SymmDiff: Symmetry-Guided Diffusion for Fracture Localization in Pelvic Radiographs

Abdul Rahman[1], Afnan Ghafoor[2], and Bumshik Lee[1⋆]

[1] Energy AI Track, Korea Institute of Energy Technology (KENTECH), South Korea
[2] Department of Information and Communication Engineering, Chosun University, South Korea
bslee@kentech.ac.kr

Abstract. Unsupervised anomaly detection (UAD) offers a promising direction for fracture analysis where abnormal annotations are scarce, yet pelvic X-ray fracture detection remains exceptionally challenging due to overlapping anatomy, projection artifacts, and the inherently asymmetric presentation of fracture patterns. We introduce SymmDiff, a symmetry-aware diffusion framework that embeds contralateral anatomical consistency directly into the generative correction process, mirroring the clinical practice of bilateral comparison. By pairing anatomically corresponding left-right regions and enforcing symmetry-consistent reconstruction during diffusion, SymmDiff generates a normal-corrected radiograph that selectively suppresses fracture-induced asymmetries while preserving global anatomical structure, enabling precise fracture localization through residual discrepancies between the input image and its symmetry-corrected counterpart. Extensive evaluations on pelvic radiographs spanning subtle to severe fractures show that SymmDiff consistently outperforms existing anomaly-detection baselines, highlighting the effectiveness of symmetry-guided generative modeling for reliable fracture localization in clinical imaging. Code is available at https://itsabdulrahman.github.io/symmdiff.

Keywords: Diffusion model · Unsupervised anomaly detection · Pelvic fracture.

1 Introduction

Recent advances in deep learning have substantially improved X-ray image analysis and enabled AI systems that can provide practical decision support in radiology work-flows [7,9]. In musculoskeletal imaging, pelvic X-rays are fast, inexpensive, and widely used for fracture screening, yet detection/localization remains challenging due to overlapping anatomy, projection/position variability, and low-contrast fracture cues.

A key limitation of many current clinical AI tools is that they are typically trained for predefined disease labels [12], which is often brittle for pelvic radiographs, where subtle fractures appear as small deviations within highly variable

⋆ Corresponding author

but normal anatomy. This motivates unsupervised anomaly detection (UAD), not merely for reducing annotation burden, but also to model normal structure and highlights deviations. Feature-based UAD methods have shown strong results in general anomaly detection [14], but may miss fine grained, anatomy specific cues. Generative methods can reconstruct a normal looking counterpart of the input and localize pathology via the reconstruction discrepancy, which is well-suited to radiographs where abnormalities appear as breaks from expected structural continuity.

Diffusion models have become a strong foundation for reconstruction-based anomaly detection [4,17,18,10,11]. However, pelvic radiographs pose a key challenge: normal anatomical and acquisition variability can dominate differences, causing false positives around normal structures and weak sensitivity to microfractures when "normal" is not anchored to a stable structural reference. Importantly, radiologists often mitigate this ambiguity via side-to-side comparison: contrasting the suspected side with the contralateral side. Prior work has shown that explicitly modeling contralateral comparison can improve fracture diagnosis [6]. Motivated by this clinical practice, we argue that pelvic X-ray UAD should not be anatomy-agnostic; instead, it should incorporate contralateral consistency as an inductive bias. However, post-hoc residual/feature comparisons are insufficient for diffusion-based UAD, as iterative denoising can drift anatomically unless contralateral coupling is embedded in the reverse transition dynamics.

We therefore propose SymmDiff, a diffusion-based UAD framework for pelvic radiographs that embeds contralateral comparison into the denoising/correction process to generate a contralateral-consistent pseudo-healthy reconstruction. Fractures and implant-related deviations are localized via input-reconstruction discrepancies, reducing projection/anatomy-induced false positives while improving sensitivity to subtle fracture cues. The contributions of this paper can be summarized as follows:

1. A symmetry conditioned latent diffusion model that establishes a patient specific representation of normal pelvic anatomy through contralateral correspondence, providing a structural prior for fracture localization.
2. A symmetry conditioned reverse diffusion formulation in which contralateral information guides the reverse transition at each step, stabilizing the anatomical trajectory during iterative denoising.
3. An ROI-guided reverse diffusion scheme with contralateral latent swap initialization that softly restricts updates to regions exhibiting elevated bilateral discrepancy, reducing projection-induced over-correction while preserving global anatomy.

2 Background

2.1 Generative Diffusion for Medical UAD

Diffusion models learn to invert a gradual corruption process, as popularized by DDPMs [8], and have become a strong backbone for reconstruction-based

UAD. Trained on normal anatomy, they produce a pseudo-normal correction at test time and localize anomalies via input-correction discrepancies, as explored in early diffusion-UAD [18] and later refinement-based variants that better preserve unaffected regions [4]. To improve stability and efficiency under acquisition variability, many radiograph pipelines operate in latent space following latent diffusion principles [13]. Concretely, an encoder E maps an image x to a latent $z = E(x)$, and diffusion is applied to z rather than x. With a variance schedule $\{\beta_t\}$, the forward process is

$$q(z_t \mid z_{t-1}) = \mathcal{N}(\sqrt{1-\beta_t} z_{t-1}, \beta_t \mathbf{I}), \tag{1}$$

and a denoiser is trained to predict the injected noise at arbitrary timesteps [8]. At inference, deterministic DDIM sampling is often used to obtain stable and efficient reconstructions without stochastic variation [15].

Despite strong performance, diffusion-based UAD in radiographs remains vulnerable: projection, overlap, and contrast variability can suppress residual cues and hide subtle pathology, creating a mismatch between generic denoising and localization needs, where preserving normal anatomy is critical [1]. Moreover, most diffusion-UAD methods encode anatomy only implicitly via global density, which is often insufficient to stabilize "normality" in highly symmetric regions such as the pelvis.

2.2 Anatomy of Symmetry in Radiographic Assessment

Pelvic X-rays are often read via side-to-side comparison, using the contralateral side as an internal control under similar projection to disambiguate subtle cues from overlap and normal variants. While large-scale supervised pelvic X-ray systems demonstrate strong performance for trauma findings, they rely on predefined labels [5]. Mathematically, let Ω_{L} and Ω_{R} denote left/right domains and $\mathcal{T}$ denote the approximate reflection/alignment mapping. For a healthy pelvis, corresponding features are highly dependent: $f(p) \approx f(\mathcal{T}(p))$, $p \in \Omega_L$, which implies that conditioning on the opposite side can reduce uncertainty about what is anatomically plausible at a given location. Formally, learning a normality model conditioned on a patient-specific reference gives the generative process a structural anchor, reducing the space of plausible "normal correction".

Cross-side comparison modules have improved fracture diagnosis by reducing ambiguity from local textures and edges [6]. However, most diffusion-based UAD methods do not impose such cross-side consistency within the correction dynamics, which can over-correct normal structures and increase false responses in overlap-heavy pelvic regions. Unlike direct left-right subtraction, SymmDiff enforces structural consistency along the reverse trajectory, enabling correction rather than mere comparison.

2.3 Structural Inductive Biases in Latent Diffusion

A recurring challenge in reconstruction-style UAD is the "reconstruction paradox": highly expressive models can also reconstruct abnormal patterns, weak-

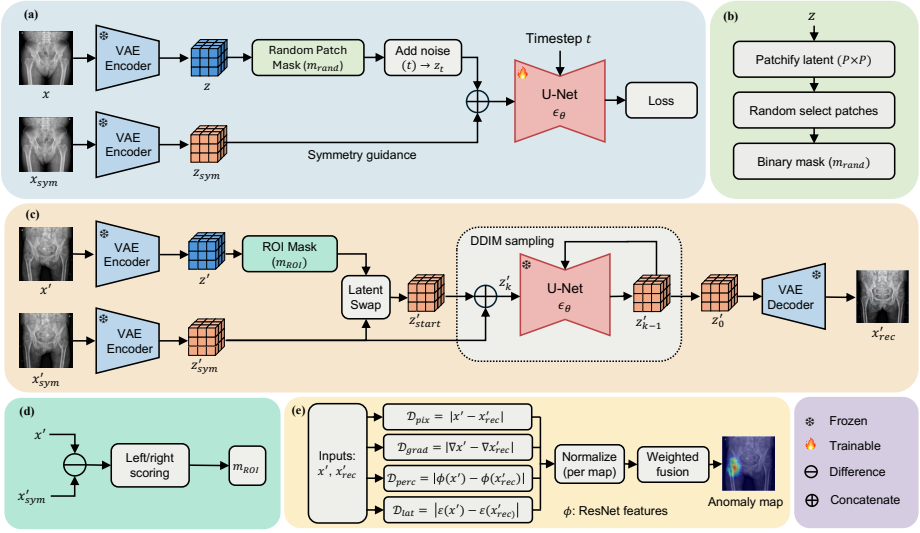

Fig. 1. (a) Training with input x and its flipped counterpart x_{sym}; (b) mask generation; (c) inference pipeline; (d) left/right ROI scoring; (e) anomaly map fused differences.

ening residual-based localization. This has motivated constrained and masking-based diffusion strategies in medical UAD [16]. In latent diffusion, selectivity can be expressed by restricting reverse updates to a subset of latent coordinates. Let $m \in \{0,1\}^{h \times w}$ be a binary mask and Φ_t a deterministic reverse update [15]. A generic masked reverse step is

$$z_{t-1} = (1 - m) \odot z_t + m \odot \Phi_t(z_t), \tag{2}$$

where $\odot$ denotes element-wise multiplication, enforcing invariance outside the mask while allowing correction inside it. Such constrained correction reduces the risk of altering normal structure and makes residual maps more interpretable, complementing feature-based anomaly scoring approaches that do not explicitly enforce preservation constraints [14].

3 Methodology

We propose SymmDiff, a symmetry-guided latent diffusion framework for pelvic X-ray fracture localization that uses the patient-specific contralateral side as a structural prior to generate a pseudo-healthy reconstruction and localize anomalies via input-reconstruction discrepancies. Figure 1 shows: SymmDiff consists of (i) symmetry-aware latent encoding of the input and flipped reference, (ii) symmetry-constrained reverse diffusion with ROI-targeted, cross-side coupled correction, and (iii) multi-scale discrepancy fusion, enforcing coupling during the correction trajectory rather than post-hoc.

3.1 Latent Space and Symmetry Mapping

Let $x \in \mathbb{R}^{H\times W}$ denote an input pelvic radiograph and $x_{sym} = \mathcal{T}(x)$ its horizontally flipped reference, where $\mathcal{T}$ denotes reflection. Both are encoded by a pretrained VAE adapted from [13] and fine-tuned on pelvic radiographs into latent codes $z = \varepsilon(x)$, $z_{sym} = \varepsilon(x_{sym})$, $z, z_{sym} \in \mathbb{R}^{h\times w\times c}$. While $\mathcal{T}$ provides a coarse structural alignment, latent encoding ε abstracts local textures into higher-level features, improving robustness to minor projection mismatch and tilt.

Pelvic anatomy exhibits strong bilateral correspondence, so reference conditioning can be more specific than an unconditional model. Let Z and R denote latent random variables of z and z_{sym} under normal data, where $I(\cdot;\cdot)$ denotes mutual information and $H(\cdot)$ denotes entropy. Since contralateral anatomy is not independent, $I(Z;R) > 0$, thus conditioning reduces uncertainty, i.e., $H(Z \mid R) < H(Z)$. This reduction holds under slight positional variability as diffusion reconciles correspondence on the latent manifold, narrowing plausible normal reconstructions and stabilizing the reverse trajectory, motivating conditioning on z_{sym} to reduce ambiguity from projection and overlap.

3.2 Training: Symmetry-Constrained Diffusion

We define a latent diffusion process over T steps with a variance schedule $\{\beta_t\}_{t=1}^{T}$, $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{s=1}^{t} \alpha_s$. Given a clean latent $z_0 = z$, we sample:

$$z_t = \sqrt{\bar{\alpha}_t} z_0 + \sqrt{1-\bar{\alpha}_t}\epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}). \tag{3}$$

A conditional denoiser ϵ_θ(U-Net) is trained to predict ϵ given z_t, t, and the reference latent z_{sym}. To encourage the network to use global structure rather than memorizing local textures, we apply random patch masking to the noisy latent, following the masking strategy used in [2], while z_{sym} provides a patient-specific anatomical prior that restores globally consistent anatomy. Let $m_{rand} \in \{0,1\}^{h\times w}$ be a binary mask generated by selecting random patches of size $P \times P$ and setting their entries to 0; we broadcast m_{rand} across channels. We form a masked latent

$$\tilde{z}_t = m_{rand} \odot z_t + (1 - m_{rand}) \odot \xi, \tag{4}$$

where ξ denotes a replacement value (0 or Gaussian noise). The training objective is:

$$\mathcal{L}_{\mathrm{SD}}(\theta) = \mathbb{E}_{z_0,\epsilon,t}\left[\|\epsilon - \epsilon_\theta(\tilde{z}_t, t; z_{sym})\|_2^2\right] \tag{5}$$

The denoiser parameterizes the conditional reverse transition

$$z_{t-1} = \Phi_{\theta,t}(z_t, z_{sym}) \Leftrightarrow p_\theta(z_{t-1} \mid z_t, z_{sym}) \tag{6}$$

Therefore, the contralateral reference explicitly shapes the reverse trajectory $z_t \to \cdots \to z_0$. In practice, the conditioning is implemented by concatenating z_{sym} with $\tilde{z}_t$ channel-wise, enabling ϵ_θ to inpaint masked regions using patient-specific bilateral context.

3.3 Inference: ROI-Targeted Reconstruction

Unlike the random masking used during training, the inference ROI mask is anatomically driven and constrains reverse updates to suspicious regions. Given a test image x', we compute its reference $x'_{sym} = \mathcal{T}(x')$ and latents $z' = \varepsilon(x')$, $z'_{sym} = \varepsilon(x'_{sym})$. We estimate a region of interest mask $m_{\text{ROI}} \in \{0,1\}^{h\times w}$ from a side-to-side discrepancy score: $m_{\text{ROI}} = \mathcal{M}(S(x', x'_{sym}))$, where S is a left-right scoring function and $\mathcal{M}$ binarizes it. Although S is based on intensity residuals, it is used only as a soft spatial gate (not a hard pre-filter), reducing the correction search space and limiting projection-induced over-correction. We initialize correction via latent swap:

$$z'_{\text{start}} = (1 - m_{\text{ROI}}) \odot z'_t + m_{\text{ROI}} \odot z'_{sym} \tag{7}$$

$$\|z'_{\text{start}} - z_N\|_2 \leq \|z'_t - z_N\|_2 \tag{8}$$

where z_N denotes the normal latent. Critically, this swap biases the trajectory toward a symmetry-consistent manifold by anchoring the ROI to the patient's own contralateral anatomy, making the pipeline robust to minor mask errors. Although the ROI is derived from a simple side-to-side intensity discrepancy, it is used only as a soft spatial prior rather than strict segmentation. The reverse update remains fully data-driven, and the ROI merely narrows the correction search space without dictating the reconstruction. We then denoise z'_{start} using deterministic DDIM ($\eta = 0$) to obtain z'_0, conditioned on z'_{sym}. Let $\Phi_{\theta,t}(\cdot; z'_{sym})$ denote one DDIM update induced by ϵ_θ. We explicitly constrain the reverse trajectory to update only within the ROI:

$$z'_{t-1} = (1 - m_{\text{ROI}}) \odot z'_t + m_{\text{ROI}} \odot \Phi_{\theta,t}(z'_t; z'_{sym}) \tag{9}$$

where $\Phi_{\theta,t}$ represents the symmetry-conditioned reverse update approximating $p(z_{t-1} \mid z_t, z_{sym})$. For any spatial index p with $m_{\text{ROI}}(p) = 0$, Eq. (9) implies $z'_{t-1}(p) = z'_t(p)\,\forall t$, i.e., the reverse process is identity outside the ROI. This identity constraint outside the ROI ensures that imperfect mask estimation cannot induce global structural drift. Finally, the pseudo-healthy reconstruction is obtained by decoding $x_{\text{rec}} = \mathcal{D}(z'_0)$, where $\mathcal{D}$ is the frozen VAE decoder.

3.4 Multi-scale Anomaly Fusion

A single residual is unreliable in pelvic radiographs since abnormalities appear at different scales and different cues (edges, texture, global shape). We therefore use multi-view discrepancy x' and x_{rec}: pixel residual (intensity deviation), gradient residual (cortical edge breaks), perceptual distance (context/texture change robust to photometric shifts), and latent discrepancy (global mismatch in the generative space). These cues are complementary, thin fractures are often stronger in gradient perceptual maps, while broader deformities dominate pixel/latent. We fuse them via a convex combination:

$$\mathcal{A} = \sum_{k\in\{\text{pix,grad,perc,lat}\}} \omega_k \cdot \text{Norm}(\mathcal{D}_k), \quad \omega_k \geq 0, \sum_k \omega_k = 1. \tag{10}$$

Table 1. Quantitative comparison on Chosun University Hospital pelvic X-ray dataset, best results are highlighted in **bold**.

Method	Image-level			Pixel-level		
	AUROC	AUPRC	$F1_{max}$	AUROC	AUPRC	$F1_{max}$
DDPM [8]	84.5	83.0	80.1	76.5	23.3	24.6
AnoDDPM [18]	87.5	85.4	83.0	79.5	26.1	27.8
AutoDDPM [3]	89.0	85.9	84.2	80.6	27.6	27.0
THOR [4]	89.3	89.1	88.3	84.5	36.2	37.2
SymmDiff (ours)	**94.2**	**91.6**	**91.1**	**88.4**	**43.9**	**42.7**

Values in %

where $\mathcal{A}$ denotes the fused pixel-wise anomaly map, D_k each discrepancy map, and ω_k are fixed on validation for stability and interpretability.

4 Experimental Results

4.1 Experimental Setting

Dataset. We evaluate on a retrospective pelvic X-ray dataset collected at Chosun University Hospital, South Korea, from 2015-2025, with 1,268 normal radiographs (80/10/10 train/val/test) and 534 abnormal cases used only for evaluation, with pixel-level masks available for localization. Images are resized to 256×256, diffusion is performed in the VAE latent space with spatial size 32×32.

Evaluation Metrics. We report AUROC, AUPRC, and $F1_{max}$ at both the image level and pixel level. $F1_{max}$ is obtained by sweeping thresholds, with operating points selected on the validation split.

Implementation Details. SymmDiff is trained in the latent space with DDPM objective using $T = 1000$ diffusion steps. The VAE uses an $8\times$ downsampling latent representation, and the denoiser is an attention U-Net with timestep embedding and channel-wise z_{sym} conditioning. Random patch masking uses $P = 4 \times 4$ and masking ratio $\rho \in [0.6, 0.75]$. Inference uses deterministic DDIM with $K = 25$ steps. We compare against the diffusion-based UAD baselines: AnoDDPM, THOR, DDPM, and AutoDDPM under the same preprocessing and comparable sampling budgets

4.2 Results

Quantitative Results. Table 1 reports a quantitative comparison against diffusion-based UAD baselines on the Chosun University Hospital pelvic X-ray dataset. SymmDiff consistently achieves the best performance at both image and pixel levels, demonstrating clear gains over recent methods. Compared to

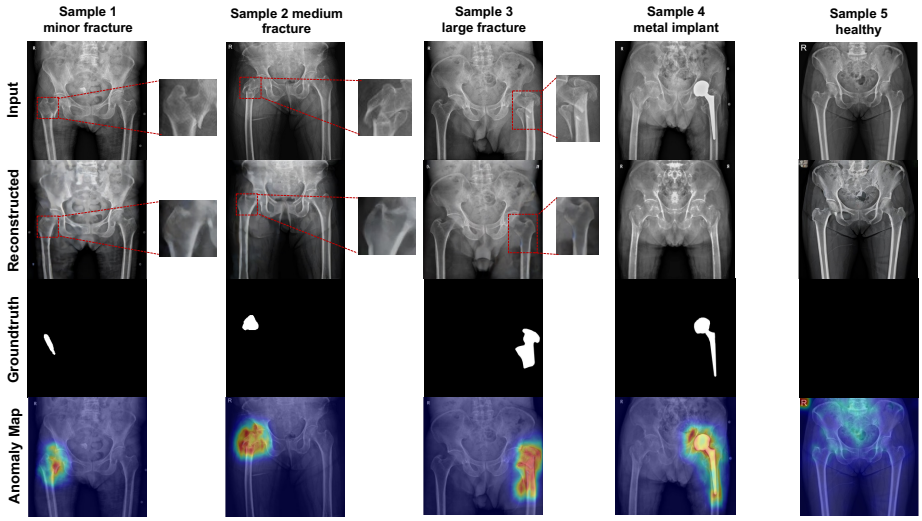

Fig. 2. Qualitative localization results on pelvic X-rays. Columns show five different test samples. The top row is the input radiographs, second row is the reconstructed corrected image, third and final rows show binary ground-truth masks and predicted anomaly heatmaps, respectively.

the second-best approach (THOR), SymmDiff improves image-level AUROC/AUPRC/$F1_{max}$ by 4.9/2.5/2.8, and pixel-level AUROC/AUPRC/$F1_{max}$ by 3.9/7.7/5.5. The improvements are even larger relative to the vanilla DDPM baseline, confirming that SymmDiff offers improved screening performance, greater stability, and clearer localization cues.

Qualitative Results. Figure 2 shows representative cases. Across minor-to-large fractures (Samples 1–3), SymmDiff yields anatomically plausible reconstructions and localizes anomalies consistent with the masks. For the implant

Table 2. Ablation study of SymmDiff components on the Chosun University Hospital pelvic X-ray dataset.

Method	Image-level			Pixel-level		
	AUROC	AUPRC	$F1_{max}$	AUROC	AUPRC	$F1_{max}$
SymmDiff (full)	**94.2**	**91.6**	**91.1**	**88.4**	**43.9**	**42.7**
w/o reference conditioning	89.4	88.2	86.2	79.9	32.4	30.6
w/o random patch masking	93.7	90.8	89.7	86.3	36.0	35.7
w/o ROI constrained de-noising	91.8	90.3	89.8	83.2	35.7	35.1
w/o latent swap	92.6	90.5	89.8	84.5	38.5	37.7
w/o multi-scale fusion (pixel only)	93.9	91.2	90.7	86.6	36.8	35.9

Values in %

case (Sample 4), responses concentrate around the implant with minimal spillover, and for a healthy case (Sample 5), activations remain low.

Ablations. Table 2 shows that each SymmDiff component contributes to performance. Removing symmetry conditioning causes the largest drop, indicating the side-to-side prior stabilizes reconstruction and reduces contralateral false positives. Disabling ROI-constrained denoising also degrades localization due to global drift and spurious activations outside the suspicious side. Random patch masking and latent swap provide additional, consistent gains, improving both screening and pixel-level localization over their respective ablations. Finally, replacing multi-scale fusion with a single pixel residual reduces pixel-level AUPRC/$F1_{max}$, confirming that this fusion provides complementary evidence for thin fractures that are weak in intensity residuals alone.

5 Conclusion

We introduced SymmDiff, a symmetry-guided latent diffusion framework that leverages a patient-specific contralateral reference to guide pseudo-healthy reconstruction and improve anomaly localization in pelvic radiographs. By combining reference-conditioned denoising with ROI-constrained correction, SymmDiff helps reduce anatomy- and projection-driven false responses and provides clearer localization cues. Experiments on the Chosun University Hospital pelvic X-ray dataset show consistent performance gains over diffusion-based UAD baselines at both image and pixel levels, and ablations indicate that each component contributes meaningfully to the overall framework. Future work will explore improving robustness under strong pose and projection mismatch, refining ROI estimation and alignment, and evaluating cross-site generalization. In bilateral abnormal cases, ROI gating with identity constraints outside the mask provides an additional stabilizing bias during reverse diffusion, reducing sensitivity to imperfect contralateral references. Looking ahead, we also plan to extend beyond contralateral priors by incorporating population-level healthy atlases for symmetry-violating scenarios.

Acknowledgments. This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea Government (MSIT) (No. RS-2023-00217471)

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

References

1. Beizaee, F., Hajimiri, S., Ben Ayed, I., Lodygensky, G., Desrosiers, C., Dolz, J.: Reflect: Rectified flows for efficient brain anomaly correction transport. In: Medical Image Computing and Computer-Assisted Intervention – MICCAI (2025)

2. Beizaee, F., Lodygensky, G.A., Desrosiers, C., Dolz, J.: Correcting deviations from normality: A reformulated diffusion model for multi-class unsupervised anomaly detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 19088–19097 (2025)
3. Bercea, C.I., Neumayr, M., Rueckert, D., Schnabel, J.A.: Mask, stitch, and re-sample: Enhancing robustness and generalizability in anomaly detection through automatic diffusion models. In: ICML 3rd Workshop on Interpretable Machine Learning in Healthcare (IMLH) (2023)
4. Bercea, C.I., Wiestler, B., Rueckert, D., Schnabel, J.A.: Diffusion models with implicit guidance for medical anomaly detection. In: Medical Image Computing and Computer-Assisted Intervention – MICCAI (2024)
5. Cheng, C.T., , et al.: A scalable physician-level deep learning algorithm detects universal trauma on pelvic radiographs. Nature Communications **12**(1) (2021)
6. Gao, Y., , et al.: Cce-net: A rib fracture diagnosis network based on contralateral, contextual, and edge enhanced modules. Biomedical Signal Processing and Control **75**, 103620 (2022)
7. Ghafoor, A., Moon, S.Y., Lee, B.: Multiclass segmentation using teeth attention modules for dental x-ray images. IEEE Access **11**, 123891–123903 (2023). https://doi.org/10.1109/ACCESS.2023.3329364
8. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: Advances in Neural Information Processing Systems (2020)
9. Jamal, M.A., Chao, G., Lee, B.: Dentseg-eda: A 3d tooth segmentation framework for cbct images using enhanced dual attention. Neurocomputing **655** (2025)
10. Liang, Z., Guo, X., Noble, J.A., Kamnitsas, K.: Itermask2: Iterative unsupervised anomaly segmentation via spatial and frequency masking for brain lesions in mri. In: Medical Image Computing and Computer-Assisted Intervention – MICCAI (2024)
11. Marimont, S.N., Siomos, V., Baugh, M., Tzelepis, C., Kainz, B., Tarroni, G.: Ensembled cold-diffusion restorations for unsupervised anomaly detection. In: Medical Image Computing and Computer-Assisted Intervention – MICCAI (2024)
12. Muhammad, I., Rao, R.S., Lee, B.: Bone-net: A novel hybrid deep-learning model for effective osteoporosis detection. PLOS ONE **20**(10) (2025)
13. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10674–10685 (2022)
14. 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 (CVPR). pp. 14318–14328 (2022)
15. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: International Conference on Learning Representations (ICLR) (2021)
16. Wolleb, J., Bieder, F., Friedrich, P., Zhang, P., Durrer, A., Cattin, P.C.: Binary noise for binary tasks: Masked bernoulli diffusion for unsupervised anomaly detection. In: Medical Image Computing and Computer-Assisted Intervention – MICCAI (2024)
17. Won, J., Kim, J., Oh, J., Yoo, Y., Yum, J., Lee, J., Park, J.H., Jo, W., Nam, Y., Lee, H., Hong, G., Kim, N.: Generative unsupervised anomaly detection with coarse-fine ensemble for workload reduction in 3d non-contrast brain ct of emergency room. In: Medical Image Computing and Computer-Assisted Intervention – MICCAI (2025)

18. 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 (CVPR) Workshops. pp. 650–656 (2022)