly detection, we consider metrics commonly used in the state-of-the-art: Image-level AUROC (I-AUROC) and Pixel-level AUROC (P-AUROC). Additionally, we employ the Dice score as it more accurately evaluates the anomaly segmentation accuracy [4].

3.2 Anomaly detection and segmentation results

Table 1 compares **LoRIS** with state-of-the-art unsupervised anomaly detection approaches and a previously proposed supervised technique [3]. Considering the unsupervised techniques, recent ones (PatchCore [24], Simplenet [16] and Cflow [8]) yield the best results, with PatchCore having an I-AUROC of 0.701 and a P-AUROC of 0.871. However, unsupervised techniques have a Dice score lower than 0.2, showing that they do not obtain a relevant segmentation of the anomalous region. This is also exemplified in Fig.2 that shows, for three sample images, the segmentation results of various techniques. As shown in the figure, the unsupervised techniques fail in most of the cases to detect the recess area, and, even when they do, they do not approximate the recess accurately. The *multi-task* supervised technique [3] achieves a higher I-AUROC value of 0.780 but it cannot compute the recess segmentation.

Table 1 also shows the results of two variants of **LoRIS**, when using MSGMSD (**LoRIS+MSGMSD**) and DD (**LoRIS+DD**) as image similarity deviation metrics. The former achieves the best performance in terms of image-level AUROC (0.783) when compared with all other techniques, including *multi-task*.

[4] `https://github.com/warpcut/LoRIS`

Table 1: Anomaly detection and segmentation results

Model	Setting	I-AUROC	P-AUROC	Dice
RIAD [30]	Unsupervised	0.583 ± 0.083	0.682 ± 0.016	0.051 ± 0.017
InTrans [21]	Unsupervised	0.581 ± 0.053	0.574 ± 0.033	0.028 ± 0.009
Ganomaly [1]	Unsupervised	0.573 ± 0.035	-	-
FAIR [15]	Unsupervised	0.544 ± 0.035	0.668 ± 0.021	0.102 ± 0.012
Cflow [8]	Unsupervised	0.645 ± 0.125	0.864 ± 0.011	0.124 ± 0.049
Draem [29]	Unsupervised	0.547 ± 0.066	0.626 ± 0.041	0.033 ± 0.007
UAE [18]	Unsupervised	0.621 ± 0.068	0.699 ± 0.014	0.061 ± 0.018
Simplenet [16]	Unsupervised	0.68 ± 0.104	0.818 ± 0.01	0.144 ± 0.047
PatchCore [24]	Unsupervised	0.701 ± 0.090	0.871 ± 0.009	0.193 ± 0.066
Multi-task [3]	Supervised	0.780 ± 0.050	-	-
LoRIS+MSGMSD	Weakly-supervised	$\mathbf{0.783 \pm 0.050}$	$\mathbf{0.932 \pm 0.018}$	0.263 ± 0.042
LoRIS+DD	Weakly-supervised	0.750 ± 0.100	0.746 ± 0.047	$\mathbf{0.353 \pm 0.034}$

It also outperforms all other unsupervised techniques in terms of pixel-level AUROC (0.932). Taking into account the segmentation ability, **LoRIS+DD** achieves the best results, with a Dice score of 0.353. However, we note that the Dice score is still relatively low, indicating that accurate anomaly segmentation in this domain is particularly challenging. This observation is also supported by the results obtained using UAE [18] which, despite being designed for medical imaging, yields poor results (AUROC of 0.699 and dice of 0.061). Nevertheless, as shown in Fig.2, while segmentation is not extremely accurate, it approximates the actual recess shape well.

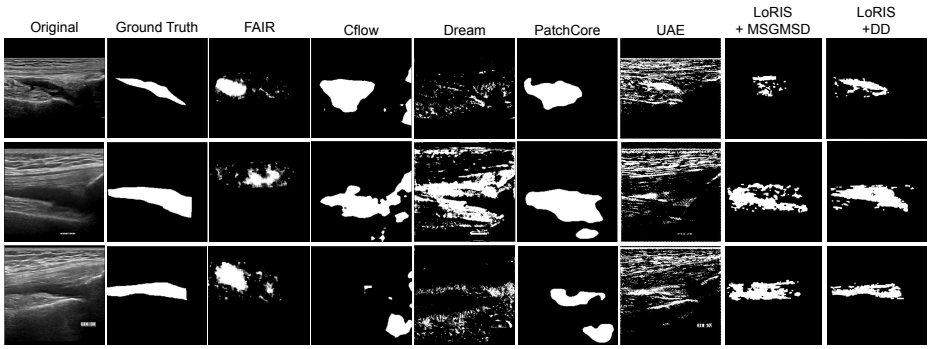

Fig. 2: Comparison of the anomaly segmentations generated by different techniques

3.3 Automated detection of the recess bounding box

For the automated detection of the bounding box location prior, we examine two object detection approaches, Yolo (V5) [11] and Co-DETR [32]. We trained the two models on the non-distended images in the training set and measured

the performance of **LoRIS+MSGMSD** with the location prior automatically computed by the trained object detection model at test time.

As shown in Table 2, YOLO fails to achieve results comparable to the upper baseline represented by the Ground Truth (GT) annotations. Indeed, there is a significant drop in performance: −4.4% in I-AUROC, −5.4% in P-AUROC and −6.2 in Dice score. Instead, using CoDETR, shown to perform better in several domains [32], the results remain comparable with those obtained with GT: −0.7% in I-AUROC, −2.2% in P-AUROC and −2.3% in Dice score.

Table 2: Performances of the object detection algorithms and their impact.

	precision	**map@50**	mAP@75	**I-AUROC**	**P-AUROC**	**Dice**
GT	-	-	-	0.783 ± 0.050	0.932 ± 0.018	0.263 ± 0.042
Yolo-V5	0.954 ± 0.044	0.796 ± 0.052	0.254 ± 0.063	0.773 ± 0.038	0.872 ± 0.033	0.223 ± 0.030
CoDETR	1.0 ± 0.0	0.9 ± 0.035	0.41 ± 0.066	0.776 ± 0.029	0.910 ± 0.038	0.240 ± 0.031

4 Conclusions

The approach proposed in this paper is the first anomaly detection technique to use a location prior and to adopt the reconstruction-by-inpainting approach on US images, which are noisy and have high variability. Experimental results show that the technique can separate normal images from anomalous ones better than state-of-the-art unsupervised approaches, achieving results comparable to a fully supervised approach proposed previously [3], when **LoRIS+MSGMSD** is used. Instead, **LoRIS+DD** yields the best results for the purpose of anomaly segmentation.

Furthermore, **LoRIS** has two additional benefits with respect to the supervised approach. First, it is trained using non-anomalous data only, and therefore it is more suitable to the target problem domain in which anomalous data is scarce. Second, it provides more anatomically reasonable anomaly segmentations, only requiring the recess bounding box as a location prior. We also show that this information can be obtained using a state-of-the-art object detection technique, achieving results comparable to the use of the manually annotated data and thus achieving a fully automated SQR distension detection pipeline.

As a future work, we will experiment different image similarity deviation metrics for different tasks: **MSGMSD** can be used to achieve a better SQR distension detection accuracy, while **DD** can be used to compute more accurate anomaly segmentation. We will also investigate the possibility of using the bounding box location prior in existing anomaly detection techniques, for example for delimiting the areas targeted with synthetic anomalies in Draem [29]. Finally, we will integrate the location prior detection to form an end-to-end solution, and we will explore how our anomaly detection approach performs on other medical imaging sources as well as other application domains.

Acknowledgments. This research is partially supported by the MUSA (Multilayered Urban Sustainability Action) and by the FAIR (Future Artificial Intelligence Research) projects both funded by the NextGeneration EU program. It is also partially supported by the Italian Ministry of Health – Bando Ricerca Corrente. The Hemostasis & Thrombosis Unit of the Fondazione IRCCS Ca' Granda Ospedale Maggiore Policlinico is member of the European Reference Network on Rare Haematological Diseases EuroBloodNet-Project ID No 101157011. ERN-EuroBloodNet is partly co-funded by the European Union within the framework of the Fourth EU Health Programme.

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

References

1. Akcay, S., Atapour-Abarghouei, A., Breckon, T.P.: Ganomaly: Semi-supervised anomaly detection via adversarial training. In: Asian Conference on Computer Vision. Springer (2019)
2. Asgari Taghanaki, S., Abhishek, K., Cohen, J.P., Cohen-Adad, J., Hamarneh, G.: Deep semantic segmentation of natural and medical images: a review. Artificial Intelligence Review **54**, 137–178 (2021)
3. Colussi, M., Civitarese, G., Ahmetovic, D., Bettini, C., Gualtierotti, R., Peyvandi, F., Mascetti, S.: Ultrasound detection of subquadricipital recess distension. Intelligent Systems with Applications (2023)
4. Eelbode, T., Bertels, J., Berman, M., Vandermeulen, D., Maes, F., Bisschops, R., Blaschko, M.B.: Optimization for medical image segmentation: theory and practice when evaluating with dice score or jaccard index. IEEE Transactions on Medical Imaging **39**(11), 3679–3690 (2020)
5. El Jurdi, R., Petitjean, C., Honeine, P., Cheplygina, V., Abdallah, F.: High-level prior-based loss functions for medical image segmentation: A survey. Computer Vision and Image Understanding **210**, 103248 (2021)
6. Feng, G., Zhang, L., Hu, Z., Lu, H.: Learning from box annotations for referring image segmentation. IEEE Transactions on Neural Networks and Learning Systems (2022)
7. Gualtierotti, R., Solimeno, L.P., Peyvandi, F.: Hemophilic arthropathy: current knowledge and future perspectives. Journal of Thrombosis and Haemostasis (2021)
8. Gudovskiy, D., Ishizaka, S., Kozuka, K.: Cflow-ad: Real-time unsupervised anomaly detection with localization via conditional normalizing flows. In: IEEE/CVF Winter Conference on Applications of Computer Vision (2022)
9. Hilgartner, M.W.: Current treatment of hemophilic arthropathy. Current opinion in pediatrics **14**(1), 46–49 (2002)
10. Huang, Q., Zhang, F., Li, X.: Machine learning in ultrasound computer-aided diagnostic systems: a survey. BioMed research international **2018** (2018)
11. Jocher, G., Chaurasia, A., Stoken, A., Borovec, J., NanoCode012, Kwon, Y., TaoXie, Fang, J., imyhxy, Michael, K., Lorna, V, A., Montes, D., Nadar, J., Laughing, tkianai, yxNONG, Skalski, P., Wang, Z., Hogan, A., Fati, C., Mammana, L., AlexWang1900, Patel, D., Yiwei, D., You, F., Hajek, J., Diaconu, L., Minh, M.T.: ultralytics/yolov5: v6.1 (Feb 2022). https://doi.org/10.5281/zenodo.6222936
12. Kervadec, H., Dolz, J., Wang, S., Granger, E., Ayed, I.B.: Bounding boxes for weakly supervised segmentation: Global constraints get close to full supervision. In: Medical imaging with deep learning. pp. 365–381. PMLR (2020)

13. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
14. Kulharia, V., Chandra, S., Agrawal, A., Torr, P., Tyagi, A.: Box2seg: Attention weighted loss and discriminative feature learning for weakly supervised segmentation. In: European Conference on Computer Vision. pp. 290–308. Springer (2020)
15. Liu, T., Li, B., Du, X., Jiang, B., Geng, L., Wang, F., Zhao, Z.: Fair: Frequency-aware image restoration for industrial visual anomaly detection. arXiv preprint arXiv:2309.07068 (2023)
16. Liu, Z., Zhou, Y., Xu, Y., Wang, Z.: Simplenet: A simple network for image anomaly detection and localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20402–20411 (2023)
17. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. Nature Communications **15**(1), 654 (2024)
18. Mao, Y., Xue, F.F., Wang, R., Zhang, J., Zheng, W.S., Liu, H.: Abnormality detection in chest x-ray images using uncertainty prediction autoencoders. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part VI 23. pp. 529–538. Springer (2020)
19. Martinoli, C., Alberighi, O.D.C., Di Minno, G., Graziano, E., Molinari, A.C., Pasta, G., Russo, G., Santagostino, E., Tagliaferri, A., Tagliafico, A., Morfini, M.: Development and definition of a simplified scanning procedure and scoring method for haemophilia early arthropathy detection with ultrasound (head-us). Thrombosis and haemostasis **109**(6), 1170–1179 (2013)
20. Peyvandi, F., Garagiola, I., Biguzzi, E.: Advances in the treatment of bleeding disorders. Journal of Thrombosis and Haemostasis **14**(11), 2095–2106 (2016)
21. Pirnay, J., Chai, K.: Inpainting transformer for anomaly detection. In: International Conference on Image Analysis and Processing. pp. 394–406. Springer (2022)
22. Plut, D., Kotnik, B.F., Zupan, I.P., Kljucevsek, D., Vidmar, G., Snoj, Z., Martinoli, C., Salapura, V.: Diagnostic accuracy of haemophilia early arthropathy detection with ultrasound (head-us): a comparative magnetic resonance imaging (mri) study. Radiology and oncology **53**(2), 178–186 (2019)
23. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention. Springer (2015)
24. 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)
25. Tan, J., Hou, B., Day, T., Simpson, J., Rueckert, D., Kainz, B.: Detecting outliers with poisson image interpolation. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part V 24. pp. 581–591. Springer (2021)
26. Tschuchnig, M.E., Gadermayr, M.: Anomaly detection in medical imaging-a mini review. In: International Data Science Conference. Springer (2022)
27. Tyrrell, P., Blanchette, V., Mendez, M., Paniukov, D., Brand, B., Zak, M., Roth, J.: Detection of joint effusions in pediatric patients with hemophilia using artificial intelligence-assisted ultrasound scanning; early insights from the development of a self-management tool. Res Pract Thromb Haemost **5** (2021)
28. Xue, W., Zhang, L., Mou, X., Bovik, A.C.: Gradient magnitude similarity deviation: A highly efficient perceptual image quality index. IEEE transactions on image processing **23**(2), 684–695 (2014)

29. Zavrtanik, 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. pp. 8330–8339 (2021)
30. Zavrtanik, V., Kristan, M., Skočaj, D.: Reconstruction by inpainting for visual anomaly detection. Pattern Recognition **112**, 107706 (2021)
31. Zhang, B., Sander, P.V., Bermak, A.: Gradient magnitude similarity deviation on multiple scales for color image quality assessment. In: International Conference on Acoustics, Speech and Signal Processing. pp. 1253–1257. IEEE (2017)
32. Zong, Z., Song, G., Liu, Y.: Detrs with collaborative hybrid assignments training. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6748–6758 (2023)