



Guided reconstruction with conditioned diffusion models for unsupervised anomaly detection in brain MRIs

Finn Behrendt^a , Debayan Bhattacharya^a , Robin Mieling^a , Lennart Maack^a,
Julia Krüger^b, Roland Opfer^b, Alexander Schlaefer^a

^a Hamburg University of Technology, Hamburg, Germany

^b Jung Diagnostics GmbH, Hamburg, Germany

ARTICLE INFO

Keywords:

Unsupervised anomaly detection
Segmentation
Brain MRI
Diffusion models

ABSTRACT

The application of supervised models to clinical screening tasks is challenging due to the need for annotated data for each considered pathology. Unsupervised Anomaly Detection (UAD) is an alternative approach that aims to identify any anomaly as an outlier from a healthy training distribution. A prevalent strategy for UAD in brain MRI involves using generative models to learn the reconstruction of healthy brain anatomy for a given input image. As these models should fail to reconstruct unhealthy structures, the reconstruction errors indicate anomalies. However, a significant challenge is to balance the accurate reconstruction of healthy anatomy and the undesired replication of abnormal structures. While diffusion models have shown promising results with detailed and accurate reconstructions, they face challenges in preserving intensity characteristics, resulting in false positives. We propose conditioning the denoising process of diffusion models with additional information derived from a latent representation of the input image. We demonstrate that this conditioning allows for accurate and local adaptation to the general input intensity distribution while avoiding the replication of unhealthy structures. We compare the novel approach to different state-of-the-art methods and for different data sets. Our results show substantial improvements in the segmentation performance, with the Dice score improved by 11.9%, 20.0%, and 44.6%, for the BraTS, ATLAS and MSLUB data sets, respectively, while maintaining competitive performance on the WMH data set. Furthermore, our results indicate effective domain adaptation across different MRI acquisitions and simulated contrasts, an important attribute for general anomaly detection methods. The code for our work is available at <https://github.com/FinnBehrendt/Conditioned-Diffusion-Models-UAD>.

1. Introduction

Magnetic Resonance Imaging (MRI) is an important tool for diagnosing various conditions in the human brain [1,2]. However, interpreting brain MRIs can be error-prone, time-consuming, and places a significant workload on available radiologists [3,4]. To address these challenges and improve diagnostic efficiency, deep learning techniques like convolutional neural networks (CNN) have shown great promise in assisting radiologists by automating certain aspects of the analysis [2]. A common task is the detection and delineation of pathological structures in the MRI scans such as tumors [5], white matter lesions [6] or Alzheimer's disease [7]. Supervised deep learning approaches exhibit robust performance for these tasks, given that task-specific, annotated data sets are available. However, gathering such data sets is a cumbersome and costly process. Additionally, applying supervised models for screening tasks is difficult as any pathology must be detected, even

those that are underrepresented or not included in the training data. In this context, screening tasks refer to scenarios where various conditions need to be identified without a specific target condition in mind. This includes assisting radiologists in detecting a wide range of findings, from expected abnormalities to unexpected or incidental ones, as well as applications in large population studies, such as the Hamburg City Health Study [8].

An alternative approach is unsupervised anomaly detection (UAD), which relies on healthy data instead of annotated pathologies. The goal is to learn the underlying data distribution of healthy brain MRI scans and to identify anomalies as outliers from that learned distribution. A popular approach is reconstruction-based UAD, where generative models (GM) are trained to reconstruct healthy anatomy. Given that the GMs are trained exclusively on healthy data, they should fail to generate unhealthy components of the input images and replace them

* Corresponding author.

E-mail address: finn.behrendt@tuhh.de (F. Behrendt).

<https://doi.org/10.1016/j.complbiomed.2025.109660>

Received 16 February 2024; Received in revised form 29 November 2024; Accepted 6 January 2025

Available online 22 January 2025

0010-4825/© 2025 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

by approximations of healthy anatomy. Subsequently, the discrepancy of input and pseudo-healthy reconstruction is measured, e.g., by the mean absolute error where large errors indicate anomalies [9–12]. Therefore, in theory, any deviation from the learned norm can be localized, including known conditions, unexpected anomalies such as artifacts, or previously unseen pathologies.

However, in practice, reconstruction-based UAD methods typically do not result in perfect anomaly detection and segmentation. Considering that the training task is reconstructing the input image, a key challenge is to limit the reconstruction to healthy brain anatomy. On the one hand, GMs that generate highly accurate reconstructions tend to perform a ‘copy task’, resulting in unhealthy structures still reflected in the pseudo-healthy reconstructions. These unhealthy structures do not deviate from the input image, resulting in false negatives in the segmentation mask. On the other hand, GMs with limited reconstruction accuracy may produce imperfect reconstructions everywhere, which in turn appear as differences in the anomaly map, even for healthy structures. This complicates distinguishing between actual pathologies and reconstruction errors, typically causing false positives in the segmentation map. Another related challenge in UAD is the potential domain shift between the distribution of the healthy data and the unhealthy test cases. Here, the reconstruction may appear different due to the domain shift; e.g., in simple cases, it is generally brighter than the input, and this shift reflects differences in the anomaly map. In summary, the challenge is to avoid copying unhealthy input features to the reconstruction while still adapting the general appearance of the reconstruction to the input.

Recently, denoising diffusion probabilistic models (DDPMs) [13] have shown promise as GMs for reconstruction-based UAD in brain MRI [14–16]. DDPMs generate images by denoising images corrupted by artificial noise, leveraging a high-dimensional latent space to preserve spatial context and achieve high-fidelity reconstructions. While the spatial latent space enables the accurate reconstruction of healthy structures, the additional denoising task aims to prevent the DDPMs from solely copying the content of the input image. Therefore, DDPMs can achieve a reasonable trade-off between the reconstruction accuracy of healthy and unhealthy structures [15,16]. However, a significant challenge remains in accurately reconstructing healthy brain anatomy that exhibits aligned intensity characteristics with the input image. The forward and backward processes of DDPMs do not adequately capture the highly variable local intensity distributions of MRI scans. This can result in discrepancies between the input and the reconstruction, which are difficult to distinguish from those arising from actual pathologies. This leads to reduced segmentation performance, particularly when facing domain shifts at test time. One potential solution could be to incorporate the input image as an additional input to the denoising process in the DDPM, e.g., as a second input channel. However, this could allow the denoising process to replicate the content of the input image, which would contradict the principles of reconstruction-based UAD.

In response to these challenges, we propose the use of conditioned DDPMs (cDDPMs) for UAD in brain MRI. Our approach includes a conditioning mechanism that guides the denoising process of the DDPM by utilizing an extra feature representation of the input image. This feature representation, derived from a CNN-based image encoder, does not capture detailed structural information. However, it contains the coarse local intensity information from the input image, which is often partially lost during the DDPM’s forward process. This way, we aim to align the local intensity distribution of the reconstructed image without providing detailed structural information that could be used to replicate unhealthy structures. We conduct a comprehensive investigation of our conditioning approach, specifically addressing the challenges associated with reconstruction-based UAD in brain MRI. Initially, we evaluate the quality of healthy brain MRI reconstructions and assess the ability to reconstruct healthy anatomy while replacing unhealthy structures. Subsequently, we examine the domain adaptation

capabilities of our method by assessing the alignment of intensity between input and reconstruction for datasets not encountered during training and by simulating varying contrast levels. Finally, we evaluate the segmentation performance of our approach on a variety of datasets, comparing established state-of-the-art UAD methods.

Our results indicate that our conditioning approach effectively aligns the local intensity distributions of input and reconstruction without supporting the replication of unhealthy structures. Furthermore, cDDPMs can effectively adapt to different intensity and contrast profiles of different MRI data sets. As a result, our proposed cDDPMs address key challenges of DDPMs in reconstruction-based UAD in brain MRI and improve or match the segmentation performance of the compared baseline models. When compared to DDPMs, the Dice score significantly increases ($p < 0.05$), rising from 50.27 to 56.30 for the BraTS21 dataset and from 20.18 to 24.22 for the ATLAS v2 dataset. For the MSLUB dataset, the performance improves from 9.71 to 14.04 and for the WMH dataset, both models report similar performance, with Dice scores of 12.06 and 11.59, respectively.

In summary, the main contributions of this work are:

- **Reconstruction Quality:** We develop a conditioning mechanism within cDDPMs that guides the denoising process of DDPMs while avoiding the replication of unhealthy structures.
- **Domain Adaptation and Intensity Alignment:** Our conditioning mechanism aligns the local intensity distribution of the reconstructed image with that of the input image. This effectively improves the generalization to intensity shifts of different MRI scans.
- **Segmentation Performance:** The accurate reconstructions and aligned local intensity distributions featured by our conditioning mechanism can significantly improve the segmentation performance and applicability of DDPMs for UAD in brain MRI.

This paper is organized as follows: In Section 2, we review relevant literature on UAD in brain MRI. In Section 3, we introduce DDPMs and subsequently explain our conditioning approach. In Section 4, we provide details of the experimental setup. In Section 5, we present the results and subsequently discuss them in Section 6. Finally, we provide a conclusion in Section 7.

2. Recent work

Autoencoders (AE) have been the primary focus of research on reconstruction-based UAD in brain MRI. Although these models exhibit potential in capturing the underlying healthy distribution, their effectiveness in UAD is limited by their blurry reconstructions [9]. To overcome this limitation, researchers have focused on improving the representations and reconstructions by adding skip connections with dropout [17], using multi-scale features [18], utilizing feature activation maps [19] or employing feature discrepancies [20,21]. Additionally, online outlier removal strategies [22] have been proposed for AEs. In parallel, Variational Autoencoders (VAE) have been investigated for the UAD task [23], focusing on enhancing the used context in 2D [24] and 3D [25,26] or utilizing restoration methods [11]. Moreover, Soft-Intro VAEs (SIVAE) [27] have been investigated. Additionally, Generative Adversarial Networks (GAN) have been explored for the task of UAD either as pure GAN [28,29] or in combination with VAEs [30,31] and VQ-VAEs with transformers [12].

While AEs with skip connections and a spatial latent space enable reconstructions of high fidelity, they tend to perform a ‘copy task’ which enables the reconstruction of unhealthy anatomy and therefore contradicts the UAD principle [9,27]. Lately, [10] have shown that AEs with skip connections can be effectively used for UAD in brain MRI if they are regularized by an additional denoising task. Taking a similar direction, DDPMs have shown promise in the field of UAD [14, 15,32,33]. DDPMs provide high reconstruction fidelity, but important information about the input image can be lost due to the noising

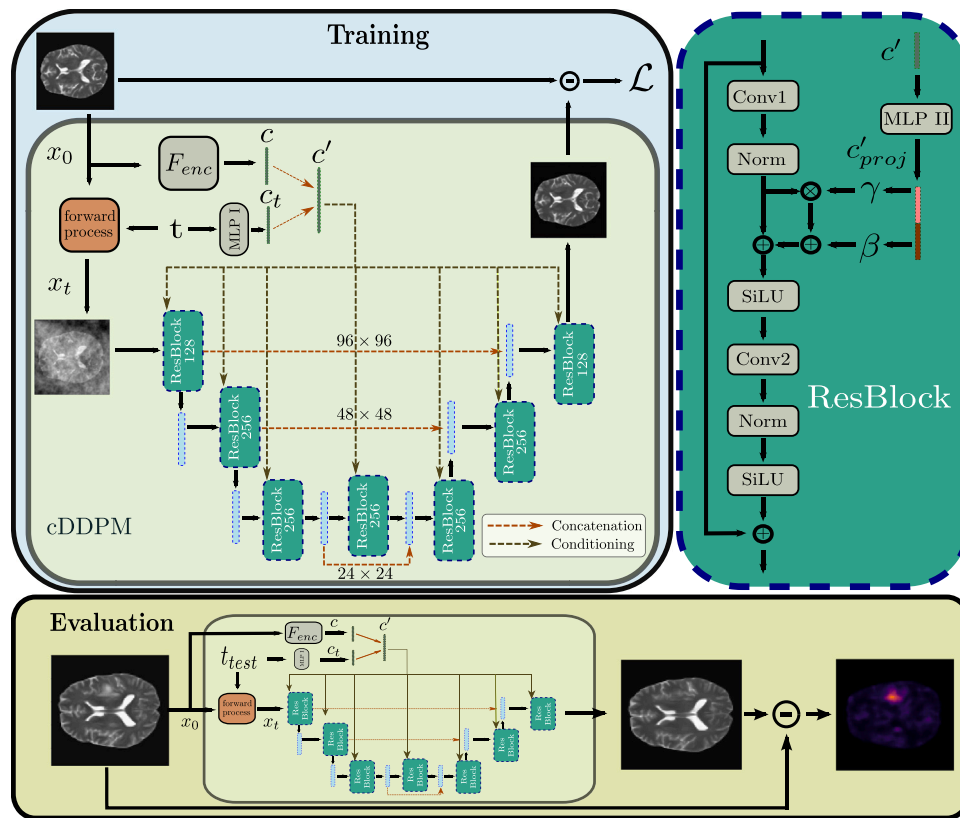


Fig. 1. Overview of our proposed approach. The encoder representations are learned along with the DDPM in the main training stage to condition the denoising process. The timestep embedding c_t is concatenated with the input image’s projected encoder representation c . The resulting conditioning vector c' is used to scale and shift feature maps of the denoising Unet in the residual blocks. Each residual block consists of two convolution operators (conv1, conv2), group normalization (Norm) and Sigmoid Linear Units (SiLU). During evaluation, the residual map between unhealthy brain images and their healthy reconstructions is used for anomaly detection.

process. To address this, [16] proposed patch-based DDPMs that allow the use of parts of the original image content to provide information for the reconstruction of the input image. However, this patching strategy increases complexity and computational effort and can lead to artifacts in overlapping patch regions. A more efficient approach is seen in conditioning the denoising process of DDPMs with knowledge of the input image. Conditioned DDPMs have been successful in text-to-image synthesis tasks [34] and image-guided synthetic image generation [35, 36]. However, in the specific case of UAD, the objective is not to generate new images or to transfer styles but to accurately estimate a given input image while ensuring that unhealthy anatomy is absent in the estimation. Directly conditioning DDPMs with information from the input image can pose a risk of reconstructing unhealthy anatomy. Recent studies leverage different approaches to fuse information from a target image into the generation process of DDPMs. [37] introduce a classifier model trained to differentiate between healthy and tumorous brain MRI scans. At its core, this model employs classifier guidance during the sampling process, targeting the denoising towards a ‘healthy’ classification. Similarly, [38] apply classifier conditioning for counterfactual generation to remove tumors by conditioning the model to generate ‘healthy’ images. Despite showing promise, these methods rely on sample-level label information of given pathologies, contradicting the unsupervised setting. Moreover, their reliance on a scalar variable for conditioning limits the capacity to capture complex image characteristics. The authors [39] approach this challenge by incorporating synthetic anomalies in two separate diffusion processes. Their method involves feeding concatenated noised images (normal and anomalously altered) into a denoising Unet. This approach relies on synthetic defects that can be added to the input images. While effective in industrial defect detection, as demonstrated on the MVTEC dataset, this approach’s dependency on synthetic anomalies raises concerns

about its applicability to medical datasets. [40] propose an approach to fuse information of a target image into the generation process. By estimating and applying the same noise pattern to both the input and target images, their model aims to minimize the variance between these noised versions. While guiding the generation towards a target image, this method results in the loss of input image information, as the same noise is applied to the target image.

Our approach, detailed in the following sections, addresses these limitations by introducing a simple conditioning mechanism applied to the denoising Unet within DDPMs. In contrast to [37,38], we utilize spatial conditioning information and rely on unlabeled data. Unlike the proposed conditioning in [39], our approach does not rely on simulated anomalies and we directly condition the denoising process without the need for additional sub-networks. This direct conditioning of the denoising process in the Unet diverges from [40]’s approach, ensuring that valuable information from the target image is retained and utilized effectively throughout the diffusion process.

3. Methods

In our proposed cDDPM we use an image encoder network and embed the input image in a context vector $c \in \mathbb{R}^d$ to condition the denoising Unet on meaningful features of the input image. Our motivation is that the additional information in c guides the generation process towards consistent intensity characteristics across the input image and its reconstruction. Hence, by introducing the context vector c we aim to recover local intensity information lost during the forward (noising) process of DDPMs. We utilize an image encoder with a dense latent space to extract information regarding the coarse shape and local intensity information of the noise-free input image. This latent representation can then be used to condition the denoising process and

supplement the individual context of the input image without providing detailed pixel-wise information that could be used to perform a ‘copy task’. A general depiction of our approach is shown in Fig. 1.

3.1. DDPMs

DDPMs are generative models that learn the underlying data distribution of images $\mathbf{x} \in \mathbb{R}^{H,W,C}$ with height H , width W and C channels, given a training set. Training of DDPMs consists of two steps. The forward process, where an input image \mathbf{x}_0 is gradually transformed to Gaussian noise $\mathbf{x}_T = \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and the backward process, where reversing the forward process is learned.

In the forward process, transforming \mathbf{x}_0 to \mathbf{x}_T follows a predefined schedule β_1, \dots, β_T , where intermediate versions \mathbf{x}_t are derived as

$$\mathbf{x}_t \sim q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\sqrt{\bar{\alpha}_t} \mathbf{x}_0, (1 - \bar{\alpha}_t) \mathbf{I}),$$

$$\text{with } \bar{\alpha}_t = \prod_{s=0}^t (1 - \beta_s).$$

The time step t controls the amount of added noise and is sampled from $t \sim \text{Uniform}(1, \dots, T)$. For edge cases, the image \mathbf{x}_t is transformed to pure noise ($t = T$) or no transformation is applied ($t = 0$). In the backward process, the reconstructed image \mathbf{x}_0^{rec} is recovered from \mathbf{x}_t by

$$\mathbf{x}_0^{rec} \sim p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t),$$

$$\text{with } p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(\boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t)).$$

Here, following [13], $\boldsymbol{\mu}_\theta$ is estimated by a Unet [41] with trainable parameters θ , and $\boldsymbol{\Sigma}_\theta(t) = \boldsymbol{\Sigma}(t) = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t \mathbf{I}$ is fixed. Variational inference is used to achieve a tractable loss function and the variational lower bound (VLB) is derived as

$$\begin{aligned} \mathcal{L}_{VLB} &= -\log(p_\theta(\mathbf{x}_0)) \\ &+ D_{KL}(q(\mathbf{x}_{1:T} | \mathbf{x}_0) \| p_\theta(\mathbf{x}_{1:T} | \mathbf{x}_0)). \end{aligned}$$

which can be reformulated to

$$\mathcal{L}_{simple} = \|\epsilon - \epsilon_\theta(\mathbf{x}_t, t)\|^2$$

by applying simplifications and by conditioning the denoising step on \mathbf{x}_0 , as shown in [13]. In our work, instead of predicting the noise ϵ we perform the equivalent task of directly estimating $\mathbf{x}_0^{rec} = \mathbf{x}_t - \epsilon$. Hence, we derive our loss function as

$$\mathcal{L}_{rec} = |\mathbf{x}_0 - \mathbf{x}_0^{rec}|.$$

Typically, to generate new images with DDPMs, the backward step is applied step-wise to gradually denoise a random noise vector. For the given UAD task, we do not aim to generate new images but to estimate healthy brain anatomy given an input image. Therefore, we directly estimate \mathbf{x}_0^{rec} given \mathbf{x}_t at test time as it is done in [16]. The time step $t_{test} < T$ controls the level of noise to remove from \mathbf{x}_t at test time. Optionally, to become agnostic to the noise magnitude, we use an ensemble of different values $t_{test} = [250, 500, 750]$ and average the reconstructions of each noise level, similar to [32].

3.2. Conditioned DDPMs (cDDPMs)

A general depiction of our conditioning approach is provided in Fig. 1. Formally, we condition the backward process of DDPMs on a context vector c as follows

$$\mathbf{x}_0^{rec} \sim p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t, c),$$

$$\text{with } p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t, c) = \mathcal{N}(\boldsymbol{\mu}_\theta(\mathbf{x}_t, t, c), \boldsymbol{\Sigma}(t)).$$

We use an image encoder F_{enc} to achieve a latent representation $c = F_{enc}(\mathbf{x}_0)$ of the input image \mathbf{x}_0 where $c \in \mathbb{R}^d$ with d as conditioning

dimension.

To integrate the context vector c , we manipulate the denoising Unet of the DDPM. Therefore, we individually adapt the features $f_i \in \mathbb{R}^{H_i, W_i, C_i}$ at each level of the denoising Unet based on c where H_i , W_i and C_i are the respective feature map dimensions. To achieve this, we adapt the time step conditioning of DDPMs as follows. First, the time step is encoded using a sinusoidal position embedding. Next, t is projected to a vector $c_t \in \mathbb{R}^d$ by a multi-layer perceptron (MLP I). Subsequently, we concatenate the context vector c and the time step vector c_t , resulting in a conditioning vector $c' \in \mathbb{R}^{2 \cdot d}$. Finally, c' is projected to $c'_{proj} \in \mathbb{R}^{2 \cdot C_i}$ by another multi-layer perceptron (MLP II) at each feature level i . The vector c'_{proj} is then split into half, where the first and last C_i elements resemble the scaling factor γ and the shift value β . Inspired by [42], the variables γ and β are used to transform the individual feature maps as $f'_i = f_i * (\gamma + 1) + \beta$ in each residual block. This transformation adaptively scales and shifts the feature maps at each level of the denoising UNet, based on the context vector $c = F_{enc}(\mathbf{x}_0)$, which encodes relevant information from the input image. The purpose of this transformation is to allow the model to dynamically adjust feature representations in response to both the image features and the conditioning information. By modulating the feature maps through context-based scaling and shifting, the model can more effectively preserve critical details of the clean target image while denoising the noisy input.

Optionally, to achieve a meaningful starting point for the calculation of the context vector $c = F_{enc}(\mathbf{x}_0)$, we pre-train the feature extraction of the image encoder F_{enc} which is described in the next section.

3.3. Pre-training

We utilize a generative pre-training strategy for F_{enc} . More precisely, we utilize masked pre-training where typically transformer-based AEs are trained to reconstruct an image where a significant fraction of patches are masked out [43]. We adopt the SparK framework [44], where sparse convolutions and hierarchical features are used to enable the masked pre-training for CNNs. We pre-train the encoder with the same healthy training set as the main training task to learn the general feature representations required to capture important information from the MRI scans. After the pre-training stage, we discard the decoder and only use F_{enc} and fine-tune it along with the denoising Unet during the main training stage of the cDDPM. A schematic description of the pre-training stage is provided in Fig. 2.

4. Experimental setup

4.1. Data sets

Following the principle of UAD, we train our models for the reconstruction task on healthy data only (IXI). At test time, we evaluate the models’ anomaly detection ability on unhealthy test sets of various pathologies (BraTS21, ATLAS, MSLUB, WMH). We provide an overview of all available MRI scanner details for the different data sets in Table 5.

4.1.1. Training data

We use the publicly available IXI data set¹ as our healthy reference data set for training. This data set includes 560 3D brain MRI scans collected from three medical facilities. Of the training data, 158 samples are set aside for testing, while the remaining data is divided into 5 folds, each containing 358 training samples and 44 validation samples for cross-validation.

Note that only the healthy training and validation data is used during pre-training and training stage.

¹ <https://brain-development.org/ixi-dataset/>




Fig. 2. Overview of the pre-training strategy. Random patches are erased from the input image. F_{enc} is used to derive a latent representation and F_{dec} is used to reconstruct the removed patches. After pre-training, the encoder F_{enc} is then fine-tuned along with the $DDPM$ in the main training stage to condition the denoising process.

4.1.2. Evaluation data

For evaluation, we utilize four different publicly available data sets that contain different types of pathologies and the corresponding manual expert annotations:

1. Multimodal Brain Tumor Segmentation Challenge 2021 (BraTS21) [45–47]
2. Multiple Sclerosis data set from the University Hospital of Ljubljana (MSLUB) [48]
3. Anatomical Tracings of Lesions After Stroke v2.0 (ATLAS) [49]
4. White Matter Hyperintensity (WMH) [50]

The BraTS21 data set includes 2040 3D brain routine MRI scans of patients with glioma with a pathologically confirmed diagnosis. Accompanying the MRI scans, annotations from expert neuroradiologists are provided for 1251 scans that delineate tumor sub-regions as categorical masks. We fuse all sub-regions to obtain a binary segmentation mask to evaluate the anomaly detection task. All scans are available as T1-weighted volumes with and without contrast enhancement (T1-CE, T1) and T2-weighted or T2 fluid-attenuated inversion recovery (T2, FLAIR) volumes. The MSLUB data set includes 3D brain MRI scans of 30 patients with multiple sclerosis (MS) lesions. For each patient, along with the T1, T2 and FLAIR MRI scans, ground truth annotations are available derived based on multi-rater consensus. The ATLAS data set consists of 655 T1-weighted MRI scans of stroke patients collected from 44 research cohorts. The stroke lesions are annotated by domain experts and binary segmentation masks are provided. The WMH data set consists of 60 MRI scans of patients with white matter hyperintensities from three different institutions and scanner types. WMH segmentation masks are derived from the consensus of two expert radiologists. The

datasets contain images acquired with different MRI parameters. While BraTS21 and MSLUB include T1, T2, and FLAIR data, WMH contains only T1 and FLAIR data, and ATLAS only T1 data. The IXI data set used for training contains T1 and T2 data. Hence, we train the models separately on T1 and T2 data, and evaluate on the respective data sets.

4.2. Pre-processing

We pre-process the images according to established pre-processing strategies for UAD in brain MRI [9]. First, we resample all MRI scans to the isotropic resolution of $1\text{ mm} \times 1\text{ mm} \times 1\text{ mm}$ using cubic spline interpolation. Second, we register all MRI scans to the SRI24-Atlas. Third, we remove the skull from the MRI scans by skull stripping with HD-BET [51]. Subsequently, we crop each brain scan using its corresponding brain mask, removing unnecessary background while preserving relevant brain tissue. We then apply N4 bias field correction [52]. To standardize image sizes, we identify the largest brain volume in our dataset, add a safety margin, and pad all images to a unified size of $192 \times 192 \times 160$. Finally, to save computational resources, we reduce volume resolution by a factor of two and remove the 15 top and bottom slices parallel to the transverse plane, leading to a final resolution of $96 \times 96 \times 50$ voxels.

4.3. Post-processing

At test time, we derive a binary segmentation map from the residual map $\mathbf{R} = |\mathbf{x}_0 - \mathbf{x}_0^{rec}|$ where regions of high residuals indicate anomalies. To binarize \mathbf{R} , we first apply the following post-processing steps that are commonly used in the field of UAD in brain MRI [9,10,16,24]. First, a median filter with a kernel size of $K = 5 \times 5 \times 5$ is applied to smooth the residual map and to remove false positives. Subsequently, we perform brain mask eroding for 3 iterations. This step is mainly applied to filter out residuals due to poor reconstructions at sharp edges near the brain mask [9]. After binarization, we use connected component filtering to remove areas that include less than 7 voxels. This post-processing step removes false-positive predictions smaller than the anomalies considered in this study. In the supplemental material, we systematically compare the post-processing strategies for each data set.

4.4. Implementation details

This study compares our proposed cDDPM method with multiple established baselines for UAD in brain MRI. The baselines include AE, VAE [9], SVAE [26], Reverse Anomaly RA [27], PHANES [31] and denoising AEs DAE [10]. We also compare with simple thresholding *Thresh* [53]. Furthermore, we compare the feature-based reverse distillation method RD [21], Feature-Autoencoders FAE [20], the Encoder-Decoder Contrast method EDC [54] and the self-supervised Poisson Image Interpolation PII [55]. For a direct comparison, we also include DDPM [15] and patched DDPMs pDDPM [16] as a counterpart to our proposed method.

We adapt the baseline implementations by tuning hyper-parameters based on the validation set to improve training stability and performance. We set β_{VAE} to 0.001 for VAE and SVAE. For EDC, we use a lr of $1e-5$ and $5e-5$ for the encoder and decoder, respectively. RA and PHANES are trained with $\beta_{rec} = 16$ for the Soft-Intro VAE. For PHANES, we use a masking threshold of 0.15 and 0.10 for training and testing, respectively.

For DDPM, pDDPM and cDDPM, we use structured simplex noise, which has been shown to improve the UAD performance on MRI images [15]. Furthermore, we uniformly sample $t \in [1, T]$ with $T = 1000$ during training. At test time, we either use a fixed value of $t_{test} = \frac{T}{2} = 500$ or an ensemble of different values $t_{test} = [250, 500, 750]$ and average the individual reconstructions of each noise level. The denoising network for all DDPM-based approaches is an Unet similar to [34], with channel dimensions of [128, 256, 256]. As encoder network F_{enc} ,

we utilize a ResNet-backbone with a fully connected layer to match the target dimension of $c \in \mathbb{R}^d$ with $d = 128$ as conditioning dimension. We evaluated several conditioning dimensions, $d \in \{32, 64, 128, 256, 512\}$, and observed that the model achieved optimal performance when $d = 128$. A mask-out ratio of 65% is used during pre-training. For data augmentation, we utilize random -blur ($p = 0.25$), -bias ($p = 0.25$), -gamma ($p = 0.5$) and -ghosting ($p = 0.5$) from the torchio library [56]. We train for a maximum of 1600 epochs on NVIDIA V100 (32 GB) GPUs, using Adam as an optimizer, a learning rate of $1e-4$, and a batch size of 32. The best model checkpoint, as determined by performance on the validation set, is used for testing. The volumes are processed slice-wise, uniformly sampling slices with replacement during training and iterating over all slices to reconstruct the full volume at test time. We implement all models in Pytorch (v0.10).

4.5. Experiments and evaluation

We conduct various experiments to assess the individual features of our proposed cDDPMs. First, we evaluate the reconstruction quality for both healthy and unhealthy structures. Second, we investigate the generalization to real and simulated domain shifts between the training and test data. Lastly, we assess the final segmentation performance. For all experiments, we compare our cDDPMs to established baseline models on different data sets. The following subsections provide the detailed experimental settings for each individual evaluation.

4.5.1. Reconstruction quality

To evaluate the overall reconstruction quality, we utilize the held-out test set of the healthy IXI data set and calculate similarity metrics between input and reconstruction. We consider the Structural Similarity Index Measure (SSIM) [57], the Peak Signal To Noise Ratio (PSNR) and the Learned Perceptual Image Patch Similarity (LPIPS) as metrics to assess the reconstruction quality. For the feature-based LPIPs metric [58], features are extracted by a ResNet-based network, pre-trained on 3D medical data [59]. Furthermore, we report the reconstruction error ($l1$ -error) for the healthy data set. For UAD, only healthy anatomy should be reconstructed. Hence, it is interesting to consider the $l1$ -error of healthy and unhealthy anatomy separately, given the unhealthy evaluation data sets. Therefore, we calculate the $l1$ -error for both healthy and unhealthy anatomy, as indicated by the annotation masks and calculate an $l1$ -ratio as follows:

$$l1\text{-ratio} = \frac{l1_{unhealthy}}{l1_{healthy}}.$$

A higher value for the $l1$ -ratio indicates that the model successfully reconstructs the healthy anatomy while struggling to reconstruct the unhealthy parts of the input, and vice versa.

4.5.2. Domain adaptation and intensity alignment

We investigate our proposed approach's generalization to domain shifts and the capability of adequately reconstructing the local intensity patterns of a given input image. We utilize data sets from different domains. For training, we utilize in-domain data from the IXI data set. As out-of-domain data, we utilize the BraTS21 data set. Notably, we only consider the content of the BraTS21 data set that is marked as healthy by the provided annotation masks, thereby ensuring the evaluation of domain shifts regarding scanner and brain diversity, not domain shifts arising from unhealthy brain MRI structures. Furthermore, we simulate different contrast levels ranging from $cl \in [0.3, 0.7, 1.0, 1.5, 2.0]$. The images of different contrast levels are derived by potentiating the gray values by the respective contrast level, i.e., $x_0^{cl=2} = x_0^2$. To evaluate the effect of the conditioning mechanism, we utilize the IXI data set and simulate different levels of conditioning information to investigate the reconstructions qualitatively. Therefore, we alter the available information in the image fed to the image encoder to condition the cDDPM. To achieve this, we crop the conditioning image at a given

width of 50% and 100% where 100% indicates that the entire input image is used as the conditioning image. In addition to the qualitative evaluation of reconstructed, simulated data, we quantitatively assess the domain shift across input and reconstruction by comparing their intensity histograms. Therefore, we first calculate and plot the histograms. We partition the intensity values into 500 bins and divide the raw count by the total number of counts and the bin width. For a quantitative analysis of the distance between the intensity distributions, we calculate the Kullback–Leibler Divergence (KLD) as follows:

$$KLD = \left[- \sum_i p_{input} \log(p_{input}) \right] - \left[- \sum_i p_{reconstruction} \log(p_{reconstruction}) \right]$$

where $p = [p_1, p_2, \dots, p_n]$ represents each intensity distribution.

4.5.3. Segmentation performance

To assess the segmentation performance for the UAD task, we utilize all unhealthy test sets and consider two different segmentation metrics. First, we report the best possible Dice score ($\lceil DICE \rceil$). The formula for the Dice score is given by:

$$DICE = \frac{2 \cdot |A \cap B|}{|A| + |B|}$$

where A and B are the predicted anomaly map and the ground truth annotation, respectively. The best possible $\lceil DICE \rceil$ is estimated by a greedy search of the threshold that optimizes the Dice score, similar to [60].

Second, we calculate the Area Under Precision–Recall Curve (AUPRC) as follows:

$$AUPRC = \sum_r (R(r) - R(r-1)) \cdot P(r).$$

Here, $R(r)$ represents the recall at a given threshold or rank r , and $P(r)$ represents the precision at the corresponding recall $R(r)$. The sum is taken over all thresholds or ranks r at which the precision and recall are computed.

4.5.4. Statistical testing

To conduct significance tests, we utilize the MLxtend library's permutation test [61] with 10,000 rounds of permutations and a significance level of $\alpha = 5\%$. This test computes the mean difference of the considered scores of two models for each permutation, and the resulting p -value is computed by counting the number of times the mean differences were equal to or greater than the sample differences, divided by the total number of permutations.

5. Results

We first compare the overall reconstruction quality of healthy and unhealthy structures. Second, we evaluate the domain adaptation properties of our approach and investigate the effect of our conditioning mechanism. Lastly, we evaluate the segmentation performance for different data sets, comparing our approach to established state-of-the-art UAD methods.

5.1. Reconstruction quality

In Table 1 we compare baseline models regarding their ability to reconstruct the healthy anatomy, given the held-out test set of the IXI data set. Overall, DAEs, pDDPMs and cDDPMs show high performance regarding the image-based similarity metrics. In contrast, lower reconstruction quality is reported for the dense autoencoder-based baselines. Comparing the DDPM-based approaches, both pDDPM and cDDPM outperform the baseline DDPM in terms of reconstruction quality with statistical significance ($p < 0.05$). We also analyze the unhealthy-to-healthy error ratio ($l1$ -ratio) based on the unhealthy test

Table 1

Comparison of the reconstruction quality of the different models with the best results highlighted in bold. The asterisk * denotes superior performance with statistical significance compared to all other models ($p < 0.05$). For all metrics, the mean \pm standard deviation across the different folds are reported. The arrows \uparrow and \downarrow indicate that higher and lower values are favorable, respectively. The $l1$ -ratio is derived by dividing the $l1$ -error of unhealthy anatomy by the $l1$ -error of healthy anatomy. DDPM-based models are evaluated by ensembling different values for $t_{test} = [250, 500, 750]$.

Model	IXI (T2)				BraTS21 (T2)	MSLUB (T2)	ATLAS (T1)	WMH (T1)
	SSIM \uparrow	PSNR \uparrow	LPIPS (e-3) \downarrow	$l1$ -error (e-3) \downarrow	$l1$ -ratio \uparrow	$l1$ -ratio \uparrow	$l1$ -ratio \uparrow	$l1$ -ratio \uparrow
VAE [9]	74.98 \pm 0.54	23.38 \pm 0.14	4.03 \pm 0.50	32.32 \pm 0.64	3.52 \pm 0.08	2.92 \pm 0.06	4.43 \pm 0.03	2.36 \pm 0.04
SVAE [26]	77.87 \pm 0.15	23.94 \pm 0.06	3.31 \pm 0.24	29.08 \pm 0.16	3.90 \pm 0.05	3.13 \pm 0.05	3.38 \pm 0.11	2.07 \pm 0.01
AE [9]	76.11 \pm 0.27	23.41 \pm 0.14	3.19 \pm 0.54	31.67 \pm 0.41	3.84 \pm 0.17	3.26 \pm 0.18	4.40 \pm 0.07	2.36 \pm 0.04
RA [27]	75.46 \pm 0.35	23.92 \pm 0.23	2.18 \pm 0.41	34.36 \pm 1.43	3.10 \pm 0.16	2.56 \pm 0.11	3.93 \pm 0.25	2.42 \pm 0.19
PHANES [31]	69.04 \pm 1.23	21.39 \pm 0.32	1.08 \pm 0.09	38.70 \pm 1.74	3.54 \pm 0.13	2.73 \pm 0.07	4.01 \pm 0.07	2.20 \pm 0.05
DAE [10]	98.69 \pm 0.15*	36.69 \pm 0.38*	0.14 \pm 0.01	8.14 \pm 0.17*	7.17 \pm 0.63	2.69 \pm 0.15	4.51 \pm 0.15	2.99 \pm 0.14
DDPM [15]	93.96 \pm 0.37	31.79 \pm 0.26	0.49 \pm 0.14	14.29 \pm 0.32	6.16 \pm 0.53	3.37 \pm 0.24	5.00 \pm 0.23	3.16 \pm 0.15
pDDPM [16]	96.62 \pm 0.25	34.58 \pm 0.39	0.09 \pm 0.04*	9.70 \pm 0.43	7.16 \pm 0.15	4.34 \pm 0.13	5.58 \pm 0.28	3.00 \pm 0.16
cDDPM (Ours)	96.80 \pm 0.19	34.87 \pm 0.23	0.11 \pm 0.05	9.68 \pm 0.16	7.43 \pm 0.17	4.49 \pm 0.18	5.69 \pm 0.27	3.12 \pm 0.08

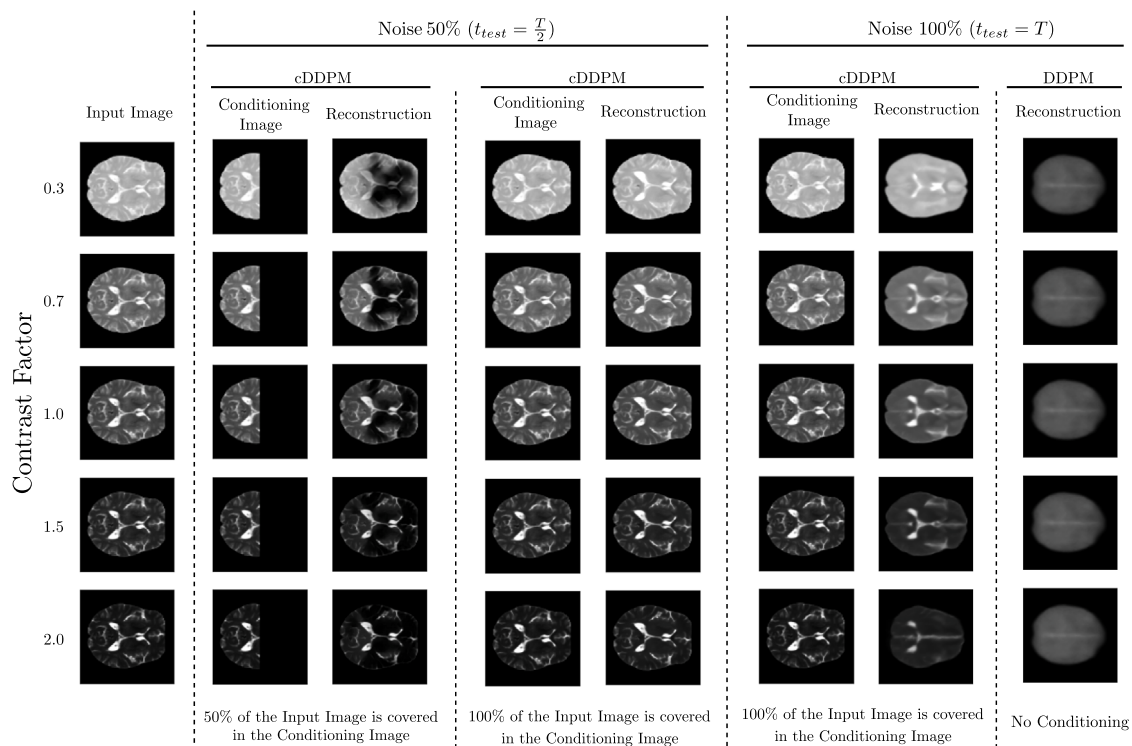


Fig. 3. Simulating the conditioning effect of the cDDPM for 50% of noise and 100% noise in the input image. In the first block, the input image that is fed to the DDPM or cDDPM is shown. In the second block, the reconstructions of cDDPM for different conditioning inputs are shown when a noise level of 50% is applied. In the third block, the reconstructions of cDDPMs and DDPMs are compared at a noise level of 100%. From top to bottom, the contrast level of the conditioning and input image is increased, respectively, for all columns.

sets. Generally, a higher $l1$ -ratio is preferable as this indicates that the model successfully reconstructs the healthy anatomy without replicating the unhealthy parts of the input. Overall, the $l1$ -ratio is highest for cDDPM across almost all data sets, except for the WMH, where DDPM shows competitive performance. While the $l1$ -ratio of DAEs is high for the BraTS data set, it is substantially lower for all other data sets, indicating limited generalization. Additional results on the $l1$ -error for all datasets are presented in Table 4 in the supplementary material.

5.2. Conditioning effect

To evaluate the effect the additional conditioning input has on individual reconstructions, we simulate different conditioning inputs, varying in the amount of used image information. Furthermore, we apply artificial contrast levels for the input images to mimic domain shifts. For each conditioning input and contrast level, we provide the

reconstructions generated by cDDPMs in Fig. 3 for a noise level of $t_{test} = 500$ (50%). Moreover, we compare the reconstruction of DDPMs and cDDPMs in the extreme case of pure noise as input ($t_{test} = T = 1000$ (100%)). From Fig. 3, it becomes evident that the overall shape of the brain to reconstruct is preserved across the different (masked) conditioning images. However, the local intensity information in the reconstruction is only captured at regions covered in the conditioning image. In the extreme case of a noise level of 100%, the reconstructions of cDDPMs still coarsely follow the intensity distribution provided by the conditioning image, leading to a blurry reconstruction of the input image. In contrast, for unconditioned DDPMs, only a generic reconstruction that shares low similarity with the given input image can be obtained. Therefore, the conditioning mechanism primarily guides the generation process of cDDPMs to maintain local intensity distributions of the input image while detailed structural information is not captured.

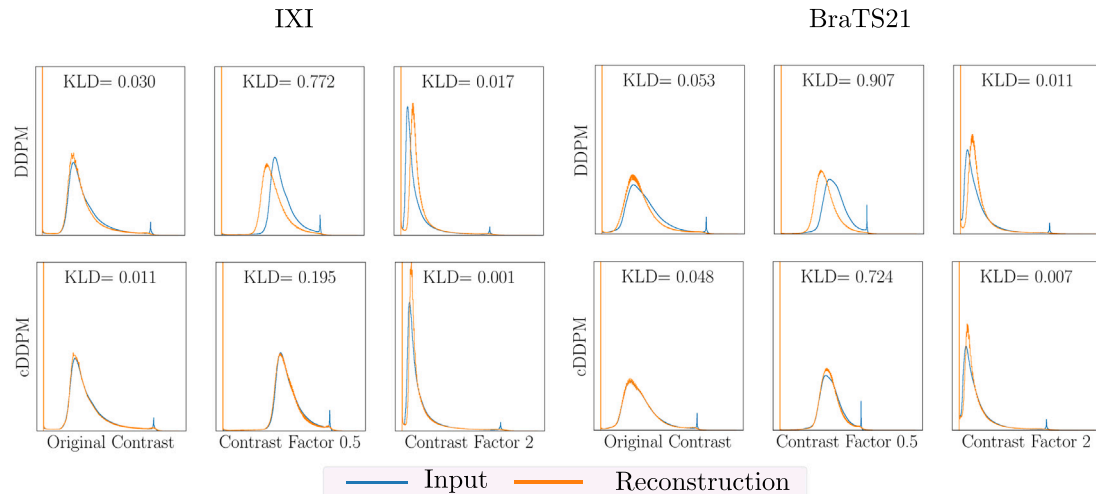


Fig. 4. Comparison of the histograms for input-reconstruction pairs of the healthy IXI (left) and the unhealthy BraTS21 (right) data set with original and augmented contrast. The top row shows the baseline DDPM without conditioning and the bottom row our proposed cDDPM with conditioning. The Kullback–Leibler divergence (KLD) for both histograms is indicated within each plot (lower values are preferable). Both models are evaluated by ensembling different values for $t_{test} = [250, 500, 750]$.

Table 2

Comparison of the evaluated models with the best results highlighted in bold. The asterisk * denotes superior performance with statistical significance compared to all other models ($p < 0.05$). For all metrics, the mean \pm standard deviation across the different folds are reported. A checkmark at SSL denotes that a pre-trained encoder is used. A checkmark at ENS denotes the ensembling of different values for $t_{test} = [250, 500, 750]$. Otherwise, a fixed value of $t_{test} = 500$ is used for DDPM-based models.

Model	Modification		BraTS21 (T2)		MSLUB (T2)		ATLAS (T1)		WMH (T1)	
	ENS	SSL	[DICE] [%]	AUPRC [%]	[DICE] [%]	AUPRC [%]	[DICE] [%]	AUPRC [%]	[DICE] [%]	AUPRC [%]
<i>Thresh</i> [53]			30.26	20.27	7.65	4.23	4.66	1.71	10.32	4.72
VAE [9]			33.12 \pm 1.12	25.74 \pm 1.37	8.10 \pm 0.18	4.48 \pm 0.18	15.63 \pm 0.73	11.44 \pm 0.5	7.60 \pm 0.28	3.86 \pm 0.40
SVAE [26]			36.43 \pm 0.36	30.3 \pm 0.45	8.55 \pm 0.11	4.8 \pm 0.09	10.32 \pm 0.53	6.84 \pm 0.44	7.18 \pm 0.07	2.97 \pm 0.06
AE [9]			36.04 \pm 1.73	28.8 \pm 1.72	9.65 \pm 0.97	5.71 \pm 0.80	14.04 \pm 0.6	10.16 \pm 0.53	7.34 \pm 0.08	3.43 \pm 0.14
DAE [10]			48.82 \pm 3.68	49.38 \pm 4.18	7.57 \pm 0.61	4.47 \pm 0.69	15.95 \pm 0.69	13.37 \pm 0.62	12.02 \pm 1.01	8.54 \pm 1.02
RA [27]			16.75 \pm 0.51	9.98 \pm 0.43	3.96 \pm 0.03	1.92 \pm 0.04	12.21 \pm 0.98	8.75 \pm 0.93	6.04 \pm 0.45	3.15 \pm 0.31
PHANES [31]			28.42 \pm 0.91	21.29 \pm 1.06	6.11 \pm 0.27	2.98 \pm 0.07	17.62 \pm 0.41	13.81 \pm 0.48	7.55 \pm 0.17	3.87 \pm 0.13
RD [21]			32.57 \pm 0.15	27.11 \pm 0.15	6.48 \pm 0.20	3.32 \pm 0.06	19.69 \pm 0.26	15.51 \pm 0.20	7.48 \pm 0.10	3.89 \pm 0.07
FAE [20]			44.59 \pm 2.19	43.63 \pm 0.47	6.85 \pm 0.65	3.85 \pm 0.08	17.76 \pm 0.16	13.91 \pm 0.10	8.81 \pm 0.38	4.77 \pm 0.26
EDC [54]			36.66 \pm 3.03	30.47 \pm 4.25	7.23 \pm 0.29	3.88 \pm 0.4	18.67 \pm 1.02	14.34 \pm 0.86	8.62 \pm 0.47	4.67 \pm 0.37
PII [55]			40.83 \pm 2.18	36.52 \pm 2.66	9.46 \pm 0.43	5.22 \pm 0.37	9.73 \pm 1.89	7.31 \pm 1.64	6.59 \pm 1.87	3.36 \pm 1.03
DDPM [15]			49.43 \pm 1.94	50.00 \pm 2.13	9.63 \pm 1.33	6.81 \pm 1.54	17.57 \pm 1.05	15.64 \pm 0.90	11.56 \pm 0.93	8.65 \pm 0.87
DDPM [15]	✓		50.27 \pm 2.67	50.61 \pm 2.92	9.71 \pm 1.29	6.27 \pm 1.58	20.18 \pm 0.58	17.77 \pm 0.47	12.06 \pm 0.97	8.89 \pm 0.89
pDDPM [16]			53.25 \pm 0.50	54.73 \pm 0.52	12.40 \pm 0.36	10.14 \pm 0.44	19.20 \pm 0.45	17.31 \pm 0.34	10.14 \pm 0.50	7.78 \pm 0.56
pDDPM [16]	✓		53.61 \pm 0.51	55.08 \pm 0.54	12.83 \pm 0.40	10.02 \pm 0.36	19.92 \pm 0.24	17.84 \pm 0.10	10.13 \pm 0.53	7.52 \pm 0.56
cDDPM (Ours)			54.49 \pm 1.63	56.81 \pm 1.96	12.79 \pm 1.08	10.07 \pm 1.07	22.6 \pm 1.67	20.65 \pm 1.52	11.21 \pm 0.54	9.05 \pm 0.56
cDDPM (Ours)		✓	55.67 \pm 1.05	58.14 \pm 1.28	13.52 \pm 0.91	10.89 \pm 1.08	22.66 \pm 1.20	20.85 \pm 1.28	11.15 \pm 0.8	9.03 \pm 0.90
cDDPM (Ours)	✓	✓	56.30 \pm 1.25*	58.82 \pm 1.56*	14.04 \pm 1.16	10.97 \pm 1.17	24.22 \pm 1.10*	22.22 \pm 1.15*	11.59 \pm 0.93	9.26 \pm 1.07

5.3. Domain adaptation

To evaluate the domain adaptation to different data sets, we consider the healthy IXI data set as an in-domain data set and the unhealthy BraTS21 data set as an out-of-domain one. Note that for the BraTS21 data set, we only consider regions annotated as healthy. Thereby, we ensure the evaluation of domain shifts regarding scanner and brain diversity, not domain shifts arising from unhealthy brain MRI structures. In Fig. 4, we provide the histograms of input and reconstructions of the healthy IXI data set (left) and the unhealthy BraTS21 data set (right). We observe that DDPMs show substantial discrepancies across input and reconstruction intensity distributions. Particularly for simulated contrast levels, the histograms deviate. In contrast, the intensity distribution of images reconstructed by cDDPMs exhibits higher similarities with the input intensity distribution for in- and out-of-domain data. Compared to DDPMs, cDDPMs decrease the KLD by a factor of 2.3, 4.0 and 17.0 across the contrast factors, respectively, considering the IXI data set.

5.4. Segmentation performance

We report the segmentation performance in Table 2 considering all unhealthy test sets. For cDDPMs, improved performance is reported compared to all baselines across all data sets, except for the WMH data set, where the performance is on par with DDPMs. While the improvements for the cDDPM are statistically significant for the BraTS21 and ATLAS data sets ($p < 0.05$), no significant difference can be observed for the MSLUB and WMH data sets. Furthermore, we report enhanced performance of cDDPMs when pre-training the encoder (SSL checkmark in Table 2) and ensembling the reconstructions of different noise levels (ENS checkmark in Table 2) for most data sets. Notably, the inference time of cDDPMs is reduced by $\sim 37\%$ compared to pDDPMs and increased by $\sim 2\%$ compared to DDPMs.

The reconstruction-based AEs, (S)VAEs, RA and PHANES as well as the self-supervised PII and the feature-based EDC are outperformed by a margin when compared to DDPM-based approaches. While the feature-based methods RD and FAE and the reconstruction-based DAEs show competitive performance considering individual data sets, limited

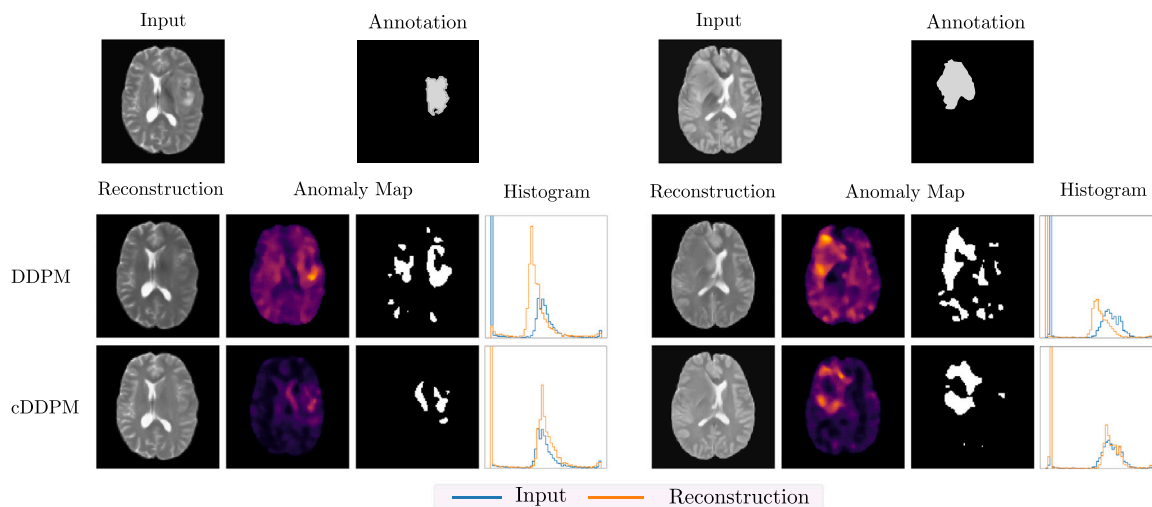


Fig. 5. Exemplary reconstructions and anomaly maps for DDPMs (second row) and cDDPMs (third row). The input and the corresponding ground truth annotation are provided in the first row. For each case, the reconstruction, the anomaly map and the histograms of intensity values in input and reconstruction are shown. Note that for histogram calculation, only healthy areas are considered. For visualization purposes, we provide segmentation maps next to the anomaly maps. We derive the binarization threshold by optimizing for the best possible dice score.




Fig. 6. Comparison of different noise levels t_{test} regarding the AUPRC. Top row: MSLUB (left) and BraTS21 (right) data sets. Bottom row: ATLAS (left) and WMH data set. The superscript *ens* denotes the ensembling of reconstructions from different noise levels t_{test} in [250, 500, 750].

generalization across all evaluated data sets is shown. For a qualitative assessment of the conditioning mechanism, we provide reconstructions and the corresponding anomaly maps in Fig. 5, comparing unconditioned DDPMs to cDDPMs. cDDPMs provide accurate reconstructions of the target image with aligned intensity distributions across input and reconstructions. In contrast, the reconstruction of DDPMs shows less details with a slight intensity shift. Hence, the anomaly map of cDDPMs shows a higher contrast across normal and abnormal regions, which facilitates the delineation of the present pathology.

In Fig. 6 we provide an ablation study considering different noise levels applied during the diffusion process at test time. The achieved

AUPRC score is dependent on the noise level. Applying the ensembling strategy reduces this dependency and achieves consistent performance without specifying an individual noise level.

We supply a collection of exemplary residual maps for different baseline methods in Fig. 7. Compared to the baselines, our proposed cDDPM demonstrates the low occurrences of false positives and a pronounced contrast in the residual maps. This observation is further supported by Figure 8 in the supplementary material.

Additionally, we include ablation studies on the applied post-processing steps in the supplementary material.

6. Discussion

UAD in brain MRI has gained significant attention due to its potential to identify abnormalities without costly data annotation. Compared to supervised approaches that rely on annotated data sets, UAD methods take a different approach by learning the underlying data distribution of healthy brain anatomy and identifying anomalies as outliers. This is a crucial property for screening tasks, where any pathology has to be detected, even if it is not represented in an annotated training set.

In this study, we focus on reconstruction-based UAD with DDPMs. DDPMs generate images by reconstructing an input corrupted by noise, leveraging the high-dimensional latent space to achieve high-fidelity reconstructions of fine-grained structures. However, we show that the forward and backward processes of DDPMs do not sufficiently capture the highly variable intensity characteristics of MRI scans. This results in differences between input and reconstructions that are difficult to distinguish from differences that arise from actual pathologies. This limitation becomes especially prominent in the presence of domain shifts at test time.

To address this challenge, we propose conditioned DDPMs (cDDPMs) for UAD in brain MRI. We train a DDPM to reconstruct healthy brain anatomy and condition the denoising process by a latent feature representation of the input image derived by an additional image encoder. While the additional dense feature representation can capture local intensity information of the image to reconstruct, it is unsuitable for the reconstruction of detailed structures [9,27]. This is important, as providing detailed structural information of the unhealthy input image could lead to copying abnormal structures, which would prevent the detection of pathologies. As demonstrated in our results, our approach facilitates the effective utilization of the conditioning signal, contributing to improved reconstructions that adapt locally to the input intensity distribution. Therefore, we observe enhanced domain adaptation capabilities to both real and simulated intensity profiles with our conditioning mechanism. Moreover, our results indicate that the additional conditioning signal does not support the replication of unhealthy structures. Finally, these individual features of our approach lead to an improved segmentation performance across various data sets.

In the following, we systematically evaluate our approach regarding reconstruction quality, domain adaptation, and segmentation performance based on five different data sets.

6.1. Reconstruction quality

We compare the reconstruction quality of our method with baseline models on the healthy IXI data set in Table 1. For the AE and (S)VAE, overall the worst reconstruction quality is reported. A reason for this is seen in the strict bottleneck enforced by the dense latent space as it inhibits information flow [9]. In contrast, methods like DAE or DDPMs that are not constrained by a dense latent space but by a denoising task [10] show improved reconstruction performance. While pDDPMs and cDDPMs outperform the baseline DDPM in terms of reconstruction quality, we observe that all models are outperformed by the DAE. We note that the overall training objective of the compared generative models is to reconstruct the image with high accuracy and copying the input image would be a trivial solution. However, for the UAD task, it is crucial that the given input image is not solely copied but that pseudo-healthy representatives replace unhealthy anatomy. Hence, comparing only the reconstruction quality of healthy anatomy does not necessarily reflect the usefulness of the UAD task. Therefore, we utilize the l_1 -ratio where high values indicate a better trade-off between the reconstruction of healthy and unhealthy anatomy and vice-versa. While DDPMs and particularly the cDDPM achieve a high l_1 -ratio, across all unhealthy data sets, it becomes evident that the DAE fails to generalize to different pathology types, a crucial property of UAD methods. A reason for that is seen in the chosen noise type in DAE that mimics the visual appearance of tumors [27,60]. In summary, the

cDDPM shows improved reconstruction quality compared to DDPMs while preserving a high l_1 -ratio. We conclude that the conditioning mechanism effectively captures intensity information from the input image for an improved reconstruction without providing too much detailed structural information that would enable the cDDPM to solely copy the input image.

6.2. Domain adaptation

We evaluate the domain adaptation capabilities of cDDPMs by simulating different contrast levels and conditioning inputs in Fig. 3. The reconstructions show that while the overall shape is preserved across different conditioning masks, meaningful reconstructions are achieved only in regions covered by the conditioning image. Particularly, the conditioning image is critical in capturing local intensity information, demonstrating the ability of cDDPMs to adapt to different contrast levels and effectively capture intensity information. This becomes even more evident when high noise levels are considered (100%), where the only source of information concerning the given input image is the conditioning image. Here, the reconstruction becomes totally dependent on the shape and intensity characteristics of the conditioning image. The conditioning facilitates a blurred reconstruction of prominent local intensity changes. This indicates that the dense latent representation used for conditioning guides the reconstruction with local intensity details from the input image, while detailed structural information is not captured. Similar results are reported in the literature comparing AE architectures with dense and spatial latent spaces [9]. Furthermore, these findings indicate that cDDPMs effectively learn to balance information from the noisy input image and the conditioning encoder features during training, adapting according to the input's noise level.

We explore the domain adaptation capabilities of cDDPMs in real-world scenarios where a different, out-of-domain data set is used for testing. To assess the domain adaptation ability, we investigate the deviation between the intensity distributions of the input and reconstruction by plotting histograms and calculating the Kullback–Leibler Divergence (KLD) as a proxy in Fig. 4. Our results show that cDDPMs exhibit improved performance in capturing and estimating the intensity distribution. Particularly when simulating contrast levels considerably different from the training distribution, cDDPMs demonstrate superior alignment of the histograms and lower KLD values compared to unconditioned DDPMs. This analysis highlights the potential of the conditioning mechanism in cDDPMs to effectively adapt to unseen variations in intensity profiles and improve the coherence between input and reconstruction, which can contribute to improved domain adaptation.

6.3. Segmentation performance

Overall, cDDPMs demonstrate competitive or superior results compared to traditional autoencoder-based approaches, as well as the baseline DDPM and pDDPM, as presented in Table 2. Additionally, the feature-based and self-supervised methods are outperformed by our cDDPM. We conclude that the enhanced reconstruction quality, global and local intensity information capture, and effective domain adaptation capabilities attributed to our conditioning approach contribute to strong segmentation performance. Furthermore, we observe that pre-training the encoder F_{enc} slightly enhances the segmentation performance in most cases, indicating that starting from an already learned representation space has the potential to improve the overall integration of the conditioning features, compared to simultaneously training the parameters of both, DDPM and F_{enc} from scratch. Our results demonstrate improved performance of cDDPMs over pDDPMs, indicating that the proposed conditioning mechanism is more effective compared to the patching strategy in pDDPMs. A reason for this is seen in potential artifacts, introduced by the patching strategy. Furthermore, cDDPMs reduce complexity and inference time as there is no need for a




Fig. 7. Exemplary residual maps from the BraTS21 data set (rows 1 and 2), the WMH data set (rows 3 and 4) and the ATLAS data set (rows 5 and 6). The input image, residual maps, and ground truth are shown from left to right. Corresponding segmentation maps are provided in Figure 8.

costly patching strategy, making them a practical and efficient solution for UAD in brain MRI. Despite the superior reconstruction quality of the *DAE* on healthy data, *cDDPMs* show stronger segmentation results given the final UAD task. A reason for this is seen in the *DAEs*' ability to reconstruct unhealthy anatomy, particularly for pathologies differing from the appearance of tumors, as discussed in Section 6.1. Notably, the performance of *AEs* and *VAEs* in this work is lower compared to previous studies such as [9]. This discrepancy can be attributed to differences in the MRI sequences used. Unlike [9], which uses FLAIR images, we rely on T2- and T1-weighted images. Previous research has shown that hyperintense lesions in FLAIR images can often be detected using simple thresholding [53]. This indicates that the high performance of *AEs* and *VAEs* in [9] may be largely due to the hyperintensity of the lesions, rather than the model's reconstruction ability.

To further analyze the effect of the conditioning mechanism, we compare reconstructions and anomaly maps of *DDPMs* and *cDDPMs* in Fig. 5. In contrast to *DDPMs*, *cDDPMs*' reconstructions follow the local

intensity information of the respective input images. This is in line with the observations in Fig. 4, where a lower distribution shift is reported for *cDDPMs*. This leads to a reduction of false positives, facilitating the delineation of anomalies.

In Fig. 6, we explore the impact of noise levels on the segmentation performance. We demonstrate that *cDDPMs* outperform the baseline models across different noise levels for most data sets. However, we also observe that the noise level is a crucial hyper-parameter. High noise levels tend to result in more blurry and generic reconstructions, whereas low noise levels enable sharper reconstructions, including unhealthy anatomy. Thus, selecting an appropriate noise level is essential to achieve reasonable performance. However, the optimal value for the applied noise depends on the evaluated data set, as shown in Fig. 6. We assume that the main reason for this dependency is the different sizes of pathologies in the considered data sets, as also stated in [33]. We apply different noise levels and average the resulting reconstructions to address this dependency. Thereby, we utilize complementary information

of different reconstructions and effectively mitigate the dependency on the noise level, which enhances the model's generalization abilities.

6.4. Limitations and future work

Despite showing promising results, our approach has limitations. Specifically, the segmentation performance falls below that of supervised algorithms. However, compared to supervised networks, UAD methods offer a crucial advantage in detecting any pathologies unseen during training. While additional efforts are required for the practical clinical application of UAD methods, our approach demonstrates superior performance to state-of-the-art UAD methods, adding a valuable contribution to the field of UAD in brain MRI.

Another avenue for future work is the incorporation of multi-scale image encodings into our conditioning mechanism. Our study does not utilize multi-scale analysis, which could be advantageous in capturing fine-grained details of the intensity patterns and contextual information at different resolutions. By carefully integrating multi-scale image encodings without allowing a copy of the conditioning image, we see the potential to enhance the performance of our cDDPMs in capturing both global and local features of the input images.

We conduct our studies using a downsampled resolution. To improve performance, especially for detecting smaller lesions, future work could explore and evaluate our proposed method at higher resolutions. A promising approach to balance the increased computational cost is the use of latent diffusion models [62]. To further enhance the conditioning mechanism, an additional direction for future work is to explore the use of 3D image encoders. In earlier studies, we have shown that 3D information can improve the reconstruction quality for VAEs [25,26] and expect similar improvements for DDPMs. Currently, our approach operates on 2D slices of the MRI data, which may limit the preservation of 3D context and spatial relationships between slices. By utilizing a 3D encoder to condition the 2D DDPM, we can potentially capture and preserve 3D contextual information.

7. Conclusion

UAD in brain MRI presents a promising alternative to supervised models, especially considering clinical screening tasks. DDPMs have demonstrated their utility for UAD, largely due to their high reconstruction accuracy of fine-grained structures. However, accurately reconstructing the intensity characteristics of a given MRI scan remains a challenge, especially when facing domain shifts. To address this, we propose cDDPMs and introduce a conditioning mechanism that incorporates an additional feature representation of the input image into the DDPMs' denoising process. Our findings indicate that this conditioning mechanism effectively addresses challenges of DDPMs regarding capturing accurate intensity capture and domain adaptation. As a result, our approach outperformed state-of-the-art architectures for UAD in brain MRI on various publicly available data sets. Our work addresses challenges of applying DDPMs for UAD in brain MRI and has practical implications for detecting and segmenting pathologies in scenarios where domain shifts are likely.

CRediT authorship contribution statement

Finn Behrendt: Writing – original draft, Software, Methodology, Conceptualization. **Debayan Bhattacharya:** Writing – review & editing, Methodology, Conceptualization. **Robin Mieling:** Writing – review & editing, Methodology, Conceptualization. **Lennart Maack:** Writing – review & editing, Methodology, Conceptualization. **Julia Krüger:** Writing – review & editing, Funding acquisition, Conceptualization. **Roland Opfer:** Writing – review & editing, Funding acquisition, Conceptualization. **Alexander Schlaefer:** Writing – review & editing, Supervision, Funding acquisition, Conceptualization.

Ethics statement

This study was conducted using publicly available, fully anonymized datasets, and no new data were collected or experiments involving human subjects were performed. All procedures adhered to relevant laws, institutional guidelines, and ethical standards for research.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was partially funded by the “Zentrales Innovationsprogramm Mittelstand, Germany” [grant numbers KK5208101KS0 and ZF4026303TS9] and by the Free and Hanseatic City of Hamburg (Interdisciplinary Graduate School), Germany.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.combiomed.2025.109660>.

References

- [1] M.W. Vernooij, M.A. Ikram, H.L. Tanghe, A.J. Vincent, A. Hofman, G.P. Krestin, W.J. Niessen, M.M. Breteler, A. van der Lugt, Incidental findings on brain MRI in the general population, *N. Engl. J. Med.* 357 (18) (2007) 1821–1828.
- [2] A.S. Lundervold, A. Lundervold, An overview of deep learning in medical imaging focusing on MRI, *Z. Med. Phys.* 29 (2) (2019) 102–127, <http://dx.doi.org/10.1016/j.zemedi.2018.11.002>, URL <https://www.sciencedirect.com/science/article/pii/S0939388918301181>, Special Issue: Deep Learning in Medical Physics.
- [3] M.A. Bruno, E.A. Walker, H.H. Abujudeh, Understanding and confronting our mistakes: the epidemiology of error in radiology and strategies for error reduction, *Radiographics* 35 (6) (2015) 1668–1676.
- [4] R.J. McDonald, K.M. Schwartz, L.J. Eckel, F.E. Diehn, C.H. Hunt, B.J. Bartholmai, B.J. Erickson, D.F. Kallmes, The effects of changes in utilization and technological advancements of cross-sectional imaging on radiologist workload, *Academic Radiol.* 22 (9) (2015) 1191–1198.
- [5] M. Perkuhn, P. Stavrinou, F. Thiele, G. Shakirin, M. Mohan, D. Garpis, C. Kabbasch, J. Borggrefe, Clinical evaluation of a multiparametric deep learning model for glioblastoma segmentation using heterogeneous magnetic resonance imaging data from clinical routine, *Invest. Radiol.* 53 (11) (2018) 647.
- [6] P. Moeskops, J. de Bresser, H.J. Kuijff, A.M. Mendrik, G.J. Biessels, J.P. Pluim, I. Išgum, Evaluation of a deep learning approach for the segmentation of brain tissues and white matter hyperintensities of presumed vascular origin in MRI, *NeuroImage: Clin.* 17 (2018) 251–262.
- [7] J. Islam, Y. Zhang, Brain MRI analysis for alzheimer's disease diagnosis using an ensemble system of deep convolutional neural networks, *Brain Informat.* 5 (2018) 1–14.
- [8] A. Jagodzinski, C. Johansen, U. Koch-Gromus, G. Aarabi, G. Adam, S. Anders, M. Augustin, R.B. der Kellen, T. Beikler, C.-A. Behrendt, et al., Rationale and design of the hamburg city health study, *Eur. J. Epidemiol.* 35 (2020) 169–181.
- [9] C. Baur, S. Denner, B. Wiestler, N. Navab, S. Albarqouni, Autoencoders for unsupervised anomaly segmentation in brain MR images: a comparative study, *Med. Image Anal.* (2021) 101952.
- [10] A. Kascenas, N. Pugeault, A.Q. O'Neil, Denoising autoencoders for unsupervised anomaly detection in brain MRI, in: *Medical Imaging with Deep Learning*, in: *Proceedings of Machine Learning Research*, PMLR, 2022.
- [11] X. Chen, S. You, K.C. Tezcan, E. Konukoglu, Unsupervised lesion detection via image restoration with a normative prior, *Med. Image Anal.* 64 (2020) 101713.
- [12] W.H. Pinaya, P.-D. Tudosiu, R. Gray, G. Rees, P. Nachev, S. Ourselin, M.J. Cardoso, Unsupervised brain imaging 3D anomaly detection and segmentation with transformers, *Med. Image Anal.* 79 (2022) 102475.
- [13] J. Ho, A. Jain, P. Abbeel, Denoising diffusion probabilistic models, *NIPS* 33 (2020) 6840–6851.
- [14] W.H. Pinaya, M.S. Graham, R. Gray, P.F. Da Costa, P.-D. Tudosiu, P. Wright, Y.H. Mah, A.D. MacKinnon, J.T. Teo, R. Jager, et al., Fast unsupervised brain anomaly detection and segmentation with diffusion models, 2022, arXiv preprint [arXiv:2206.03461](https://arxiv.org/abs/2206.03461).

- [15] J. Wyatt, A. Leach, S.M. Schmon, C.G. Willcocks, AnoDDPM: Anomaly detection with denoising diffusion probabilistic models using simplex noise, in: CVPR, 2022, pp. 650–656.
- [16] F. Behrendt, D. Bhattacharya, J. Krüger, R. Opfer, A. Schlaefer, Patched diffusion models for unsupervised anomaly detection in brain MRI, in: Medical Imaging with Deep Learning, PMLR, 2024, pp. 1019–1032.
- [17] C. Baur, B. Wiestler, S. Albarqouni, N. Navab, Bayesian skip-autoencoders for unsupervised hyperintense anomaly detection in high resolution brain MRI, in: IEEE ISBI, IEEE, 2020, pp. 1905–1909.
- [18] C. Baur, B. Wiestler, S. Albarqouni, N. Navab, Scale-space autoencoders for unsupervised anomaly segmentation in brain MRI, in: Computer Assisted Radiology and Surgery, Springer, 2020, pp. 552–561.
- [19] J. Silva-Rodríguez, V. Naranjo, J. Dolz, Constrained unsupervised anomaly segmentation, *Med. Image Anal.* 80 (2022) 102526.
- [20] F. Meissen, J. Paetzold, G. Kaissis, D. Rueckert, Unsupervised anomaly localization with structural feature-autoencoders, 2022, arXiv preprint arXiv:2208.10992.
- [21] H. Deng, X. Li, Anomaly detection via reverse distillation from one-class embedding, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR, 2022, pp. 9737–9746.
- [22] F. Behrendt, M. Bengs, F. Rogge, J. Krüger, R. Opfer, A. Schlaefer, Unsupervised anomaly detection in 3D brain MRI using deep learning with impured training data, in: 2022 IEEE 19th International Symposium on Biomedical Imaging, ISBI, IEEE, 2022, pp. 1–4.
- [23] D. Zimmerer, F. Isensee, J. Petersen, S. Kohl, K. Maier-Hein, Unsupervised anomaly localization using variational auto-encoders, in: D. Shen, T. Liu, T.M. Peters, L.H. Staib, C. Essert, S. Zhou, P.-T. Yap, A. Khan (Eds.), Medical Image Computing and Computer Assisted Intervention – MICCAI 2019, Springer International Publishing, Cham, 2019, pp. 289–297.
- [24] D. Zimmerer, S. Kohl, J. Petersen, F. Isensee, K. Maier-Hein, Context-encoding variational autoencoder for unsupervised anomaly detection, in: Medical Imaging with Deep Learning, 2019.
- [25] M. Bengs, F. Behrendt, J. Krüger, R. Opfer, A. Schlaefer, Three-dimensional deep learning with spatial erasing for unsupervised anomaly segmentation in brain MRI, *Comput. Assist. Radiol. Surg.* 16 (9) (2021) 1413–1423.
- [26] F. Behrendt, M. Bengs, D. Bhattacharya, J. Krüger, R. Opfer, A. Schlaefer, Capturing inter-slice dependencies of 3D brain MRI-scans for unsupervised anomaly detection, in: Medical Imaging with Deep Learning, 2022, URL <https://openreview.net/forum?id=db8wDgKH4p4>.
- [27] C.I. Bercea, B. Wiestler, D. Rueckert, J.A. Schnabel, Generalizing unsupervised anomaly detection: Towards unbiased pathology screening, in: Medical Imaging with Deep Learning, 2023, URL <https://openreview.net/forum?id=8ojx-Ld3yJR>.
- [28] C. Han, L. Rundo, K. Murao, T. Noguchi, Y. Shimahara, Z.Á. Milacski, S. Koshino, E. Sala, H. Nakayama, S. Satoh, MADGAN: Unsupervised medical anomaly detection GAN using multiple adjacent brain MRI slice reconstruction, *BMC Bioinformatics* 22 (2) (2021) 1–20.
- [29] T. Schlegl, P. Seeböck, S.M. Waldstein, G. Langs, U. Schmidt-Erfurth, f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks, *Med. Image Anal.* 54 (2019) 30–44.
- [30] C. Baur, B. Wiestler, S. Albarqouni, N. Navab, Deep autoencoding models for unsupervised anomaly segmentation in brain MR images, in: MICCAI Brainlesion Workshop, Springer, 2018, pp. 161–169.
- [31] C.I. Bercea, B. Wiestler, D. Rueckert, J.A. Schnabel, Reversing the abnormal: Pseudo-healthy generative networks for anomaly detection, in: Medical Image Computing and Computer Assisted Intervention – MICCAI23, vol. 14224, 2023, pp. 293–303, http://dx.doi.org/10.1007/978-3-031-43904-9_29.
- [32] M.S. Graham, W.H. Pinaya, P.-D. Tudosiu, P. Nachev, S. Ourselin, M.J. Cardoso, Denoising diffusion models for out-of-distribution detection, 2022, arXiv preprint arXiv:2211.07740.
- [33] C.I. Bercea, M. Neumayr, D. Rueckert, J.A. Schnabel, Mask, stitch, and re-sample: Enhancing robustness and generalizability in anomaly detection through automatic diffusion models, in: ICML 3rd Workshop on Interpretable Machine Learning in Healthcare, IMLH, 2023, URL <https://openreview.net/forum?id=kTpaFpXrqa>.
- [34] P. Dhariwal, A. Nichol, Diffusion models beat gans on image synthesis, *NIPS* 34 (2021) 8780–8794.
- [35] C. Saharia, W. Chan, H. Chang, C. Lee, J. Ho, T. Salimans, D. Fleet, M. Norouzi, Palette: Image-to-image diffusion models, in: ACM, 2022, pp. 1–10.
- [36] T. Wang, T. Zhang, B. Zhang, H. Ouyang, D. Chen, Q. Chen, F. Wen, Pretraining is all you need for image-to-image translation, 2022, arXiv preprint arXiv:2205.12952.
- [37] J. Wolleb, F. Bieder, R. Sandkühler, P.C. Cattin, Diffusion models for medical anomaly detection, 2022, arXiv preprint arXiv:2203.04306.
- [38] P. Sanchez, A. Kascenas, X. Liu, A.Q. O’Neil, S.A. Tsafaris, What is healthy? generative counterfactual diffusion for lesion localization, in: MICCAI Workshop on Deep Generative Models, Springer, 2022, pp. 34–44.
- [39] X. Zhang, N. Li, J. Li, T. Dai, Y. Jiang, S.-T. Xia, Unsupervised surface anomaly detection with diffusion probabilistic model, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2023, pp. 6782–6791.
- [40] A. Mousakhan, T. Brox, J. Tayyub, Anomaly detection with conditioned denoising diffusion models, 2023, arXiv preprint arXiv:2305.15956.
- [41] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: Medical Image Computing and Computer Assisted Intervention, Springer, 2015, pp. 234–241.
- [42] E. Perez, F. Strub, H. De Vries, V. Dumoulin, A. Courville, Film: Visual reasoning with a general conditioning layer, in: AAAI, Vol. 32, No. 1, 2018.
- [43] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, R. Girshick, Masked autoencoders are scalable vision learners, in: CVPR, 2022, pp. 16000–16009.
- [44] K. Tian, Y. Jiang, Q. Diao, C. Lin, L. Wang, Z. Yuan, Designing BERT for convolutional networks: Sparse and hierarchical masked modeling, 2023, arXiv:2301.03580.
- [45] U. Baid, S. Ghodasara, S. Mohan, M. Bilello, E. Calabrese, E. Colak, K. Farahani, J. Kalpathy-Cramer, F.C. Kitamura, S. Pati, et al., The rsna-asnr-miccai brats 2021 benchmark on brain tumor segmentation and radiogenomic classification, 2021, arXiv preprint arXiv:2107.02314.
- [46] B.H. Menze, A. Jakab, S. Bauer, J. Kalpathy-Cramer, K. Farahani, J. Kirby, Y. Burren, N. Porz, J. Slotboom, R. Wiest, et al., The multimodal brain tumor image segmentation benchmark (BRATS), *IEEE Trans. Med. Imaging* 34 (10) (2014) 1993–2024.
- [47] S. Bakas, H. Akbari, A. Sotiras, M. Bilello, M. Rozycki, J.S. Kirby, J.B. Freymann, K. Farahani, C. Davatzikos, Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features, *Sci. Data* 4 (1) (2017) 1–13.
- [48] Ž. Lesjak, A. Galimzianova, A. Koren, M. Lukin, F. Pernuš, B. Likar, Ž. Špiclin, A novel public MR image dataset of multiple sclerosis patients with lesion segmentations based on multi-rater consensus, *Neuroinformatics* 16 (1) (2018) 51–63.
- [49] S.-L. Liew, B.P. Lo, M.R. Donnelly, A. Zavaliangos-Petropulu, J.N. Jeong, G. Barisano, A. Hutton, J.P. Simon, J.M. Juliano, A. Suri, et al., A large, curated, open-source stroke neuroimaging dataset to improve lesion segmentation algorithms, *Sci. Data* 9 (1) (2022) 320.
- [50] H.J. Kuijff, J.M. Biesbroek, J. De Bresser, R. Heinen, S. Andermatt, M. Bento, M. Berseht, M. Belyaev, M.J. Cardoso, A. Casamitjana, et al., Standardized assessment of automatic segmentation of white matter hyperintensities and results of the WMH segmentation challenge, *IEEE Trans. Med. Imaging* 38 (11) (2019) 2556–2568.
- [51] F. Isensee, M. Schell, I. Pflueger, G. Brugnara, D. Bonekamp, U. Neuberger, A. Wick, H.-P. Schlemmer, S. Heiland, W. Wick, et al., Automated brain extraction of multisequence MRI using artificial neural networks, *Hum. Brain Mapp.* 40 (17) (2019) 4952–4964.
- [52] N.J. Tustison, B.B. Avants, P.A. Cook, Y. Zheng, A. Egan, P.A. Yushkevich, J.C. Gee, N4ITK: improved N3 bias correction, *IEEE Trans. Med. Imaging* 29 (6) (2010) 1310–1320.
- [53] F. Meissen, G. Kaissis, D. Rueckert, Challenging current semi-supervised anomaly segmentation methods for brain MRI, in: MICCAI Brainlesion Workshop, Springer, 2022, pp. 63–74.
- [54] J. Guo, S. Lu, L. Jia, W. Zhang, H. Li, Encoder-decoder contrast for unsupervised anomaly detection in medical images, *IEEE Trans. Med. Imaging* 43 (3) (2023) <http://dx.doi.org/10.1109/TMI.2023.3327720>.
- [55] J. Tan, B. Hou, T. Day, J. Simpson, D. Rueckert, B. 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, Springer, 2021, pp. 581–591.
- [56] F. Pérez-García, R. Sparks, S. Ourselin, Torchio: A python library for efficient loading, preprocessing, augmentation and patch-based sampling of medical images in deep learning, *Comput. Methods Programs Biomed.* 208 (2021) 106236.
- [57] Z. Wang, A. Bovik, H. Sheikh, E. Simoncelli, Image quality assessment: from error visibility to structural similarity, *IEEE Trans. Image Process.* 13 (4) (2004) 600–612, <http://dx.doi.org/10.1109/TIP.2003.819861>.
- [58] R. Zhang, P. Isola, A.A. Efros, E. Shechtman, O. Wang, The unreasonable effectiveness of deep features as a perceptual metric, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 586–595.
- [59] S. Chen, K. Ma, Y. Zheng, Med3d: Transfer learning for 3d medical image analysis, 2019, arXiv preprint arXiv:1904.00625.
- [60] I. Lagogiannis, F. Meissen, G. Kaissis, D. Rueckert, Unsupervised pathology detection: A deep dive into the state of the art, *IEEE Trans. Med. Imaging* PP (2023) <http://dx.doi.org/10.1109/TMI.2023.3298093>.
- [61] S. Raschka, Mlxtend: Providing machine learning and data science utilities and extensions to python’s scientific computing stack, *J. Open Source Softw.* 3 (24) (2018).
- [62] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, B. Ommer, High-resolution image synthesis with latent diffusion models, in: CVPR, 2022, pp. 10684–10695.