

Denoising Diffusion Models for Anomaly Localization in Medical Images

Cosmin I. Bercea^{1,2} , Philippe C. Cattin³ , Julia A. Schnabel^{1,2,4} , Julia Wolleb^{3,5,6} 

- 1** School of Computation, Information and Technology, Technical University of Munich, Germany
2 Institute of Machine Learning in Biomedical Imaging, Helmholtz Munich, Germany
3 Department of Biomedical Engineering, University of Basel, Allschwil, Switzerland
4 School of Biomedical Engineering and Imaging Sciences, King's College London, United Kingdom
5 Department of Biomedical Informatics and Data Science, Yale University School of Medicine, New Haven, CT, USA
6 Yale Biomedical Imaging Institute, Yale University, New Haven, CT, USA

Abstract

This review explores anomaly localization in medical images using denoising diffusion models. After providing a brief methodological background of these models, including their application to image reconstruction and their conditioning using guidance mechanisms, we provide an overview of available datasets and evaluation metrics suitable for their application to anomaly localization in medical images. In this context, we discuss supervision schemes ranging from fully supervised segmentation to semi-supervised, weakly supervised, self-supervised, and unsupervised methods, and provide insights into the effectiveness and limitations of these approaches. Furthermore, we highlight open challenges in anomaly localization, including detection bias, domain shift, computational cost, and model interpretability. Our goal is to provide an overview of the current state of the art in the field, outline research gaps, and highlight the potential of diffusion models for robust anomaly localization in medical images.

Keywords

Anomaly Detection, Denoising Diffusion Models, Generative Models

Article informations

<https://doi.org/10.59275/j.melba.2025-c586>

Volume 3, Received: 2025-03-31, Published 2025-11-27

Corresponding author: julia.wolleb@yale.edu

©2025 Bercea, Cattin, Schnabel, and Wolleb. License: CC-BY 4.0



1. Introduction

Anomaly localization in medical images refers to the process of identifying abnormal areas or regions within images of various modalities, such as X-ray, computed tomography (CT), magnetic resonance imaging (MRI), or optical coherence tomography (OCT), see Fig. 1.

These anomalies may be indicative of a range of conditions or diseases, including tumors, fractures, organ malformations, and vascular abnormalities. Precisely locating and highlighting these pathological changes assists medical experts in diagnosing and monitoring diseases. However, manual analysis is a tedious task requiring significant time and expertise for accurate interpretation. Deep learning algorithms have played a key role in the automatic processing of medical images (Chen et al., 2022; Zhou et al., 2023). Specifically, fully supervised segmentation algorithms like

U-Nets (Ronneberger et al., 2015) have been instrumental in automatic lesion segmentation, providing detailed delineation of abnormal areas. In addition, generative models such as Variational Autoencoders (VAEs) (Kingma et al., 2019) and Generative Adversarial Networks (GANs) (Goodfellow et al., 2020) have shown promising results in anomaly localization, even with limited labeled data. In recent years, denoising diffusion models have emerged as a class of deep generative models that estimate complex probability distributions (Ho et al., 2020). They have gained attention for their ability to generate high-quality samples, outperforming GANs in image synthesis (Dhariwal and Nichol, 2021). While diffusion models have primarily been proposed for image generation, their ability to generate healthcare data has enabled real-world applications that address data scarcity and privacy concerns (Giuffrè and Shung, 2023). In addition, they have been adapted to data augmentation and to solve

downstream tasks like denoising, inpainting, and image restoration. In this review, we focus on the methodological aspects of denoising diffusion models (DDMs) for lesion localization, a direction with strong potential to support radiologists in the future (Prasad et al., 2024). Regulatory frameworks are still evolving, with consensus guidelines introduced in Lekadir et al. (2025) and the EU Artificial Intelligence Act (Edwards, 2021) representing the first comprehensive legislation regulating AI across the European Union.

In this work, we investigate the potential and limitations of diffusion models for anomaly localization in medical imaging under various data and label availability scenarios. Anomaly segmentation refers to delineating pathological regions in an image, which is a pixel-level classification task that aims to precisely trace the extent of abnormalities. This fully supervised approach typically assumes an anomaly is present and focuses on where and how it manifests. In contrast, anomaly detection determines whether an anomaly exists, and can be applied at the image, bounding-box, or pixel level. In this review, we use the term “anomaly localization” to span both segmentation and detection. Following Baugher and Rosenfeld (1986), localization refers to estimating the approximate location of an anomaly, while delineation involves precisely tracing its boundaries. Our use of “localization” thus covers a range of approaches depending on label availability and model design — from coarse localization to fine-grained segmentation.

This review aims to provide a comprehensive overview of the current state of the art, identify research gaps, and highlight open challenges. For a more detailed comparison of denoising diffusion models to GANs and VAEs, we refer to related work by Bercea et al. (2025b) and Friedrich et al. (2024a). We begin with a brief methodological background

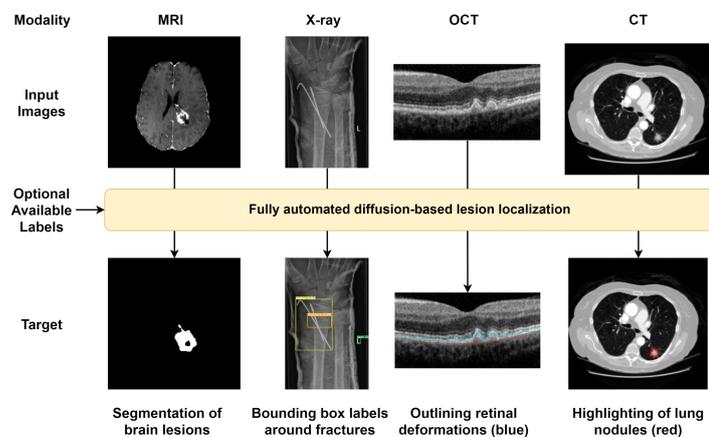


Figure 1: Anomaly localization is the task of outlining pathological regions in medical images. In this work, we present an overview of automated solutions based on denoising diffusion models.

on diffusion models, emphasizing their application in image reconstruction and conditioning through guidance mechanisms. We then present an overview of available datasets and evaluation metrics tailored for anomaly localization in medical images. Next, we explore various supervision schemes, ranging from fully supervised segmentation to semi-supervised, weakly supervised, self-supervised, and unsupervised methods, and provide insights into the effectiveness and limitations of diffusion models in these scenarios, as well as their generalizability to unknown or rare diseases. Finally, we address open challenges, such as detection bias, domain shift, clinical validation, computational cost, and model interpretability.

2. Background

In this section, we first provide a brief theoretical background on the basic framework of denoising diffusion models. In Section 2.2, we explain how diffusion models can be applied to reconstruction-based anomaly detection. Finally, in Section 2.3, we present the various guidance mechanisms necessary for conditioning diffusion models for medical anomaly localization.

2.1 Denoising Diffusion Models

Denoising diffusion models are a class of generative models inspired by thermodynamic processes (Sohl-Dickstein et al., 2015). The core idea is to gradually corrupt an input signal into pure noise using a forward diffusion process, then learn to reverse this process step by step. This learned reverse process transforms random noise into samples that resemble the original data distribution. As proposed in Ho et al. (2020) and Nichol and Dhariwal (2021a), the core idea of denoising diffusion models is that for many timesteps T , Gaussian noise is added to an input image x . This results in a series of noisy images x_0, x_1, \dots, x_T , where the noise level is steadily increased from 0 (no noise) to T (maximum noise). The forward noising process q with variances β_1, \dots, β_T is defined by

$$q(x_t|x_{t-1}) := \mathcal{N}(\sqrt{1 - \beta_t}x_{t-1}, \beta_t\mathbf{I}). \quad (1)$$

This recursion can be written explicitly as

$$x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, \quad \text{with } \epsilon \sim \mathcal{N}(0, \mathbf{I}), \quad (2)$$

$$\alpha_t := 1 - \beta_t \text{ and } \bar{\alpha}_t := \prod_{s=1}^t \alpha_s.$$

The mathematical principles, including appropriate choices of the hyperparameters T and β_t , are thoroughly derived by Ho et al. (2020), Song et al. (2020), and Dhariwal and Nichol (2021). The denoising process p_θ is learned by a deep learning model that predicts x_{t-1} from x_t for any step $t \in \{1, \dots, T\}$. It is given by

$$p_\theta(x_{t-1}|x_t) := \mathcal{N}(\mu_\theta(x_t, t), \Sigma_\theta(x_t, t)), \quad (3)$$

where μ_θ and Σ_θ are learned by a time-conditioned neural network following a U-Net architecture (Ronneberger et al., 2015). A memory-efficient U-Net implementation was presented by Bieder et al. (2024). In an alternative transformer-based architecture, Diffusion Transformers replace the commonly-used convolutional U-Net backbone with a transformer that operates on latent patches (Peebles and Xie, 2023). By using the reparametrization trick presented by Kingma et al. (2013), the output of the U-Net ϵ_θ is a noise prediction, and the MSE loss used for training simplifies to

$$\mathcal{L} := \|\epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, t)\|_2^2, \quad \text{with } \epsilon \sim \mathcal{N}(0, \mathbf{I}). \quad (4)$$

The mathematical principles, including appropriate choices of the hyperparameters T and β_t , are thoroughly derived by Ho et al. (2020), Song et al. (2020), and Dhariwal and Nichol (2021)

For image generation, we start from $x_T \sim \mathcal{N}(0, \mathbf{I})$ and iteratively go through the denoising process by predicting x_{t-1} for $t \in \{T, \dots, 1\}$. This sampling procedure of denoising diffusion models can be divided into denoising diffusion probabilistic models (DDPMs) and denoising diffusion implicit models (DDIMs). Training as described in Eq. (4) remains the same for both approaches.

DDPM sampling scheme: As shown in (Song et al., 2020), we use the DDPM formulation to predict x_{t-1} from x_t with

$$x_{t-1} = \sqrt{\bar{\alpha}_{t-1}} \left(\frac{x_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(x_t, t)}{\sqrt{\bar{\alpha}_t}} \right) + \sqrt{1 - \bar{\alpha}_{t-1} - \sigma_t^2} \epsilon_\theta(x_t, t) + \sigma_t \epsilon, \quad (5)$$

where we choose $\sigma_t = \sqrt{(1 - \bar{\alpha}_{t-1}) / (1 - \bar{\alpha}_t)} \sqrt{1 - \bar{\alpha}_t / \bar{\alpha}_{t-1}}$, and $\epsilon \sim \mathcal{N}(0, \mathbf{I})$. DDPMs thereby have a stochastic element ϵ in each sampling step.

DDIM sampling scheme: In denoising diffusion implicit models, we set $\sigma_t = 0$ in, Eq. (5), which results in a deterministic sampling process. Then, as derived in (Song et al., 2020), Eq. (5) can be viewed as the Euler method to solve an ordinary differential equation (ODE).

DDIM noise encoding: We can reverse the generation process by using the reversed ODE of the DDIM sampling scheme. Using enough discretization steps, we can encode

x_{t+1} given x_t with

$$x_{t+1} = x_t + \sqrt{\bar{\alpha}_{t+1}} \left[\left(\sqrt{\frac{1}{\bar{\alpha}_t}} - \sqrt{\frac{1}{\bar{\alpha}_{t+1}}} \right) x_t + \left(\sqrt{\frac{1}{\bar{\alpha}_{t+1}} - 1} - \sqrt{\frac{1}{\bar{\alpha}_t} - 1} \right) \epsilon_\theta(x_t, t) \right]. \quad (6)$$

We denote Eq. (6) as the DDIM noise encoding scheme. By applying Eq. (6) for $t \in \{0, \dots, T-1\}$, we can encode an image x_0 in a noisy image x_T . Then, we recover the identical x_0 from x_T by using Eq. (5) with $\sigma_t = 0$ for $t \in \{T, \dots, 1\}$.

Both the DDPM and DDIM processes are designed to generate synthetic data from random noise, i.e., it is a purely generative model. To adapt this model class to medical downstream task, conditioning mechanisms will be necessary to tailor the model output to the task at hand. This will be discussed in Section 2.3.

2.2 Reconstruction-based Anomaly Localization

We can leverage image-to-image translation to restore a pseudo-healthy image \hat{x}_H from an input image x_P potentially showing a pathology, as presented in Fig. 2. In this work, we consider diffusion-based approaches, where L steps of noise are added to an input image x_P , following Eq. (2) or Eq. (6) for the DDPM and DDIM schemes, respectively. The hyperparameter L indicates the level of noise added, as well as the number of denoising step taken during the denoising process. We emphasize the importance of this parameter L , as it directly governs how much the denoising process can alter the reconstruction, thereby affecting both sensitivity and specificity; its impact is systematically analyzed by Bercea et al. (2023a) and Wolleb et al. (2022a). A diffusion model is trained or guided to generate healthy samples. Various techniques discussed in Section 5 ensure that only pathological regions are altered while subject-specific anatomical structures are preserved. Finally, by taking the difference between an input image x_P and its pseudo-healthy reconstruction \hat{x}_H , a pixel-wise anomaly map $a = |x_P - \hat{x}_H|$ can be defined.

2.3 The Evolution of Guidance

While the denoising diffusion models described in Section 2.1 are designed for unconditional image generation, many medical downstream tasks require adherence to specific requirements, such as being conditioned on an input image or conforming to the characteristics of a particular disease. In addition to scalar conditioning demonstrated in (Nichol and Dhariwal, 2021b), more advanced conditioning mechanisms can be explored:

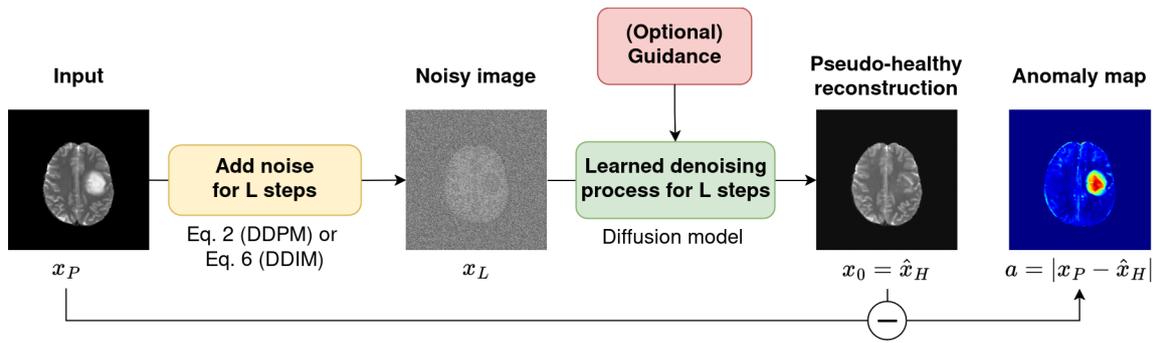


Figure 2: In reconstruction-based anomaly localization using diffusion models, we first add noise to an input image x_P . During denoising, we translate an input image to its pseudo-healthy counterpart \hat{x}_H . The pixel-wise anomaly map is given by the difference between input x_P and output \hat{x}_H .

- Conditioning through Concatenation:** Saharia et al. (2022) proposed *Palette*, where conditioning through concatenation involves appending the conditioning image (e.g., a grayscale or corrupted version) as extra input channels to the noisy image at each denoising step. This allows the model to consistently use spatial guidance from the conditioning input throughout the reverse diffusion process. The approach is simple, stable to train, and broadly applicable to paired image-to-image translation tasks like inpainting, super-resolution, and colorization.
- Gradient Guidance:** An external guidance scheme, as presented in (Dhariwal and Nichol, 2021), involves training a classification network C to distinguish between class labels. This network $C(x_t, t)$ shares the encoder architecture of the diffusion model and is trained on noisy input images x_t obtained via Eq. (2). During sampling, following the DDIM scheme, at each denoising step, ϵ_θ in Eq. (5) is modified to $\hat{\epsilon}_\theta = \epsilon_\theta(x_t, t) - s \times \sqrt{1 - \bar{\alpha}_t} \nabla_{x_t} \log C(c|x_t, t)$, with the hyperparameter s scaling the guidance. This steers image generation towards the desired class c . The diffusion model ϵ_θ and the classification model C are trained separately, offering the flexibility to combine gradients from different networks in a plug-and-play fashion without having to retrain the diffusion model (Wolleb et al., 2022b). However, training two separate networks can be unstable, and output quality depends on the performance of the classification network, introducing potential bias.
- Classifier-free guidance:** As proposed in (Ho and Salimans, 2022), guidance can be incorporated directly into the latent space of the diffusion model by training a class-conditional diffusion model $\epsilon_\theta(x_t, t, c)$. This model can omit class information by selecting $c = \emptyset$ with a certain probability, allowing it to train on both conditional and unconditional objectives by randomly dropping c during training. This method provides robust gradients and the efficiency of training a single model. However, it requires retraining the model for each new classification task.
- Implicit Guidance:** Guidance can be provided by the input image itself. Patch-based methods (Behrendt et al., 2024a) condition the model on input image patches while predicting masked areas. AutoDDPMs (Bercea et al., 2023a) further employ masking to remove highly probable anomalous tissues, stitching to integrate healthy parts with pseudo-normal complements, and resampling from their joint noised distribution. In THOR (Bercea et al., 2024a), the backward process is guided through intermediate anomaly maps to project pseudo-normal images closer to the inputs. Bernoulli diffusion models (Wolleb et al., 2024) follow the idea of binary latent diffusion models and introduce a masking algorithm based on anomaly scores that can directly be extracted from the model output.

A summary of these approaches is provided in Table 1.

3. Datasets

Diffusion-based models have been increasingly employed in the field of medical anomaly localization, leveraging various datasets encompassing multiple imaging modalities and anatomical regions. These datasets, summarized in Table 2, are critical for developing and validating models that can accurately identify anomalies. Much of the research on diffusion-based anomaly localization methods has centered around brain imaging, notably MR scans. This emphasis is primarily attributed to the abundance of well-annotated datasets and the clinical significance of anomaly localization. Datasets such as BraTS (Menze et al., 2014; Bakas et al., 2017, 2018), ATLAS (Liew et al., 2022), and WMH (Kuijff et al., 2019) are extensively utilized in this domain, offering multi-modal MRI data with pixel-wise annotations crucial for evaluating anomaly localization models. The expansion into different modalities and anatomies like chest X-rays, retinal optical coherence tomography (OCT) scans, and

Table 1: Comparison of guidance strategies in diffusion models.

Key Feature	Description	Examples	Advantages	Disadvantages
Guidance through Concatenation	The conditioning on the input image is performed through concatenation at every step of the denoising process.	Saharia et al. (2022); Amit et al. (2021); Durrer et al. (2023)	Stable to train; can be applied to any paired image-to-image translation task	Requires paired images
Gradient Guidance	A separate task-specific model is trained on the dataset. During the denoising process, the gradient of this model is used to update the denoising process, guiding it towards desired characteristics.	Dhariwal and Nichol (2021); Wolleb et al. (2022a)	The two models are independent and can be flexibly used in a plug-and-play manner	Sampling can be unstable and heavily relies on hyperparameters such as noise level and guidance scale
Classifier-free Guidance	The diffusion model is trained using an optional class label as input in every block of the U-Net. During sampling, the desired class label can be provided to guide generation.	Sanchez et al. (2022); Liu et al. (2024)	Only one model to train; stable gradients	The model needs to be retrained for each new task
Implicit Guidance	Guidance is provided by the input image itself through self-supervised training.	Bercea et al. (2023a); Wyatt et al. (2022); Behrendt et al. (2024a)	No external conditioning labels required	The chosen self-supervision scheme must be carefully adapted to the task

Table 2: Datasets used for diffusion-based medical anomaly detection, covering various anatomies, modalities, and anomaly types. The table also lists the size of each dataset, the type of annotations provided, and the methods applied.

Anatomy	Dataset	Size	Modality	Anomaly type	Annotation Type	Methods
Brain	CROMIS Wilson et al. (2018)	311 volumes	CT	Stroke	Pixel-wise	Pinaya et al. (2022a)
	KCH/CHRONIC Mah et al. (2020)	1704 volumes	CT	Stroke	Pixel-wise	Pinaya et al. (2022a) Behrendt et al. (2024a)
	BraTS Menze et al. (2014); Bakas et al. (2017, 2018)	1251 volumes	Multi-Modal MRI	Tumors	Pixel-wise	Marimont et al. (2023) Wolleb et al. (2022a) Wolleb et al. (2024)
	<i>Private Clinic</i>	22 volumes	T1w MRI	Tumors	Pixel-wise	Wyatt et al. (2022) Bercea et al. (2023a)
	ATLAS Liew et al. (2022)	655 volumes	Multi-Modal MRI	Stroke	Pixel-wise	Bercea et al. (2024a) Marimont et al. (2023)
	WMH Kuijf et al. (2019)	60 volumes	T1w + FLAIR MRI	White matter Hyper-intensities	Pixel-wise	Pinaya et al. (2022a)
	MSLub Lesjak et al. (2018)	30 volumes	Multi-Modal MRI	MS lesions	Pixel-wise	Behrendt et al. (2024a)
FastMRI+ Zbontar et al. (2018); Zhao et al. (2021)	1001 volumes	Multi-Modal MRI	30 anomaly types	Bounding Boxes	Pinaya et al. (2022a) Bercea et al. (2025b)	
Chest	Chexpert Irvin et al. (2019)	224316 scans	X-ray	14 anomaly types	Bounding boxes	Wolleb et al. (2022a)
Retina	OCT2017 Kermany et al. (2018)	30K healthy 3000 pathological	OCT	3 anomaly types	Image-wise	Wolleb et al. (2024)
Wrist	GRAZPED Nagy et al. (2022)	20327 scans	X-ray	8 anomaly types	Bounding boxes	Bercea et al. (2024a)

pediatric wrist X-rays indicates the growing versatility and potential of diffusion-based methods in the broader field of medical imaging.

However, a significant limitation remains: most datasets traditionally designed for supervised methods focus on a single disease, which is suboptimal for testing the broad detection capabilities of anomaly localization systems. For instance, BraTS is specifically geared towards brain tumors, while ATLAS focuses on stroke detection. These narrow focuses limit the ability to generalize findings across multiple conditions. Additionally, associated non-pathological

changes or unrelated diseases, such as mass effects from tumors or atrophy following strokes, are often not annotated. Unsupervised methods can detect these areas as deviations from the norm, identifying them as anomalies, adversely affecting quantitative performance even when the detection is correct.

A few datasets provide multiple labels for different diseases, allowing for a broader spectrum for evaluation. For example, the FastMRI+ (Pinaya et al., 2022a; Zhao et al., 2021) dataset includes annotations for 30 different types of anomalies such as tumors, lesions, edemas, en-

larged ventricles, resections, post-treatment changes, and more, all marked with bounding boxes. Similarly, the Chexpert (Irvin et al., 2019) dataset, primarily focused on thoracic anomalies, contains annotations for 14 different types of conditions, including pleural effusions, lung opacity, lesions, and cardiomegaly, offering a more comprehensive dataset for evaluation. Additionally, the GRAZPEDWRI-DX (GRAZPED) (Nagy et al., 2022) dataset challenges methods to simultaneously detect and classify various pediatric wrist injuries, including bone anomalies, bone lesions, foreign bodies, fractures, metallic artifacts, periosteal reactions, pronator signs, or soft tissue abnormalities. This diversity in labeling across multiple pathologies enhances the robustness and generalizability of diffusion-based anomaly localization models.

Recent work has begun to explore multimodal extensions to diffusion models in medical imaging. While most current datasets focus solely on image data, some studies (Zhang et al., 2024b; Dong et al., 2024) have incorporated text descriptions to guide generation or segmentation, marking early steps toward true multi-modal pipelines. In addition, Wang et al. (2024) proposed a diffusion-based approach for cross-modality segmentation, demonstrating the potential to leverage complementary information across MR sequences or imaging modalities. This shift toward multimodal learning could enhance model robustness and clinical relevance, making it a promising direction for future research.

In clinical practice, radiologists and domain experts usually specialize in specific organs or imaging modalities, such as neuroradiology or chest imaging. Furthermore, medical imaging datasets are often limited in size, making it challenging to train broadly generalized models effectively. Consequently, adopting a domain-specific focus can be advantageous, as it allows for the development of more accurate models tailored to particular clinical contexts, especially when data are scarce.

As shown in Table 2, the number of training samples used for diffusion models varies greatly across studies, reflecting differences in both the task and dataset availability. While some works, such as Wyatt et al. (2022), evaluated models on as few as 22 MR volumes, others, including Pinaya et al. (2022a) and Marimont et al. (2023), relied on datasets with over 1,000 volumes. Importantly, the stochastic nature of diffusion training through the progressive addition of noise (see Section 2.1) acts as an inherent and powerful data augmentation mechanism, which can partially compensate for limited data. However, given that medical imaging datasets are typically smaller than those used in general computer vision tasks, the risk of memorization becomes more pronounced. This can lead to overfitting, reduced generalization, and potential privacy concerns (Gu et al., 2023). Conditional Denoising Diffusion

Probabilistic Models (DDPMs) are significantly more prone to memorization than their unconditional counterparts, because conditioning (e.g., on image features, segmentations, or class labels) introduces informative labels that cause tighter clustering in the model’s latent space. This clustering reduces variance and increases the model’s tendency to reproduce training samples, which is especially problematic in medical imaging, where data scarcity and patient privacy are critical (Chen et al., 2024).

Table 3 and Table 4 in the Appendix provide links to publicly available datasets and to GitHub repositories containing the source code. Furthermore, we refer to the MONAI Generative framework (Pinaya et al., 2023), which includes a lot of these methods. While all of these datasets can be used to develop diffusion-based anomaly localization algorithms, proper comparison and evaluation is crucial to assess the performance and clinical value of the proposed approaches. In Section 4, we discuss different evaluation metrics for the task of anomaly detection and localization in medical images.

4. Evaluation Metrics for Anomaly Detection and Localization in Medical Images

Evaluating the performance of diffusion-based models in medical anomaly detection and localization involves various metrics depending on the available ground truth annotations.

Image-level Anomaly Detection For image or volume level, several metrics evaluate a model’s ability to classify anomaly presence or absence. The Area Under the Receiver Operating Characteristic Curve (AUROC) is popular due to its threshold-independence and single-value performance summary. However, AUROC may mislead in imbalanced datasets dominated by the majority class. The Area Under the Precision-Recall Curve (AUPRC) and Average Precision (AP) focus on precision-recall trade-offs and offer nuanced evaluations, particularly in imbalanced datasets where anomalies are scarce. The F1 Score, which combines precision and recall into a single metric by calculating their harmonic mean, is useful for balancing false positives and false negatives. However, it requires a specific threshold to be set, which can vary depending on the application, making it sensitive to the balance between precision and recall. These metrics do not provide spatial information about the detected anomalies but merely indicate whether an anomaly is present, which is a significant limitation when spatial accuracy is critical for clinical applications.

Pixel/voxel-wise Localization is measured with metrics such as pixel-wise AUROC, AUPRC, and the Dice Similarity Coefficient (DSC). The DSC measures the overlap

between the predicted and ground truth segmentations, providing a single value that balances false positives and false negatives. However, it requires setting a thresholded anomaly map to determine whether a pixel is classified as anomalous, which can significantly impact results. Selecting an appropriate threshold is challenging and can vary between datasets and applications. Moreover, these metrics might not be well-calibrated for detecting small anomalies, preferring over-segmentation, where the model identifies larger regions as anomalous than they truly are (Maier-Hein et al., 2024). Also, they might fail to capture clinically relevant information, such as total lesion load, lesion count, or detection rate.

Bounding-box Localization is commonly evaluated with the Intersection over Union (IoU). IoU measures the overlap between the predicted bounding box and the ground truth, with higher IoU indicating better localization. Precision and recall at various IoU thresholds offer a nuanced view of the performance across different levels. However, generating bounding boxes from continuous anomaly heatmaps produced by diffusion models is challenging. This process can introduce errors and reduce the accuracy of traditional IoU-based metrics. Bercea et al. (Bercea et al., 2024b) proposed alternative metrics to address these challenges, such as counting the amount of response inside and outside bounding boxes. However, further development is needed to create metrics that accurately capture clinically relevant localization features.

5. Types of Supervision for Diffusion-based Anomaly Localization

Diffusion models offer a versatile framework for anomaly detection, ranging from targeted pathology localization with supervised models to broad anomaly screening with unsupervised techniques, adapting to the diverse needs of healthcare applications, as shown in Fig. 3. Supervised diffusion models use pixel-wise labels provided by expert clinicians to detect certain pathologies, enabling highly accurate identification and localization of specific anomalies. In weakly-supervised settings, which use image-level labels, models transform known pathologies into pseudo-healthy images, allowing for precise localization of these specific anomalies as outlined in Section 2.2. Self-supervised models estimate classes of anomaly types through synthetic augmentations. These augmentations, such as coarse noise patterns, are removed to create pseudo-healthy variants, thereby simulating a broader range of anomaly types. However, these models may struggle with unknown or rare anomalies that fall outside the learned mappings. In these cases, unsupervised approaches that learn a normative distribution are effective

in identifying rare or unknown anomalies not present during training due to a lack of information on the expected anomaly distribution.

5.1 Fully Supervised Lesion Segmentation

Fully supervised lesion segmentation involves segmenting lesions in medical images using annotated training data, where each pixel is labeled as either belonging to the lesion or not. The training dataset consists of pairs (i, l) , with i being the image and l the corresponding label mask. Diffusion-based approaches use the generative denoising process to produce pixel-wise segmentation masks, leveraging the stochastic nature of diffusion models for output diversity. We want to highlight that fully supervised segmentation approaches directly predict the segmentation mask, and do not rely on the reconstruction pipeline described in Section 2.2.

The technical innovation is presented in Fig. 4, where the diffusion process is used to generate a segmentation mask. To condition this random generation process on the image to be segmented, we need to incorporate the anatomical information through concatenation. This conditioning technique was first presented in Palette (Saharia et al., 2022), and then adapted for image segmentation (Wolleb et al., 2022c; Amit et al., 2021). This setup enables paired image-to-image translation is represents a major milestone for any medical image-to-image translation task. For this, the technical details are given as follows: Starting from random noise l_T , the iterative denoising scheme in Eq. (5) is used to obtain a segmentation mask $\hat{l} = l_0$. To generate a label mask for a specific input image i , conditioning on the anatomical features of i is required. Various conditioning algorithms have been proposed for DDPMs, one of the first in *Palette* (Saharia et al., 2022). Given an input image $i \in \mathbb{R}^{c \times h \times w}$, where c is the channel dimension and h and w are the spatial dimensions, the label mask l is of dimension $s \times h \times w$, with s being the number of segmentation classes.

In MedSegDiff (Wu et al., 2024a), conditioning is achieved using attention modules combined with a feature frequency parser. MedSegDiff-V2 (Wu et al., 2024b) extends this by introducing a spectrum-space transformer to learn interactions between semantic features and diffusion noise in the frequency domain, outperforming previous diffusion-based methods. DermoSegDiff (Bozorgpour et al., 2023) leverages diffusion-based segmentation for skin lesions. To improve performance across diverse skin tones and malignancy conditions with limited data, Carrión and Norouzi (2023) propose a diffusion-based feature extractor followed by a multilayer perceptron for pixel-wise segmentation. Zhang et al. (2024b) introduces a text- and edge-guided diffusion model that generates diverse, anatomically meaningful synthetic medical images to significantly enhance segmenta-

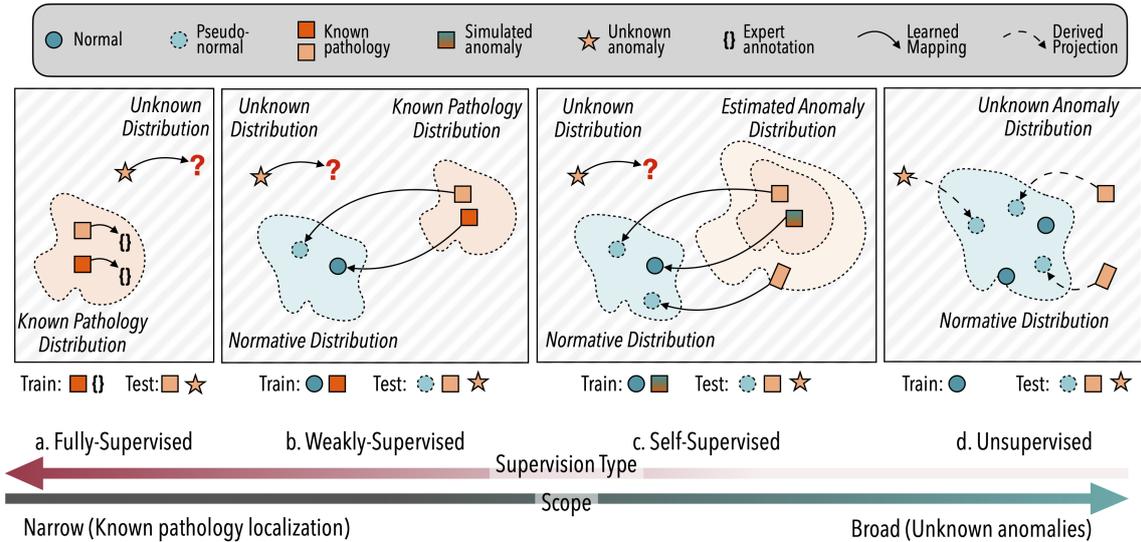


Figure 3: This figure demonstrates the impact of varying levels of supervision on anomaly detection, progressing from targeted pathology localization to the detection of broad, unknown anomalies. All symbols are explained in the grey top bar. In (a), fully supervised methods use expert annotation to directly predict the pathology segmentation masks. In (b), weakly-supervised methods learn to transform known pathologies (e.g., brain tumors) into pseudo-normal images. In (c), self-supervised methods simulate anomaly types using synthetic augmentations (e.g., coarse noise patterns) that are eliminated to produce pseudo-normal variants. In (a), (b), and (c) behavior outside the learned mappings, such as unknown or rare anomalies depicted as a star, is not defined. In (d), unsupervised methods do not train to learn specific mappings but instead estimate the normative distribution, enabling anomaly detection as deviations from this norm. Semi-supervised methods combine elements of both fully supervised and unsupervised approaches.

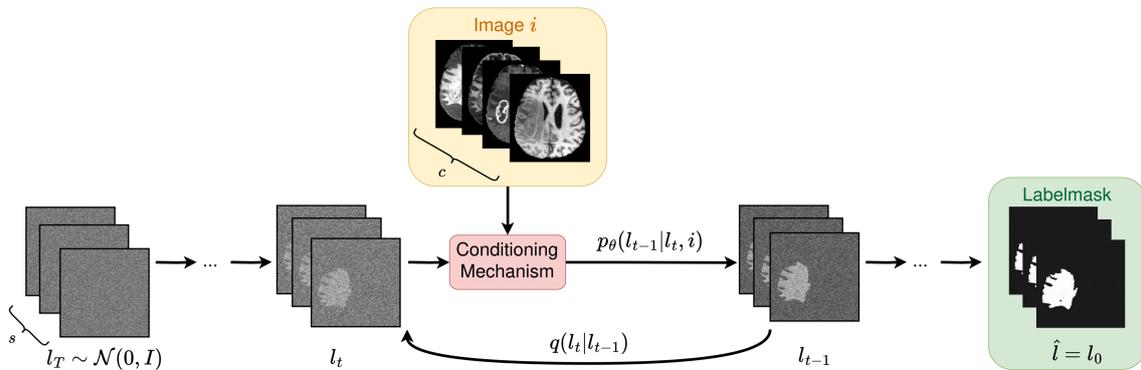


Figure 4: The denoising process of DDPMs is leveraged for the generation of a segmentation mask conditioned on the input image i . The conditioning mechanism relies on channel-wise concatenation in every step during the denoising process, as described in Saharia et al. (2022).

tion performance across multiple imaging modalities. Tan et al. (2023) introduces *DiffSS* a novel few-shot semantic segmentation paradigm by leveraging diffusion models to generate diverse, intra-class auxiliary support images conditioned on support masks, scribbles, or boundary maps, thereby enriching support diversity and significantly improving segmentation performance without altering the underlying segmentation network.

Ensembling A key advantage of DDPM-based segmentation is its probabilistic nature, allowing for the sampling of multiple segmentation masks \hat{l} for each input image i ,

resulting in an ensemble of segmentation masks $\{\hat{l}_1, \dots, \hat{l}_n\}$ (Amit et al., 2021; Wolleb et al., 2022c). The pixel-wise variance map of this ensemble highlights image areas where the model predictions are not consistent, serving for model interpretability. This model uncertainty aligns with the variability between human experts, as explored in (Rahman et al., 2023; Amit et al., 2023). The role of uncertainty in diffusion-based segmentation is further explored in Zbinden et al. (2023), where a categorical diffusion model generates multiple label maps accounting for aleatoric uncertainty, i.e., uncertainty originating from inherent randomness or

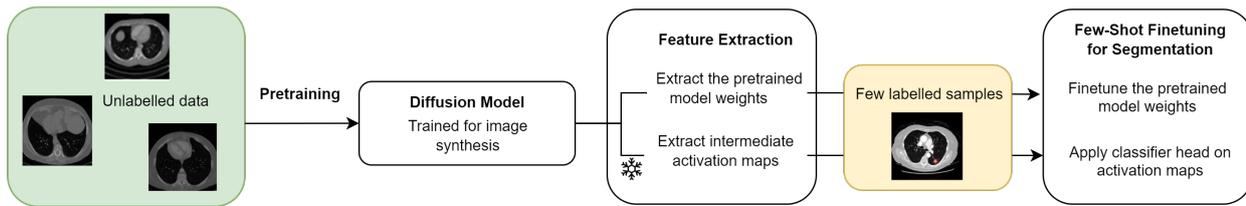


Figure 5: Diffusion models excel at distribution learning from unlabeled data. This facilitates effective few-shot fine-tuning on a limited number of labeled samples for lesion segmentation tasks.

noise in a dataset (Hüllermeier and Waegeman, 2021), from divergent ground truth annotations.

Going beyond Gaussian Noise Diffusion models are not limited to Gaussian noise. While Fig. 4 illustrates the generation process from Gaussian noise along Gaussian trajectories to binary one-hot-encoded masks, this can be improved by using categorical noise. BerDiff (Chen et al., 2023) proposes using Bernoulli noise as the diffusion kernel instead of Gaussian noise to enhance binary segmentation. Additionally, Zaman et al. (2023) suggests a cold diffusion approach for medical image segmentation, involving image perturbations through shifting and horizontal rotation of the segmentation surface. Yan et al. (2024) replace random noise with segmentation mask conditioning, enhance image features via frequency-domain contrast boosting, and improve focus on target regions using a conditional cross-attention mechanism.

Limitations While diffusion-based segmentation algorithms perform well in medical lesion segmentation, they suffer from long sampling times compared to classic U-Net approaches. LSegDiff (Vu Quoc et al., 2023) addresses high memory consumption and long sampling times by training a VAE combined with a latent diffusion model. Additionally, the acquisition and annotation of large, high-quality datasets with pixel-wise lesion annotations is a labor-intensive and time-consuming process. This limits the scope of the model to annotated lesions and introduces potential human bias. Additionally, as described by Wolleb et al. (2022c), there is a higher risk of diffusion-based segmentation to miss small anomalies by predicting empty segmentation masks compared to the classic nnU-Net approach (Isensee et al., 2021). Therefore, other supervision schemes have been explored, as described in the following sections.

5.2 Semi-Supervised Lesion Segmentation

In semi-supervised learning, algorithms utilize limited labeled data alongside abundant unlabeled data to enhance segmentation accuracy. Typically, pixel-wise labeled data outlining lesion boundaries is scarce and costly to obtain,

while large amounts of unlabeled data are often available. Leveraging diffusion models, known for effectively learning data representations, provides an opportunity to learn from this unlabeled data. DDPMs (Baranchuk et al., 2021; Yang and Wang, 2023), act as representation learners for discriminative computer vision tasks, as shown in Fig. 5. First, a diffusion model is trained for image synthesis on a large unannotated medical dataset to learn a representation of the target anatomy. Feature extraction can then be performed by either extracting the activation maps of the U-Net, or the pretrained model weights. In the final step, fine-tuning with a few pixel-wise labeled images can be done by training a classifier head to predict class labels for each pixel based on the extracted activation maps (Alshenoudy et al., 2023; Rosnati et al., 2023), or by fine-tuning the extracted model weights on the segmentation task (Rousseau et al., 2023). In Behrendt et al. (2024c), reconstruction-based unsupervised anomaly detection is combined with supervised segmentation on a small annotated dataset to improve segmentation performance for known anomalies and generalization to unknown pathologies. Ciampi et al. (2025) present a semi-supervised teacher-student framework that uses diffusion models (DDPMs) to generate segmentation pseudo-labels via unsupervised cycle-consistency pretraining, and iteratively refines them through multi-round co-training to train a student model alongside labeled data. Zhang et al. (2024a) introduce a class-conditional diffusion model that generates realistic synthetic image-label pairs by conditioning on segmentation masks, enhancing the diversity of training data.

5.3 Weakly Supervised Lesion Localization

Weak supervision involves using imprecise, noisy, or limitedly labeled data to train anomaly detection models. While each image has a label, the annotations lack the detailed information necessary for full supervision. In weakly supervised anomaly detection, labels are typically derived from image-level information indicating whether the image depicts a patient or a healthy control, as shown in Fig. 6. By utilizing two distinct datasets, \mathcal{H} and \mathcal{P} , the objective is to identify the visual manifestations that differentiate between them. Following an image-to-image translation

task, we aim to answer the question “How would a patient appear if pathology X was not present?” (Sanchez et al., 2022). A crucial technical milestone was the application of DDIM noise encoding of an input image, as presented in Section 2.2. According to the reconstruction-based method presented in Section 2.2, first an input image x_P is encoded to noise using Eq. (6) for L steps. Thereby, all anatomical information is stored in a noisy image x_L . This technical innovation enables a loss-free encoding of an input image in noise without destroying information, enabling unpaired image-to-image translation tasks (Wolleb et al., 2022a; Sanchez et al., 2022). Leveraging a data setup as presented in Fig. 6, weak labels can be used for guidance during denoising.

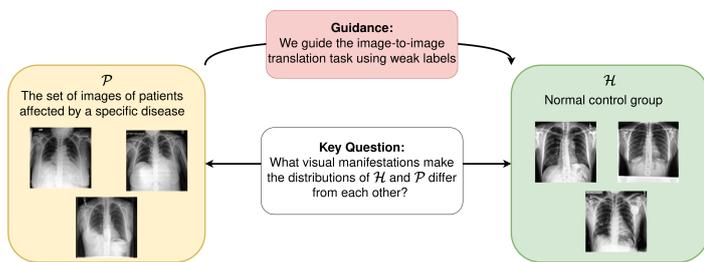


Figure 6: In weakly supervised anomaly detection scenarios, two different datasets are at hand. Dataset \mathcal{H} contains images of healthy controls, whereas dataset \mathcal{P} contains images of patients suffering from a specific disease. Using weakly supervised methods, the model learns the difference in distribution between \mathcal{H} and \mathcal{P} .

Gradient Guidance In Wolleb et al. (2022a), a classification network C is trained to distinguish between sets \mathcal{H} and \mathcal{P} . Gradient guidance towards the healthy class \mathcal{H} is applied in each denoising step, as described in Section 2.3. Hu et al. (2023) employs a similar approach but locates the desired class by approximating the derivative of the output of a class-conditional diffusion model with respect to the desired label \mathcal{H} , rather than using the gradient of an external classification model C . Another gradient-based approach is introduced in Fontanella et al. (2023), whereby a saliency map from an adversarial counterfactual attention module identifies pathological areas in medical images. During denoising, only these regions are altered using a masking and stitching algorithm.

Gradient-free Guidance As presented in Section 2.3, a class-conditional diffusion model $\epsilon(x_t, t, c)$ can be trained on the classes \mathcal{H} and \mathcal{P} . During denoising, the generation process is guided towards the healthy class by conditioning on $c = \mathcal{H}$ (Sanchez et al., 2022). Building on Ho and Salimans (2022), in Che et al. (2024), the forward process of diffusion models is employed for weakly supervised anomaly detection. Leveraging class activation maps (CAMs), Yoon

et al. (2024) enhance pseudo-mask quality by combining diffusion features with transformer-based CAMs using cross-attention, and enforces patch-level consistency to boost segmentation performance.

Limitations While weakly supervised methods are more flexible than fully supervised approaches in mimicking human experts’ outlining of a specific type of anomaly, they still have limitations in detecting diverse anomalies. Depending on the anomalies present in dataset \mathcal{P} , the model is trained to distinguish between the distributions of \mathcal{P} and \mathcal{H} , which may result in overlooking other types of anomalies in the input images. An example for the dataset bias is shown in Fig. 7. Following Wolleb et al. (2022a), a model is trained to distinguish between images with vs. without tumor. It thereby learns to remove the tumor from the image, but it does not correct from deformations caused by large tumors. On one hand, this renders the difference map more accurate, but also does not correct any anomalies caused by deformation of the brain structure. This specific problem is better handled by VAEs.

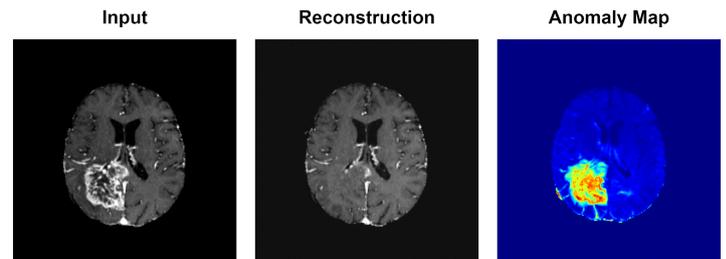


Figure 7: Illustration of dataset bias in tumor detection models. The model is trained to distinguish images with and without tumors, successfully removing tumor regions but failing to correct brain deformations caused by large tumors. This results in more accurate difference maps for tumor removal but leaves structural anomalies uncorrected.

5.4 Self-supervised Anomaly Localization

Self-supervised learning has revolutionized the pre-training of large neural networks, enabling models to learn robust representations from vast amounts of unlabeled data. The core idea behind self-supervised learning is to design auxiliary tasks that do not require manual annotations, allowing the network to learn useful features by solving these tasks. Examples include predicting the rotation angle of an image (Gidaris et al., 2018) or solving jigsaw puzzles (Noroozi and Favaro, 2016). Recent advancements have demonstrated the efficacy of self-supervised learning in enhancing feature extraction. Self-supervised models can achieve performance on par with, or even surpass, their supervised counterparts in certain tasks. Models like SimCLR (Chen

et al., 2020a), MoCo (He et al., 2020), and BYOL (Grill et al., 2020) leverage contrastive learning to maximize agreement between differently augmented views of the same data point, thereby learning powerful and discriminative features. Recently, it has been shown that diffusion models can be self-supervised representation learners, beneficial for many downstream tasks (Xiang et al., 2023; Yang and Wang, 2023).

In the context of medical imaging, self-supervised pre-training can significantly reduce the dependency on large labeled datasets, which are often difficult and expensive to obtain. By leveraging the inherent structure and properties of medical images, self-supervised methods can learn meaningful representations useful for downstream tasks such as segmentation and localization of lesions. As visualized in Fig. 8, self-supervision can be applied in two main ways: learning normative representations and enhancing feature learning, or simulating anomalies in the downstream task. Each approach offers unique advantages and limitations.

Learning Normative Representations Self-supervised learning can be used to learn representations of healthy data without requiring labeled anomalies. The goal is to simulate a set of augmentations $\bar{\mathcal{H}}$ of the normative representation and augment the normative training set. Here, the objective is to learn that $\bar{\mathcal{H}} \approx \mathcal{H}$, thereby capturing a varied distribution within \mathcal{H} . By doing so, the model becomes more adept at identifying deviations from the norm, indicative of anomalies.

Traditionally, this approach has been driven by context encoding VAEs for anomaly detection (Zimmerer et al., 2018). More recently, transformers that mask parts of the input during training, such as Masked Autoencoders (MAEs) (He et al., 2022) and Latent Transformer Models (LTMs) (Pinaya et al., 2021), have been used. Lately, this concept was also applied to diffusion-based models by noising only patches from an image and using the rest as context for the denoising process (Behrendt et al., 2024a).

Simulating Anomalies Another approach involves simulating anomalies as part of the self-supervised learning process. Here, the goal is to simulate pathologies $\bar{\mathcal{P}}$ and learn the transformation from $\bar{\mathcal{P}}$ to \mathcal{H} . Various methods have been developed to simulate anomalies. For instance, one technique involves simulating anomalies by interpolating foreign patches into images and detecting them (Tan et al., 2021). However, such strategies have not yet been adapted to diffusion models. Denoising autoencoders (DAEs) apply coarse noise to a U-Net to simulate anomalies (Kascenas et al., 2022).

AnoDDPM (Wyatt et al., 2022) extended this principle by employing diffusion models with *Simplex noise* to simulate structured, spatially coherent perturbations.

Unlike Gaussian noise, which acts independently at each pixel, Simplex noise produces continuous patterns across scales, yielding synthetic anomalies that more closely resemble lesion-like structures. A notable aspect of *AnoDDPM* is its hybrid nature. If a pathology resembles the noise distribution, the denoising process effectively suppresses it—corresponding to a self-supervised form of anomaly removal. However, if the pathology is covered by the noise during inference, the behaviour reduces to the standard unsupervised diffusion setting, where noise is reversed to match the healthy prior patterns.

Limitations Despite their potential, self-supervised anomaly-simulation methods such as *AnoDDPM* face inherent challenges. Their effectiveness depends on how well the simulated anomalies $\bar{\mathcal{P}}$ approximate the true pathological variability \mathcal{P} . Since \mathcal{P} is highly heterogeneous and encompasses many rare diseases¹, achieving broad anomaly coverage in practice is difficult. Anomalies that diverge from the e.g., Simplex noise distribution may remain undetected or only partially suppressed, as the transformation between real pathologies and their pseudo-healthy reconstructions is not explicitly learned. This limitation is evident in Fig. 9, where DDPMs with Simplex noise only partially remove stroke lesions and fails to restore enlarged ventricles in brain MRI, or suppresses fracture patterns in bone X-ray but reconstructs them with anatomically-healthy implausible, bent structures. Similar limitations have also been reported by Bercea et al. (2024a, 2023d), underscoring the challenge of generalizing beyond the simulated anomaly space.

5.5 Unsupervised Lesion Localization

Unsupervised anomaly localization methods have emerged as a powerful approach in medical imaging. These methods focus on learning the distribution of normal anatomy (\mathcal{H}) without any supervision on the expected anomaly distribution, although weak labels of only controls are required to curate a healthy-only dataset. Unlike (self-)supervised methods, unsupervised methods do not aim to learn mappings from pathological sets (\mathcal{P}) to normal sets (\mathcal{H}). Instead, they concentrate solely on modeling the distribution of normal, healthy data, as shown in Fig. 3. Deviations from this learned distribution are subsequently classified as anomalies, where pathological images are compared against pseudo-healthy reconstructions (see Fig. 2).

Variational autoencoders (VAEs) (Zimmerer et al., 2019; Chen et al., 2020b) and generative adversarial networks (GANs) (Schlegl et al., 2017) are commonly used to model the distribution of data of healthy controls. These models are trained to reconstruct healthy images accurately, and any significant deviation in the reconstruction is flagged as

1. <https://rarediseases.info.nih.gov>

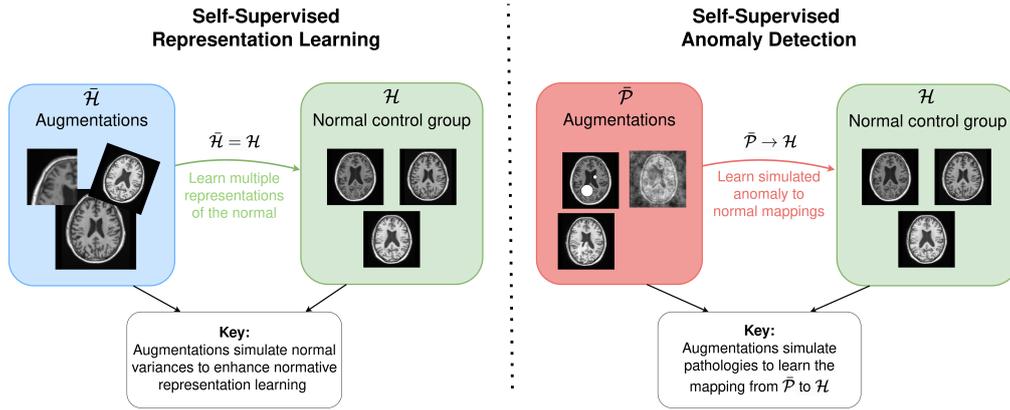


Figure 8: Illustration of self-supervised learning approaches for anomaly detection. Left: Augmentations simulate normal variances to enhance normative representation learning ($\bar{\mathcal{H}} \approx \mathcal{H}$). Right: Augmentations simulate pathologies to learn pathology-to-normal mappings ($\bar{\mathcal{P}} \rightarrow \mathcal{H}$).

a potential anomaly.

Recently, diffusion models (Ho et al., 2020; Wyatt et al., 2022; Liang et al., 2023) have been applied to this task. An innovative major algorithm was the use of latent diffusion models for unsupervised anomaly localization (Pinaya et al., 2021). In a first step, an autoencoder based on a VQ-VAE architecture (Van Den Oord et al., 2017) is trained to encode a healthy image into a low-dimensional latent space. In a second step, a denoising diffusion model is trained in the latent space. During inference, we can evaluate the algorithm on pathological images. The image is passed through the encoder into a latent representation z . As this representation will encode unseen changes, the denoising process of the diffusion model, trained for reconstruction of healthy representations, is applied for L steps, to correct for these anomalous changes. The resulting latent vector \hat{z} is then passed through the decoder to reconstruct the pseudo-healthy image. However, a main limitation is the loss of healthy tissue information during the reverse process due to the exhaustive noise needed to cover anomalies, referred to as the noise paradox in Bercea et al. (2023a). This challenge is illustrated in Fig. 9, where too little corruption leads to reconstructed anomalies being missed, while excessive corruption removes anomalies but produces false positives. To mitigate this, methods that guide the synthesis to mostly replace only the assumed pathological tissues, using the rest of the healthy context as guidance, have been developed (as discussed in Section 2.3 "Implicit Guidance"). Behrendt et al. (2025) condition the diffusion model’s denoising process using latent representations of the input image, enabling accurate reconstruction of healthy brain tissue while avoiding reconstruction of anomalies.

Limitations Despite their potential to detect arbitrary (rare) anomalies without relying on expectations about anomaly distributions, some challenges still remain. En-

suring that the learned distribution of normal anatomy is sufficiently comprehensive to delineate normal variations from subtle anomalies or even detect early alterations of tissues remains a significant challenge. The diversity and complexity of normal anatomical variations require sophisticated modeling techniques to accurately capture these nuances. Moreover, anomaly map computations often rely on pixel-wise differences, which are not ideal for detecting subtle pathological changes due to small intensity differences. Only a few works investigate other ways of computing anomaly maps using perceptual maps (Bercea et al., 2024b) or structural similarity (Behrendt et al., 2024b). Evaluating unsupervised models across diverse datasets and anomaly types remains crucial to ensure their robustness and generalizability.

6. Open Challenges

While the different types of supervision discussed in Section 5 provide opportunities to train diffusion models for anomaly localization under different data and label availability scenarios, several challenges remain. These include detection bias, high memory requirements for large 3D volumes, distribution shifts in multicentric data, the lack of comprehensive benchmark datasets for clinical validation, computational cost, and model interpretability. In the following subsections, we address these issues in detail.

6.1 Detection Bias

Traditionally, anomaly detection methods have been applied to finding multiple sclerosis (MS) lesions or tumors on Fluid-Attenuated Inversion Recovery (FLAIR) brain MR sequences. FLAIR imaging is particularly useful in diagnosing and monitoring conditions like MS, where it excels at revealing hyperintense lesions that indicate areas of demyelination or inflammation. Similarly, it is valuable in identifying other

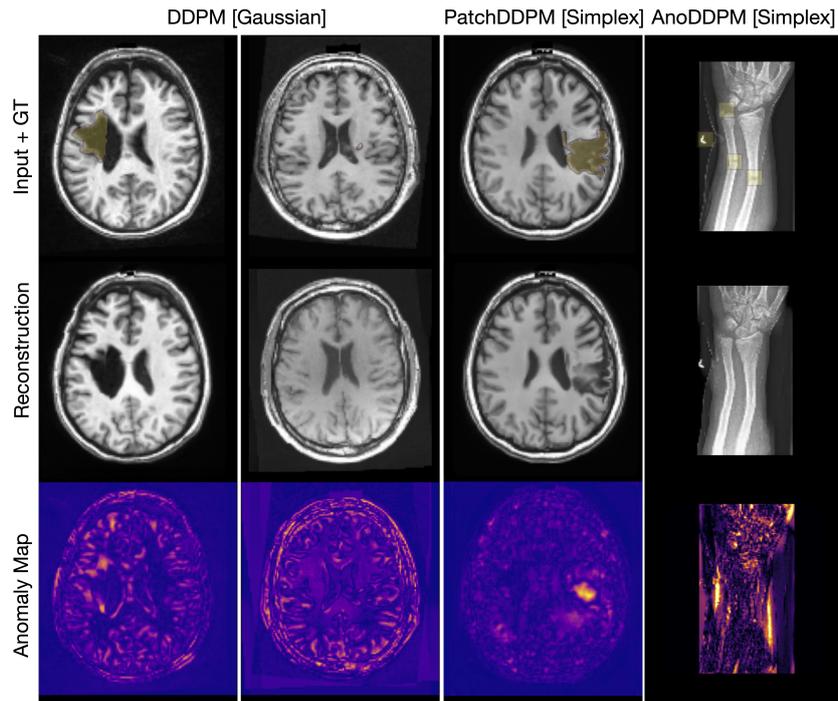


Figure 9: Limitations of diffusion-based anomaly detection in brain and bone MRI, shown from left to right. With Gaussian noise at $L = 250$ steps, a stroke lesion next to the lateral ventricle is reconstructed instead of removed, leading to a missed detection. With the same setting, a very small stroke lesion is removed, but the image contains numerous false positives that obscure the result. With PatchDDPM and Simplex noise at 250 steps, only parts of the stroke lesion are removed and the enlarged ventricle is not restored to a normal shape. With AnoDDPM and Simplex noise on bone MRI, noise-like structures such as the fracture are suppressed, yet the reconstruction produces unnaturally bent bones instead of a realistic healthy configuration.

types of brain abnormalities, including tumors and infarctions. This led to an interesting trend where methods that produced blurry reconstructions were found to be more proficient in anomaly localization (Bercea et al., 2023c). This effect was attributed to the simplicity of the task, which was later shown to be more effectively addressed using simple intensity thresholding techniques (Meissen et al., 2021). Two key lessons can be extracted from this:

1. Anomaly localization should not be evaluated in isolation. Evaluations often only report metrics such as mean absolute error (MAE) or training loss on healthy samples. A more comprehensive approach involves evaluating normative representation learning alongside anomaly localization scores, highlighting the importance of a holistic evaluation framework that provides deeper insights and more meaningful metrics (Bercea et al., 2025b).
2. Diverse and Comprehensive Datasets: Evaluated datasets should ideally contain a varied array of anomalies, including intensity-based anomalies (e.g., lesions, tumors, inflammation) and structural anomalies (e.g., atrophy, fractures, mass effects). This diversity is crucial to obtaining a comprehensive view of the detection capabilities and limitations of the methods.

Furthermore, simple baselines and classical literature on anomaly localization should not be overlooked in favor of more advanced generative modeling techniques. Methods should avoid using the vague term "anomaly detection" without clarifying the supervision type, scope, and limitations of their approach. Addressing these challenges necessitates the development of more comprehensive datasets and evaluation metrics that reflect the complexity and variety of medical anomalies. This will enable a more accurate assessment of the true performance and applicability of anomaly detection and localization methods in clinical settings.

6.2 2D/3D Anomaly Localization

While almost all approaches presented in Section 5 are implemented in 2D, when dealing with MR and CT scans, 3D approaches will be required. A current challenge are the high memory requirements and long sampling times when implementing 3D diffusion models. To address these issues, Bieder et al. (2023) proposed a memory-efficient 3D architecture, enabling a fully supervised lesion segmentation on a resolution of 256^3 by training only on patches. This reduces memory consumption during training from 78.5 GB to 10.6 GB, while maintaining a duration of 1 second

for one network evaluation. Durrer et al. (2024) further explored architectures for 3D volumes. Running inference on a 256^3 volume by stacking 2D slices from a 2D model takes around 20 minutes, with memory usage of 34.0 GB. A full 3D approach requires 78.8 GB with roughly the same sampling time. Friedrich et al. (2024b) demonstrated that 3D wavelet diffusion models can significantly reduce this to 5 minutes per volume with only 32.3 GB of memory. Another common workaround is the implementation of 3D latent diffusion models (Pinaya et al., 2022b; Khader et al., 2023), which encode the input data into a compressed latent space before applying the diffusion model. While this approach can be used for anomaly localization (Pinaya et al., 2022a; Graham et al., 2023), model performance is still limited by the lack of a well-performing 3D autoencoder (Friedrich et al., 2024b; Durrer et al., 2024).

Apart from downsampling the 3D volumes, which comes with a loss of information, applying the 2D models slice-wise brings the challenge of missing consistency between the output slices (Durrer et al., 2024). This issue is addressed in Zhu et al. (2023), where pseudo-3D volumes are generated with an additional 1D convolution into the third spatial dimension, ensuring consistent stacks of 2D slices. Friedrich et al. (2024b) proposed to apply a discrete wavelet transform to reduce the spatial dimension before applying the diffusion model, enabling processing of volumes of a resolution up to 256^3 . Applications of such architectures on an anomaly localization task still remains to be explored.

6.3 Non-pathological Distribution Shifts

A prevalent yet largely unaddressed issue in the development of anomaly localization methods is the occurrence of distributional shifts. Typically, AD methods are trained on a dataset of healthy controls (i.e., show no pathology) and subsequently evaluated on various downstream tasks involving different pathologies. However, these two datasets often originate from different hospitals or scanners, resulting in a notable disparity between the training and evaluation distributions. For instance, routine checks are not commonly conducted in cancer-specific clinics. This discrepancy and other biases can significantly impact anomaly localization performance (Bercea et al., 2023b; Meissen et al., 2024).

Some strategies to mitigate this issue involve incorporating slices without clear pathology from a pathological dataset as healthy slices in the training set. However, this approach has drawbacks. These slices may only represent partial brain volumes and might not be entirely healthy, as the non-pathological effects of surrounding lesions are not annotated. Moreover, this practice can lead to data leakage if the same patients are used for both training and evaluation, potentially skewing the results and providing an inaccurate assessment of the effectiveness of the methods.

Exploring techniques to adapt to varying distributions during inference is crucial for achieving clinical acceptance and warrants further investigation.

6.4 Clinical Integration

To compare performance on benchmark datasets, Bercea et al. (2025b) provide a comprehensive evaluation of generative AI models, including diffusion models, for detecting and correcting anomalies in brain MRI scans. From a clinical perspective, anomaly localization can add value at multiple points in the radiology workflow: (i) case triage and prioritization for suspected urgent findings; (ii) quality control (e.g., motion, missing sequences, laterality mismatches); (iii) decision support within the reporting viewport via anomaly maps and pseudo-healthy reconstructions; (iv) longitudinal change assessment; and (v) safety-net alerts for incidental or rare abnormalities. To be useful, outputs must be calibrated, time-efficient, and straightforward for radiologists to verify.

Key hurdles remain. Beyond retrospective reader studies, prospective, multi-center evaluations are needed to assess effects on reporting time, detection of subtle/rare pathologies, and downstream management. Workflow fit (PACS/RIS integration, turnaround time, human-in-the-loop acceptance/ override), robustness across scanners and protocols, transparency (uncertainty and case-level rationale), and governance (audit trails, versioning, data protection) are prerequisites for regulatory approval and safe deployment.

Existing efforts provide partial progress but are not substitutes for clinical validation. The MOOD dataset (Zimmerer et al., 2022) evaluates detection of synthetic anomalies on a hidden test set, and the NOVA benchmark (Bercea et al., 2025a) offers evaluation-only data with heterogeneous real cases to probe localization and clinical reasoning. Retrospective analyses exist (Finck et al., 2021; He et al., 2024), and broader method evaluations (including diffusion models) have been reported (Bercea et al., 2025b). What is still lacking are prospective impact studies with predefined clinical endpoints (e.g., reporting time, recall rates, time-to-treatment), standardized reporting checklists, and reference integrations (e.g., open PACS/RIS plug-ins) to lower deployment friction. Until such gaps are addressed, adoption will likely remain limited to research settings or controlled pilots.

6.5 Computational Considerations

A significant drawback of denoising diffusion models are the long sampling times due to the iterative generation process, as well as the high memory requirements, mainly due to global attention layers in the U-Net architecture. A large research field has opened to speed up diffusion-based

image generation. A first step in this direction was already presented in Song et al. (2020). As discussed in Section 2.1, the DDIM denoising schedule can be interpreted as a numerical solver of an ordinary differential equation (ODE). By skipping timesteps during denoising, the step size is increased at the cost of numerical accuracy and image quality (Song et al., 2020). Therefore, the design space of appropriate ODE solvers was explored, suggesting Heun’s method as a balanced choice between sampling speed and image quality (Karras et al., 2022; Liu et al., 2022a). Furthermore, faster sampling can be achieved via distillation of the sampling procedure (Salimans and Ho, 2022; Meng et al., 2023). Another approach is to combine diffusion models with an adversarial component (Xiao et al., 2021), or diffusion model sampling can be combined with neural operators (Zheng et al., 2023) for one-step image generation. However, the application of these approaches for anomaly localization tasks still remains to be explored. Seyfarth et al. (2024) reveals that training emissions in 2D latent diffusion models remain relatively stable across image sizes, ranging from 0.40 ± 0.16 kg CO_2 to 0.79 ± 0.20 kg CO_2 , while 3D models scale sharply with resolution, reaching up to 9.0 ± 0.9 kg CO_2 , equivalent to roughly 92 km of driving. However, synthesis is even more carbon-intensive: generating 10,000 2D samples emits up to 18.2 ± 0.3 kg CO_2 , whereas high-resolution 3D synthesis can emit as much as 362.3 ± 1.0 kg CO_2 , making it the dominant contributor to environmental impact.

6.6 Interpretability

Accurate anomaly localization accelerates the diagnostic process and effectively highlights regions of interest. However, interpretability remains a crucial aspect of anomaly localization approaches in medical imaging, particularly for clinical applications. While techniques discussed in Section 5.1 offer the possibility of estimating model uncertainty pixel-wise, they do not inherently assess the severity or urgency of the findings. Recent advances in large language models (LLMs) (Hua et al., 2024; Zhu et al., 2024) offer a promising direction for enhancing the interpretability of unsupervised anomaly localization methods. Li et al. (2024) applied visual question answering models to anomaly detection tasks and demonstrated that LLMs can enhance the interpretability of detected anomalies. Moreover, they showed that anomaly maps used as inputs for LLMs assist them in generalizing to describe unseen anomalies. Nevertheless, more research is needed to fully exploit the potential of these advancements and ensure their effectiveness in clinical practice.

7. Conclusion

In conclusion, our exploration of anomaly localization in medical images using diffusion models underscores the nuanced nature of the field. Acknowledging that not all approaches are universally effective in all scenarios is crucial. These scenarios encompass varying quantities of available data and corresponding labels, different imaging modalities, and anomaly types. To address these challenges, it is critical to carefully define and tailor evaluation metrics to the specific characteristics of each anomaly type and modality. In addition, exploring alternatives to Gaussian noise and investigating test-time adaptation techniques can help mitigate domain shifts and improve model robustness.

Moreover, the majority of approaches are currently implemented in 2D, highlighting the need for further exploration and adaptation in 3D settings to better capture volumetric anomalies. To this end, fast and memory-efficient diffusion models need to be explored. Recent advances such as flow matching (Lipman et al., 2022; Liu et al., 2022b), shortcut models (Frans et al., 2024), and improved score matching techniques (Song and Ermon, 2019; Vahdat et al., 2021) offer promising new directions for the development of diffusion-based generative models. While these methods have been designed primarily for image generation, their potential for downstream applications such as anomaly localization remains largely unexplored, with only very recent work on flow-based, pathology-aware image synthesis (Susladkar et al., 2025). In this work, in Section 5, we presented various conditioning strategies tailored to different types of denoising diffusion probabilistic or implicit models. However, such techniques will need to be carefully reviewed and adapted to align with the architectural and training differences introduced by these emerging generative paradigms.

Nevertheless, anomaly localization remains a critical challenge in the medical imaging field. Diffusion models offer a promising field of research, particularly through their ability to synthesize high-quality pseudo-healthy images. This capability opens new possibilities for clinical applications, with the potential to enhance diagnostic accuracy and streamline patient care pathways. However, to fully realize their clinical utility, further clinical validation of potential use cases of diffusion-based anomaly localization methods is needed to integrate them into routine medical practice.

Acknowledgments

C.I.B. is funded via the EVUK programme (“Next-generation AI for Integrated Diagnostics”) of the Free State of Bavaria and in part supported by Berdelle-Stiftung (grant Time-Flow). Additionally, for the completion of this project C.I.B., was partially supported by the Helmholtz Association under

the joint research school 'Munich School for Data Science - MUDS.' J.W. is funded by the Swiss National Science Foundation (Grant No. P500PT222349), and was supported by Novartis FreeNovation, and the Uniscientia Foundation.

Ethical Standards

The work follows appropriate ethical standards in conducting research and writing the manuscript, following all applicable laws and regulations regarding treatment of animals or human subjects.

Conflicts of Interest

We declare we do not have conflicts of interest.

Data availability

Datasets discussed in this review are available under the citations and links provided in Table 2 and Table 3.

References

- Ahmed Alshenoudy, Bertram Sabrowsky-Hirsch, Stefan Thumfart, Michael Giretzlehner, and Erich Kobler. Semi-supervised brain tumor segmentation using diffusion models. In *IFIP International Conference on Artificial Intelligence Applications and Innovations*, pages 314–325. Springer, 2023.
- Tomer Amit, Tal Shaharbany, Eliya Nachmani, and Lior Wolf. Segdiff: Image segmentation with diffusion probabilistic models. *arXiv preprint arXiv:2112.00390*, 2021.
- Tomer Amit, Shmuel Shichrur, Tal Shaharabany, and Lior Wolf. Annotator consensus prediction for medical image segmentation with diffusion models. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 544–554. Springer, 2023.
- Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin S Kirby, John B Freymann, Keyvan Farahani, and Christos Davatzikos. Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features. *Scientific data*, 4(1):1–13, 2017.
- Spyridon Bakas, Mauricio Reyes, Andras Jakab, Stefan Bauer, Markus Rempfler, Alessandro Crimi, Russell Takeshi Shinohara, Christoph Berger, Sung Min Ha, Martin Rozycki, et al. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the BRATS challenge. *arXiv preprint arXiv:1811.02629*, 2018.
- Dmitry Baranchuk, Ivan Rubachev, Andrey Voynov, Valentin Khruikov, and Artem Babenko. Label-efficient semantic segmentation with diffusion models. *arXiv preprint arXiv:2112.03126*, 2021.
- Ernest S Baugher and Azriel Rosenfeld. Boundary localization in an image pyramid. *Pattern Recognition*, 19(5): 373–395, 1986.
- Finn Behrendt, Debayan Bhattacharya, Julia Krüger, Roland Opfer, and Alexander Schlaefer. Patched diffusion models for unsupervised anomaly detection in brain mri. In *Medical Imaging with Deep Learning*, pages 1019–1032. PMLR, 2024a.
- Finn Behrendt, Debayan Bhattacharya, Lennart Maack, Julia Krüger, Roland Opfer, Robin Mieling, and Alexander Schlaefer. Diffusion models with ensembled structure-based anomaly scoring for unsupervised anomaly detection. *arXiv preprint arXiv:2403.14262*, 2024b.
- Finn Behrendt, Debayan Bhattacharya, Lennart Maack, Julia Krüger, Roland Opfer, and Alexander Schlaefer. Combining reconstruction-based unsupervised anomaly detection with supervised segmentation for brain mris. In *Medical Imaging with Deep Learning*, 2024c.
- Finn Behrendt, Debayan Bhattacharya, Robin Mieling, Lennart Maack, Julia Krüger, Roland Opfer, and Alexander Schlaefer. Guided reconstruction with conditioned diffusion models for unsupervised anomaly detection in brain mris. *Computers in Biology and Medicine*, 186: 109660, 2025.
- Cosmin I Bercea, Michael Neumayr, Daniel Rueckert, and Julia A Schnabel. Mask, stitch, and re-sample: Enhancing robustness and generalizability in anomaly detection through automatic diffusion models. *arXiv preprint arXiv:2305.19643*, 2023a.
- Cosmin I Bercea, Esther Puyol-Antón, Benedikt Wiestler, Daniel Rueckert, Julia A Schnabel, and Andrew P King. Bias in unsupervised anomaly detection in brain mri. In *Workshop on Clinical Image-Based Procedures*, pages 122–131. Springer, 2023b.
- Cosmin I Bercea, Daniel Rueckert, and Julia A Schnabel. What do aes learn? challenging common assumptions in unsupervised anomaly detection. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 304–314. Springer, 2023c.

- Cosmin I Bercea, Benedikt Wiestler, Daniel Rueckert, and Julia A Schnabel. Reversing the abnormal: Pseudo-healthy generative networks for anomaly detection. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 293–303. Springer, 2023d.
- Cosmin I Bercea, Benedikt Wiestler, Daniel Rueckert, and Julia A Schnabel. Diffusion models with implicit guidance for medical anomaly detection. arXiv preprint arXiv:2403.08464, 2024a.
- Cosmin I Bercea, Benedikt Wiestler, Daniel Rueckert, and Julia A Schnabel. Generalizing unsupervised anomaly detection: Towards unbiased pathology screening. In Medical Imaging with Deep Learning, pages 39–52. PMLR, 2024b.
- Cosmin I Bercea, Jun Li, Philipp Raffler, Evamaria O Riedel, Lena Schmitzer, Angela Kurz, Felix Bitzer, Paula Roßmüller, Julian Canisius, Mirjam L Beyrle, et al. Nova: A benchmark for anomaly localization and clinical reasoning in brain mri. arXiv preprint arXiv:2505.14064, 2025a.
- Cosmin I Bercea, Benedikt Wiestler, Daniel Rueckert, and Julia A Schnabel. Evaluating normative representation learning in generative ai for robust anomaly detection in brain imaging. Nature Communications, 16(1):1624, 2025b.
- Florentin Bieder, Julia Wolleb, Alicia Durrer, Robin Sandkuehler, and Philippe C Cattin. Memory-efficient 3d denoising diffusion models for medical image processing. In Medical Imaging with Deep Learning, 2023.
- Florentin Bieder, Julia Wolleb, Alicia Durrer, Robin Sandkuehler, and Philippe C Cattin. Memory-efficient 3d denoising diffusion models for medical image processing. In Medical Imaging with Deep Learning, pages 552–567. PMLR, 2024.
- Afshin Bozorgpour, Yousef Sadegheih, Amirhossein Kazerooni, Reza Azad, and Dorit Merhof. Dermosegdiff: A boundary-aware segmentation diffusion model for skin lesion delineation. In Predictive Intelligence in Medicine, pages 146–158. Springer Nature Switzerland, 2023.
- Héctor Carrión and Narges Norouzi. Fedd-fair, efficient, and diverse diffusion-based lesion segmentation and malignancy classification. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 270–279. Springer, 2023.
- Yiming Che, Fazle Rafsani, Jay Shah, Md Mahfuzur Rahman Siddiquee, and Teresa Wu. Anofpdm: Anomaly segmentation with forward process of diffusion models for brain mri. arXiv preprint arXiv:2404.15683, 2024.
- Tao Chen, Chenhui Wang, and Hongming Shan. Berdiff: Conditional bernoulli diffusion model for medical image segmentation. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 491–501. Springer, 2023.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In International conference on machine learning, pages 1597–1607. PMLR, 2020a.
- Xiaoran Chen, Suhang You, Kerem Can Tezcan, and Ender Konukoglu. Unsupervised lesion detection via image restoration with a normative prior. Medical image analysis, 64:101713, 2020b.
- Xuxin Chen, Ximin Wang, Ke Zhang, Kar-Ming Fung, Theresa C Thai, Kathleen Moore, Robert S Mannel, Hong Liu, Bin Zheng, and Yuchen Qiu. Recent advances and clinical applications of deep learning in medical image analysis. Medical Image Analysis, 79:102444, 2022.
- Yunhao Chen, Xingjun Ma, Difan Zou, and Yu-Gang Jiang. Towards a theoretical understanding of memorization in diffusion models. 2024.
- Luca Ciampi, Gabriele Lagani, Giuseppe Amato, and Fabrizio Falchi. Semi-supervised biomedical image segmentation via diffusion models and teacher-student co-training. arXiv preprint arXiv:2504.01547, 2025.
- Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. Advances in neural information processing systems, 34:8780–8794, 2021.
- Zhiwei Dong, Genji Yuan, Zhen Hua, and Jinjiang Li. Diffusion model-based text-guided enhancement network for medical image segmentation. Expert Systems with Applications, 249:123549, 2024.
- Alicia Durrer, Julia Wolleb, Florentin Bieder, Tim Sinnecker, Matthias Weigel, Robin Sandkühler, Cristina Granziera, Özgür Yaldizli, and Philippe C Cattin. Diffusion models for contrast harmonization of magnetic resonance images. arXiv preprint arXiv:2303.08189, 2023.
- Alicia Durrer, Julia Wolleb, Florentin Bieder, Paul Friedrich, Lester Melie-Garcia, Mario Alberto Ocampo Pineda, Cosmin I Bercea, Ibrahim Ethem Hamamci, Benedikt Wiestler, Marie Piraud, et al. Denoising diffusion models for 3d healthy brain tissue inpainting. In MICCAI Workshop on Deep Generative Models, pages 87–97. Springer, 2024.

- Lilian Edwards. The eu ai act: a summary of its significance and scope. Artificial Intelligence (the EU AI Act), 1:25, 2021.
- Tom Finck, David Schinz, Lioba Grundl, Rami Eisawy, Mehmet Yigitsoy, Julia Moosbauer, Franz Pfister, and Benedikt Wiestler. Automated pathology detection and patient triage in routinely acquired head computed tomography scans. Investigative Radiology, 56(9):571–578, 2021.
- Alessandro Fontanella, Grant Mair, Joanna Wardlaw, Emanuele Trucco, and Amos Storkey. Diffusion models for counterfactual generation and anomaly detection in brain images. arXiv preprint arXiv:2308.02062, 2023.
- Kevin Frans, Danijar Hafner, Sergey Levine, and Pieter Abbeel. One step diffusion via shortcut models. arXiv preprint arXiv:2410.12557, 2024.
- Paul Friedrich, Yannik Frisch, and Philippe C Cattin. Deep generative models for 3d medical image synthesis. In Generative Machine Learning Models in Medical Image Computing, pages 255–278. Springer, 2024a.
- Paul Friedrich, Julia Wolleb, Florentin Bieder, Alicia Durrer, and Philippe C Cattin. Wdm: 3d wavelet diffusion models for high-resolution medical image synthesis. In MICCAI workshop on deep generative models, pages 11–21. Springer, 2024b.
- Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In International Conference on Learning Representations, 2018.
- Mauro Giuffrè and Dennis L Shung. Harnessing the power of synthetic data in healthcare: innovation, application, and privacy. NPJ digital medicine, 6(1):186, 2023.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. Communications of the ACM, 63(11):139–144, 2020.
- Mark S Graham, Walter Hugo Lopez Pinaya, Paul Wright, Petru-Daniel Tudosiu, Yee H Mah, James T Teo, H Rolf Jäger, David Werring, Parashkev Nachev, Sebastien Ourselin, et al. Unsupervised 3d out-of-distribution detection with latent diffusion models. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 446–456. Springer, 2023.
- Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. Advances in neural information processing systems, 33:21271–21284, 2020.
- Xiangming Gu, Chao Du, Tianyu Pang, Chongxuan Li, Min Lin, and Ye Wang. On memorization in diffusion models. arXiv preprint arXiv:2310.02664, 2023.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 9729–9738, 2020.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 16000–16009, 2022.
- Yuwei He, Yuchen Guo, Jinhao Lyu, Liangdi Ma, Haotian Tan, Wei Zhang, Guiguang Ding, Hengrui Liang, Jianxing He, Xin Lou, Qionghai Dai, and Feng Xu. Disorder-free data are all you need — inverse supervised learning for broad-spectrum head disorder detection. NEJM AI, 1(4):Aloa2300137, 2024. . URL <https://ai.nejm.org/doi/abs/10.1056/AIoA2300137>.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598, 2022.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. Advances in neural information processing systems, 33:6840–6851, 2020.
- Xinrong Hu, Yu-Jen Chen, Tsung-Yi Ho, and Yiyu Shi. Conditional diffusion models for weakly supervised medical image segmentation. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 756–765. Springer, 2023.
- Liujiu Hua, Yueyi Luo, Qianqian Qi, and Jun Long. Medicalclip: Anomaly-detection domain generalization with asymmetric constraints. Biomolecules, 14(5):590, 2024.
- Eyke Hüllermeier and Willem Waegeman. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. Machine learning, 110(3): 457–506, 2021.
- Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Illcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In Proceedings of the AAAI conference on artificial intelligence, volume 33, pages 590–597, 2019.

- Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods*, 18(2):203–211, 2021.
- Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems*, 35:26565–26577, 2022.
- Antanas Kascenas, Nicolas Pugeault, and Alison Q O’Neil. Denoising autoencoders for unsupervised anomaly detection in brain mri. In *International Conference on Medical Imaging with Deep Learning*, pages 653–664. PMLR, 2022.
- Daniel S. Kermany, Kang Zhang, and Michael H. Goldbaum. Labeled optical coherence tomography (oct) and chest x-ray images for classification. 2018. URL <https://api.semanticscholar.org/CorpusID:126183849>.
- Firas Khader, Gustav Müller-Franzes, Soroosh Tayebi Arasteh, Tianyu Han, Christoph Haarbuerger, Maximilian Schulze-Hagen, Philipp Schad, Sandy Engelhardt, Bettina Baeßler, Sebastian Foersch, et al. Denoising diffusion probabilistic models for 3d medical image generation. *Scientific Reports*, 13(1):7303, 2023.
- Diederik P Kingma, Max Welling, et al. Auto-encoding variational bayes, 2013.
- Diederik P Kingma, Max Welling, et al. An introduction to variational autoencoders. *Foundations and Trends® in Machine Learning*, 12(4):307–392, 2019.
- Hugo J. Kuijf, J. Matthijs Biesbroek, Jeroen De Bresser, Rutger Heinen, Simon Andermatt, Mariana Bento, Matt Berseth, Mikhail Belyaev, M. Jorge Cardoso, Adrià Casamitjana, D. Louis Collins, Mahsa Dadar, Achilleas Georgiou, Mohsen Ghafoorian, Dakai Jin, April Khademi, Jesse Knight, Hongwei Li, Xavier Lladó, Miguel Luna, Qaiser Mahmood, Richard McKinley, Alireza Mehrtash, Sébastien Ourselin, Bo-Yong Park, Hyunjin Park, Sang Hyun Park, Simon Pezold, Elodie Puybureau, Leticia Rittner, Carole H. Sudre, Sergi Valverde, Verónica Vilaplana, Roland Wiest, Yongchao Xu, Ziyue Xu, Guodong Zeng, Jianguo Zhang, Guoyan Zheng, Christopher Chen, Wiesje van der Flier, Frederik Barkhof, Max A. Viergever, and Geert Jan Biessels. Standardized assessment of automatic segmentation of white matter hyperintensities and results of the wmh segmentation challenge. *IEEE Transactions on Medical Imaging*, 38(11):2556–2568, 2019. .
- Karim Lekadir, Alejandro F Frangi, Antonio R Porras, Ben Glocker, Celia Cintas, Curtis P Langlotz, Eva Weicken, Folkert W Asselbergs, Fred Prior, Gary S Collins, et al. Future-ai: International consensus guideline for trustworthy and deployable artificial intelligence in healthcare. *bmj*, 388, 2025.
- Žiga Lesjak, Anna Galimzianova, Ataman Koren, et al. A novel public mr image dataset of multiple sclerosis patients with lesion segmentations based on multi-rater consensus. *Neuroinform*, 16:51–63, 2018. . URL <https://doi.org/10.1007/s12021-017-9348-7>.
- Jun Li, Cosmin I Bercea, Philip Müller, Lina Felsner, Suhan Kim, Daniel Rueckert, Benedikt Wiestler, and Julia A Schnabel. Multi-image visual question answering for unsupervised anomaly detection. *arXiv preprint arXiv:2404.07622*, 2024.
- Ziyun Liang, Harry Anthony, Felix Wagner, and Konstantinos Kamnitsas. Modality cycles with masked conditional diffusion for unsupervised anomaly segmentation in mri. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 168–181. Springer, 2023.
- Sook-Lei Liew, Bethany P. Lo, ., and et al. Miarnnda R. Donnelly. A large, curated, open-source stroke neuroimaging dataset to improve lesion segmentation algorithms. *Scientific Data*, 9, 2022.
- Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. Pseudo numerical methods for diffusion models on manifolds. *arXiv preprint arXiv:2202.09778*, 2022a.
- Shiyu Liu, Fan Wang, Zehua Ren, Chunfeng Lian, and Jianhua Ma. Controllable counterfactual generation for interpretable medical image classification. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 143–152. Springer, 2024.
- Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022b.
- Y. H. Mah, P. Nachev, and A. D. MacKinnon. Quantifying the impact of chronic ischemic injury on clinical outcomes in acute stroke with machine learning. *Front Neurol.*, 11:15, Jan 2020. . URL <https://doi.org/10.3389/fneur.2020.00015>.

- Lena Maier-Hein, Annika Reinke, Patrick Godau, Minu D Tizabi, Florian Buettner, Evangelia Christodoulou, Ben Glocker, Fabian Isensee, Jens Kleesiek, Michal Kozubek, et al. Metrics reloaded: recommendations for image analysis validation. Nature methods, pages 1–18, 2024.
- Sergio Naval Marimont, Matthew Baugh, Vasilis Siomos, Christos Tzelepis, Bernhard Kainz, and Giacomo Taroni. Disyre: Diffusion-inspired synthetic restoration for unsupervised anomaly detection. arXiv preprint arXiv:2311.15453, 2023.
- Felix Meissen, Georgios Kaissis, and Daniel Rueckert. Challenging current semi-supervised anomaly segmentation methods for brain mri. In International MICCAI brainlesion workshop, pages 63–74. Springer, 2021.
- Felix Meissen, Svenja Breuer, Moritz Knolle, Alena Buyx, Ruth Müller, Georgios Kaissis, Benedikt Wiestler, and Daniel Rückert. (predictable) performance bias in unsupervised anomaly detection. Ebiomedicine, 101, 2024.
- Chenlin Meng, Robin Rombach, Ruiqi Gao, Diederik Kingma, Stefano Ermon, Jonathan Ho, and Tim Salimans. On distillation of guided diffusion models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 14297–14306, 2023.
- Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (BRATS). IEEE transactions on medical imaging, 34(10):1993–2024, 2014.
- E. Nagy, M. Janisch, F. Hržić, and et al. A pediatric wrist trauma x-ray dataset (grazpedwri-dx) for machine learning. Scientific Data, 9:222, 2022. .
- Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In Proceedings of the 38th International Conference on Machine Learning, volume 139, pages 8162–8171. PMLR, 2021a.
- Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In International conference on machine learning, pages 8162–8171. PMLR, 2021b.
- Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In European conference on computer vision, pages 69–84. Springer, 2016.
- William Peebles and Saining Xie. Scalable diffusion models with transformers. In Proceedings of the IEEE/CVF international conference on computer vision, pages 4195–4205, 2023.
- Walter HL Pinaya, Mark S Graham, Robert Gray, Pedro F Da Costa, Petru-Daniel Tudosiu, Paul Wright, Yee H Mah, Andrew D MacKinnon, James T Teo, Rolf Jager, et al. Fast unsupervised brain anomaly detection and segmentation with diffusion models. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 705–714. Springer, 2022a.
- Walter HL Pinaya, Petru-Daniel Tudosiu, Jessica Dafflon, Pedro F Da Costa, Virginia Fernandez, Parashkev Nachev, Sebastien Ourselin, and M Jorge Cardoso. Brain imaging generation with latent diffusion models. In MICCAI Workshop on Deep Generative Models, pages 117–126. Springer, 2022b.
- Walter HL Pinaya, Mark S Graham, Eric Kerfoot, Petru-Daniel Tudosiu, Jessica Dafflon, Virginia Fernandez, Pedro Sanchez, Julia Wolleb, Pedro F Da Costa, Ashay Patel, et al. Generative ai for medical imaging: extending the monai framework. arXiv preprint arXiv:2307.15208, 2023.
- Walter Hugo Lopez Pinaya, Petru-Daniel Tudosiu, Robert Gray, Geraint Rees, Parashkev Nachev, Sébastien Ourselin, and M Jorge Cardoso. Unsupervised brain anomaly detection and segmentation with transformers. In Medical Imaging with Deep Learning, pages 596–617. PMLR, 2021.
- Vivek Kumar Prasad, Ashwin Verma, Pronaya Bhattacharya, Sheryal Shah, Subrata Chowdhury, Madhuri Bhavsar, Sheraz Aslam, and Nouman Ashraf. Revolutionizing healthcare: a comparative insight into deep learning's role in medical imaging. Scientific Reports, 14(1):30273, 2024.
- Aimon Rahman, Jeya Maria Jose Valanarasu, Ilker Hacihaliloglu, and Vishal M Patel. Ambiguous medical image segmentation using diffusion models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 11536–11546, 2023.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18, pages 234–241. Springer, 2015.
- Margherita Rosnati, Mélanie Roschewitz, and Ben Glocker. Robust semi-supervised segmentation with timestep en-

- sembling diffusion models. In *Machine Learning for Health (ML4H)*, pages 512–527. PMLR, 2023.
- Jérémy Rousseau, Christian Alaka, Emma Covili, Hippolyte Mayard, Laura Misrachi, and Willy Au. Pre-training with diffusion models for dental radiography segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 174–182. Springer, 2023.
- Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 conference proceedings*, pages 1–10, 2022.
- Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*, 2022.
- Pedro Sanchez, Antanas Kascenas, Xiao Liu, Alison Q O’Neil, and Sotirios A Tsaftaris. What is healthy? generative counterfactual diffusion for lesion localization. In *MICCAI Workshop on Deep Generative Models*, pages 34–44. Springer, 2022.
- Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *International conference on information processing in medical imaging*, pages 146–157. Springer, 2017.
- Marvin Seyfarth, Salman Ul Hassan Dar, and Sandy Engelhardt. Latent pollution model: The hidden carbon footprint in 3d image synthesis. In *International Workshop on Simulation and Synthesis in Medical Imaging*, pages 146–156. Springer, 2024.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- Onkar Susladkar, Gayatri Deshmukh, Yalcin Tur, Gorkem Durak, and Ulas Bagci. Victr: Vital consistency transfer for pathology aware image synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22772–22782, 2025.
- Jeremy Tan, Benjamin Hou, Thomas Day, John Simpson, Daniel Rueckert, and Bernhard Kainz. 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*, pages 581–591. Springer, 2021.
- Weimin Tan, Siyuan Chen, and Bo Yan. Diffss: Diffusion model for few-shot semantic segmentation. *arXiv preprint arXiv:2307.00773*, 2023.
- Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. *Advances in neural information processing systems*, 34:11287–11302, 2021.
- Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
- Hung Vu Quoc, Thao Tran Le Phuong, Minh Trinh Xuan, and Sang Dinh Viet. Lsegdiff: A latent diffusion model for medical image segmentation. In *Proceedings of the 12th International Symposium on Information and Communication Technology*, pages 456–462, 2023.
- Zihao Wang, Yingyu Yang, Yuzhou Chen, Tingting Yuan, Maxime Sermesant, Herve Delingette, and Ona Wu. Diffusion based zero-shot medical image-to-image translation for cross modality segmentation. *arXiv preprint arXiv:2404.01102*, 2024.
- D. Wilson, G. Ambler, C. Shakeshaft, M. M. Brown, A. Charidimou, R. Al-Shahi Salman, G. Y. H. Lip, H. Cohen, G. Banerjee, H. Houlden, M. J. White, T. A. Yousry, K. Harkness, E. Flossmann, N. Smyth, L. J. Shaw, E. Warburton, K. W. Muir, H. R. Jäger, D. J. Werring, and CROMIS-2 collaborators. Cerebral microbleeds and intracranial haemorrhage risk in patients anticoagulated for atrial fibrillation after acute ischaemic stroke or transient ischaemic attack (cromis-2): a multicentre observational cohort study. *Lancet Neurol.*, 17(6):539–547, Jun 2018. . URL [https://doi.org/10.1016/S1474-4422\(18\)30145-5](https://doi.org/10.1016/S1474-4422(18)30145-5). Erratum in: *Lancet Neurol.* 2018 Jul;17(7):578.
- Julia Wolleb, Florentin Bieder, Robin Sandkühler, and Philippe C Cattin. Diffusion models for medical anomaly detection. In *International Conference on Medical image computing and computer-assisted intervention*, pages 35–45. Springer, 2022a.

- Julia Wolleb, Robin Sandkühler, Florentin Bieder, and Philippe C Cattin. The swiss army knife for image-to-image translation: Multi-task diffusion models. [arXiv preprint arXiv:2204.02641](#), 2022b.
- Julia Wolleb, Robin Sandkühler, Florentin Bieder, Philippe Valmaggia, and Philippe C Cattin. Diffusion models for implicit image segmentation ensembles. In [International Conference on Medical Imaging with Deep Learning](#), pages 1336–1348. PMLR, 2022c.
- Julia Wolleb, Florentin Bieder, Paul Friedrich, Peter Zhang, Alicia Durrer, and Philippe C Cattin. Binary noise for binary tasks: Masked bernoulli diffusion for unsupervised anomaly detection. [arXiv preprint arXiv:2403.11667](#), 2024.
- Junde Wu, Rao Fu, Huihui Fang, Yu Zhang, Yehui Yang, Haoyi Xiong, Huiying Liu, and Yanwu Xu. Medsegdiff: Medical image segmentation with diffusion probabilistic model. In [Medical Imaging with Deep Learning](#), pages 1623–1639. PMLR, 2024a.
- Junde Wu, Wei Ji, Huazhu Fu, Min Xu, Yueming Jin, and Yanwu Xu. Medsegdiff-v2: Diffusion-based medical image segmentation with transformer. In [Proceedings of the AAAI Conference on Artificial Intelligence](#), volume 38, pages 6030–6038, 2024b.
- Julian Wyatt, Adam Leach, Sebastian M. Schmon, and Chris G. Willcocks. Anoddpn: Anomaly detection with denoising diffusion probabilistic models using simplex noise. In [2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops \(CVPRW\)](#), pages 649–655, 2022. .
- Weilai Xiang, Hongyu Yang, Di Huang, and Yunhong Wang. Denoising diffusion autoencoders are unified self-supervised learners. In [Proceedings of the IEEE/CVF International Conference on Computer Vision](#), pages 15802–15812, 2023.
- Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with denoising diffusion gans. [arXiv preprint arXiv:2112.07804](#), 2021.
- Pengfei Yan, Minglei Li, Jiusi Zhang, Guanyi Li, Yuchen Jiang, and Hao Luo. Cold segdiffusion: A novel diffusion model for medical image segmentation. [Knowledge-Based Systems](#), 301:112350, 2024.
- Xingyi Yang and Xinchao Wang. Diffusion model as representation learner. In [Proceedings of the IEEE/CVF International Conference on Computer Vision](#), pages 18938–18949, 2023.
- Sung-Hoon Yoon, Hoyong Kwon, Jaeseok Jeong, Daehee Park, and Kuk-Jin Yoon. Diffusion-guided weakly supervised semantic segmentation. In [European Conference on Computer Vision](#), pages 393–411. Springer, 2024.
- Fahim Ahmed Zaman, Mathews Jacob, Amanda Chang, Kan Liu, Milan Sonka, and Xiaodong Wu. Surf-cdm: Score-based surface cold-diffusion model for medical image segmentation. [arXiv preprint arXiv:2312.12649](#), 2023.
- Lukas Zbinden, Lars Doorenbos, Theodoros Pissas, Adrian Thomas Huber, Raphael Sznitman, and Pablo Márquez-Neila. Stochastic segmentation with conditional categorical diffusion models. In [Proceedings of the IEEE/CVF International Conference on Computer Vision](#), pages 1119–1129, 2023.
- Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zhengnan Huang, Matthew J Muckley, Aaron Defazio, Ruben Stern, Patricia Johnson, Mary Bruno, et al. fastmri: An open dataset and benchmarks for accelerated mri. [arXiv preprint arXiv:1811.08839](#), 2018.
- Jiaying Zhang, Guibo Luo, Zi'Ang Zhang, and Yuesheng Zhu. Data augmentation in class-conditional diffusion model for semi-supervised medical image segmentation. In [2024 International Joint Conference on Neural Networks \(IJCNN\)](#), pages 1–8. IEEE, 2024a.
- Zheyuan Zhang, Lanhong Yao, Bin Wang, Debesh Jha, Gorkem Durak, Elif Keles, Alpay Medetalibeyoglu, and Ulas Bagci. Diffboost: Enhancing medical image segmentation via text-guided diffusion model. [IEEE Transactions on Medical Imaging](#), 2024b.
- Ruiyang Zhao, Burhaneddin Yaman, Yuxin Zhang, Russell Stewart, Austin Dixon, Florian Knoll, Zhengnan Huang, Yvonne W Lui, Michael S Hansen, and Matthew P Lungren. fastmri+: Clinical pathology annotations for knee and brain fully sampled multi-coil MRI data. [arXiv preprint arXiv:2109.03812](#), 2021.
- Hongkai Zheng, Weili Nie, Arash Vahdat, Kamyar Azizadenesheli, and Anima Anandkumar. Fast sampling of diffusion models via operator learning. In [International Conference on Machine Learning](#), pages 42390–42402. PMLR, 2023.
- S Kevin Zhou, Hayit Greenspan, and Dinggang Shen. [Deep learning for medical image analysis](#). Academic Press, 2023.
- Jiaqi Zhu, Shaofeng Cai, Fang Deng, and Junran Wu. Do llms understand visual anomalies? uncovering llm capabilities in zero-shot anomaly detection. [arXiv preprint arXiv:2404.09654](#), 2024.

Lingting Zhu, Zeyue Xue, Zhenchao Jin, Xian Liu, Jingzhen He, Ziwei Liu, and Lequan Yu. Make-a-volume: Leveraging latent diffusion models for cross-modality 3d brain mri synthesis. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 592–601. Springer, 2023.

David Zimmerer, Simon AA Kohl, Jens Petersen, Fabian Isensee, and Klaus H Maier-Hein. Context-encoding variational autoencoder for unsupervised anomaly detection. arXiv preprint arXiv:1812.05941, 2018.

David Zimmerer, Fabian Isensee, Jens Petersen, Simon Kohl, and Klaus Maier-Hein. Unsupervised anomaly localization using variational autoencoders. In Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part IV 22, pages 289–297. Springer, 2019.

David Zimmerer, Peter M Full, Fabian Isensee, Paul Jäger, Tim Adler, Jens Petersen, Gregor Köhler, Tobias Ross, Annika Reinke, Antanas Kascenas, et al. Mood 2020: A public benchmark for out-of-distribution detection and localization on medical images. IEEE Transactions on Medical Imaging, 41(10):2728–2738, 2022.

Appendix

In Table 3 we provide the links to publicly available datasets reported in Table 2. In Table 4, we provide links to the Github code repositories.

Table 3: Links to publicly available datasets used in diffusion-based medical anomaly detection.

Dataset	Link
BraTS	https://www.med.upenn.edu/cbica/brats2020/data.html
ATLAS	https://atlas.grand-challenge.org/
WMH	https://wmh.isi.uu.nl/
MSLub	https://pubmed.ncbi.nlm.nih.gov/29995847/
FastMRI+	https://arxiv.org/abs/2109.03812
CheXpert	https://stanfordmlgroup.github.io/competitions/chexpert/
OCT2017	https://data.mendeley.com/datasets/rscbjbr9sj/2
GRAZPED	https://figshare.com/articles/dataset/GRAZPEDWRI-DX/14825193

Table 4: Key diffusion-based anomaly detection and segmentation methods with direct links.

Method	First author
Fully supervised lesion segmentation	
SegDiff → https://github.com/tomeramit/SegDiff	Tomer Amit
Diffusion-based Segmentation Ensembles → https://github.com/JuliaWolleb/Diffusion-based-Segmentation	Julia Wolleb
MedSegDiff → https://github.com/SuperMedIntel/MedSegDiff	Junde Wu
DermoSegDiff → https://github.com/xmindflow/DermoSegDiff	Afshin Bozorgpour
BerDiff (Conditional Bernoulli Diffusion) → https://github.com/takimailto/BerDiff	Tao Chen
Self-supervised anomaly simulation	
AnoDDPM (Simplex noise) → https://github.com/Julian-Wyatt/AnoDDPM	Julian Wyatt
PatchDDPM (Patched Diffusion Models for UAD) → https://github.com/FinnBehrendt/patched-Diffusion-Models-UAD	Finn Behrendt
DISYRE (Diffusion-Inspired Synthetic Restoration) → https://github.com/snavalm/disyre	Sergio N. Marimont
Weakly supervised lesion localization	
Diffusion Models for Medical Anomaly Detection → https://github.com/JuliaWolleb/diffusion-anomaly	Julia Wolleb
What is Healthy? Generative Counterfactual Diffusion → https://github.com/vios-s/Diff-SCM	Pedro Sanchez
AnoFPDM → https://github.com/SoloChe/AnoFPDM	Yiming Che
Unsupervised lesion localization	
AutoDDPM (Mask–Stitch–Re-Sample) → https://github.com/ci-ber/autoDDPM	Cosmin I. Bercea
Fast UAD with Diffusion Models → https://github.com/Project-MONAI/GenerativeModels/tree/main/tutorials/generative/2d_ldm	Walter H. L. Pinaya
cDDPM (Conditioned DDPMs for UAD) → https://github.com/FinnBehrendt/Conditioned-Diffusion-Models-UAD	Finn Behrendt
MMCCD (Modality Cycles + Masked Conditional Diffusion) → https://github.com/ZiyunLiang/MMCCD	Ziyun Liang
Binary-noise UAD (Anomaly_berdiff) → https://github.com/JuliaWolleb/Anomaly_berdiff	Julia Wolleb
THOR (Temporal Harmonization for Optimal Restoration) → https://github.com/ci-ber/THOR_DDPM	Cosmin I. Bercea