

Leveraging the Mahalanobis Distance to enhance Unsupervised Brain MRI Anomaly Detection

Finn Behrendt¹[0000–0001–7191–6508], Debayan Bhattacharya¹, Robin Mieling¹[0000–0003–0262–2519], Lennart Maack¹, Julia Krüger², Roland Opfer²[0000–0002–9911–5478], and Alexander Schlaefer¹[0000–0001–9201–8854]

¹ Institute of Medical Technology and Intelligent Systems, Hamburg University of Technology, Hamburg, Germany
² Jung Diagnostics, Hamburg, Germany

Abstract. Unsupervised Anomaly Detection (UAD) methods rely on healthy data distributions to identify anomalies as outliers. In brain MRI, a common approach is reconstruction-based UAD, where generative models reconstruct healthy brain MRIs, and anomalies are detected as deviations between input and reconstruction. However, this method is sensitive to imperfect reconstructions, leading to false positives that impede the segmentation. To address this limitation, we construct multiple reconstructions with probabilistic diffusion models. We then analyze the resulting distribution of these reconstructions using the Mahalanobis distance to identify anomalies as outliers. By leveraging information about normal variations and covariance of individual pixels within this distribution, we effectively refine anomaly scoring, leading to improved segmentation. Our experimental results demonstrate substantial performance improvements across various data sets. Specifically, compared to relying solely on single reconstructions, our approach achieves relative improvements of 15.9%, 35.4%, 48.0%, and 4.7% in terms of AUPRC for the BRATS21, ATLAS, MSLUB and WMH data sets, respectively.

Keywords: Unsupervised Anomaly Detection · Diffusion Models · Mahalanobis Distance

1 Introduction

Deep learning (DL) methods show promise in tasks like the segmentation of brain pathologies in magnetic resonance imaging (MRI) scans [19]. However, supervised DL methods require pixel-level annotations for training. This requirement becomes a challenge, particularly for screening tasks, where any pathology has to be detected even if not represented in the training data. Unsupervised Anomaly Detection (UAD) offers an alternative approach by learning the distribution of healthy data and identifying anomalies as outliers. A prevalent strategy is using reconstruction-based techniques [2]. These methods train generative models (GM) on a data set composed solely of healthy brain MRI scans. The underlying assumption is that the GMs will fail to reconstruct anomalies or pathological

structures not present in the training data set. Therefore, anomaly maps for segmenting abnormal structures can be derived from the deviations between input and reconstruction. However, a critical challenge UAD methods face lies in their high sensitivity to errors stemming from imperfect reconstructions [23, 21, 10]. As a result, even healthy structures exhibit deviations in the anomaly map. Therefore, discriminating deviations caused by genuine pathologies from those arising due to imperfect reconstructions becomes challenging, leading to false positives in the final segmentation. While deviations from imperfect reconstructions are inevitable, analyzing multiple reconstructions of the same input can offer valuable insights into the normal variations within the distribution of pseudo-healthy reconstructions, potentially simplifying the discrimination. These multiple reconstructions can be sampled using probabilistic GMs. Previous approaches have primarily focused on comparing the average reconstruction with the corresponding input image [3, 2]. However, these approaches ignore the valuable information in the variance and covariance of pixels across different reconstructions. The inter-pixel covariance across reconstructions quantifies the relationship between pixel values at different locations. It can be utilized to achieve a more balanced decision when measuring the distance of individual input pixels to the pseudo-healthy distribution of healthy pixels. Therefore, we propose using the Mahalanobis distance (MHD) [20] to measure the divergence of pixels in the input image from the pseudo-healthy distribution of pixels across multiple reconstructions. We employ denoising diffusion probabilistic models (DDPM) [12] to generate a pseudo-healthy reference distribution of reconstructions based on an individual input image. We then calculate the MHD between the input and the pseudo-healthy distribution to refine anomaly scoring. By considering the MHD in the pixel space with a full covariance matrix, we account for inter-pixel covariance. This enables capturing spatial information of neighboring pixels and long-range dependencies across pixels, such as symmetries in the reconstructions. Our results indicate that refining anomaly scoring by the MHD can substantially enhance the segmentation performance of conditioned DDPMs (cDDPMs), particularly when considering the inter-pixel covariance of the generated pseudo-healthy distributions. Compared to cDDPMs relying on single reconstructions, our approach leads to relative improvements of 15.9%, 35.4%, 48.0%, and 4.7% considering the AUPRC for the BRATS21, ATLAS, MSLUB and WMH data sets, respectively.

1.1 Recent Work

For most reconstruction-based approaches, AEs and VAEs are employed as GMs. While these architectures are conceptually simple and show promise in capturing the underlying distribution of healthy training data, their reconstructions tend to be blurry [2], substantially mitigating the segmentation performance. Therefore, many approaches aim to improve the reconstruction quality by focusing on spatial context [34] or utilizing 3D information [6]. Also, vector quantized VAEs and soft intro VAEs are applied to UAD in brain MRI [25, 7, 8]. Recent studies have indicated the effectiveness of DDPMs for UAD in brain MRI [24, 32, 4, 5].

Overall, GMs applied to the UAD task have shown promising progress. However, a crucial requirement for reconstruction-based UAD methods is to reconstruct healthy anatomy while avoiding the trivial replication of the input image. This necessitates the regularization of GMs, such as through a bottleneck in the latent space or additional regularization tasks like dropout [3] or denoising [13]. Consequently, imperfect reconstructions become inevitable. However, probabilistic GMs offer the appealing property of sampling multiple reconstructions. The assessment of multiple reconstructions could add valuable information for discriminating anomalies from imperfect reconstructions in the anomaly map. However, only a few studies have explored using VAEs or Bayesian AEs with Monte Carlo dropout to sample multiple reconstructions [3, 2]. These studies primarily concentrate on the mean of the generated reconstructions, which has not been shown to improve performance. Other approaches utilize uncertainty estimation to normalize the anomaly map by the estimated variance of individual pixels [29, 21, 10]. While this approach can improve the segmentation performance, it does not explicitly consider covariance across pixels. However, inter-pixel dependencies could provide valuable insights for anomaly scoring. Therefore, in this work, we focus on the inter-pixel dependencies and variations across different pseudo-healthy reconstructions and employ the MHD to measure the deviation of input pixels from the distribution of pixels in healthy reconstructions. While the MHD is commonly used for outlier detection, its typical application is at the sample level within some aggregated feature space for sample-level anomaly detection [16, 31]. Furthermore, Saase et al. [28] apply the MHD in the pixel space using a healthy data set as a reference distribution, suggesting that simple statistical methods can compete with deep learning models. However, individual brains in the training data exhibit substantial differences. As a result, relying solely on these general population-based distributions could lead to a mismatch between individual test cases and the reference distribution, potentially impeding the segmentation.

2 Methods

We propose utilizing DDPMs to construct pseudo-healthy distributions specific to each individual case during evaluation. Subsequently, these case-specific distributions are employed as a reference to compute the Mahalanobis distance (MHD) in the pixel space to refine anomaly scoring.

2.1 Generating Pseudo-healthy Distributions with DDPMs

DDPMs are specialized in learning the distribution of training images $\mathbf{x} \in \mathbb{R}^{H,W,C}$, where H , W , and C represent the height, width, and the number of channels, respectively. The training involves two primary processes: a forward process and a backward process. In the forward process, an image $\mathbf{x}_0$ is incrementally transformed into Gaussian noise, represented as $\mathbf{x}_T = \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

This transformation is guided by a predetermined noise schedule $[\beta_1, \dots, \beta_T]$. The intermediate image states $\mathbf{x}_t$ are generated by

$$\mathbf{x}_t \sim q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\sqrt{\bar{\alpha}_t} \mathbf{x}_0, (1 - \bar{\alpha}_t) \mathbf{I}) \quad \bar{\alpha}_t = \prod_{s=0}^t (1 - \beta_s).$$

The noise level at each time step $t \in [1, \dots, T]$, influences $\mathbf{x}_t$, which can vary from being the original image (at $t = 0$) to complete noise (at $t = T$). In the backward process, the reconstruction of the original image $\mathbf{x}_0^{rec}$ from the noisy state $\mathbf{x}_T$ is given by

$$\mathbf{x}_0^{rec} \sim p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) \quad p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(\boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t)).$$

Following [12], $\boldsymbol{\mu}_\theta$ is estimated using a Unet [27], and $\boldsymbol{\Sigma}(t)$ is set to $\frac{1-\alpha_t-1}{1-\alpha_t} \beta_t \mathbf{I}$. The training process entails minimizing the variational lower bound, which is approximated by the straightforward objective of predicting the added noise $\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)$, as demonstrated in [12]. This yields the simplified loss function

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

In the context of reconstruction-based UAD, our objective is not to create new images from pure noise but to reconstruct healthy brain anatomy given an input image. Therefore, during testing, $\mathbf{x}_0^{rec}$ is estimated from $\mathbf{x}_t$, determining the extent of noise in $\mathbf{x}_t$ by $t_{test} < T$. To generate a distribution of multiple reconstructions, we sample N versions of $\mathbf{x}_t$ by repeatedly resampling the additional noise and reconstructing each noised image by the denoising network. As we train the model on healthy data, this leads to a pseudo-healthy distribution consisting of N different reconstructions of the given input image.

2.2 Anomaly Scoring using Pseudo-Healthy Distributions and Mahalanobis Distance

Our goal is to leverage the informative variations within the pseudo-healthy distribution of reconstructions.

Averaged Reconstructions Initially, we calculate the mean reconstruction from multiple pseudo-healthy samples, represented as: $\boldsymbol{\mu} = \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i^{rec}$. Here, $\mathbf{x}_i^{rec}$ denotes the, i -th pseudo-healthy reconstruction, and N represents the total number of reconstructions. The anomaly score is defined as the inverted pixel-wise Structural similarity index measure (SSIM) between the input image $\mathbf{x}$ and the mean reconstruction $\boldsymbol{\mu}$:

$$S_{mean}(\mathbf{x}, \boldsymbol{\mu}) = 1 - SSIM(\mathbf{x}, \boldsymbol{\mu}). \quad (1)$$

Note that we use the pixel-wise SSIM as it has been shown to improve the anomaly scoring compared to intensity-based metrics [22, 15]

Mahalanobis Distance The MHD is a statistical measure, quantifying the distance of a sample point from a multivariate reference distribution, considering its covariance. Employing the pixels of an input image as sample points that are compared to the pseudo-healthy distribution of reconstructed pixels, we can capture the degree of deviation of each pixel in the input image from what is 'expected' in the distribution of pseudo-healthy reconstructions. First, we start by calculating the MHD with a diagonal covariance matrix $\Sigma_{diag} = \text{diag}(\sigma^2) \in \mathbb{R}^{H \cdot W \times H \cdot W}$, where σ^2 is the variance of each pixel across the N reconstructions: $\sigma^2 = \frac{1}{N-1} \sum_{i=1}^N (\mathbf{x}_i^{rec} - \boldsymbol{\mu})^2$. Note that $\mathbf{x}$ and $\boldsymbol{\mu}$ are flattened to a dimension of $\mathbb{R}^{H \cdot W}$. This yields

$$MHD_{diag}(\mathbf{x}) = \sqrt{(\mathbf{x} - \boldsymbol{\mu})^\top \Sigma_{diag}^{-1} (\mathbf{x} - \boldsymbol{\mu})}. \quad (2)$$

This approach represents a standardization and allows for scaling the distance between input pixels and the mean reconstruction by the variance of individual pixels across different reconstructions. However, the diagonal covariance matrix does not consider covariance across different pixels. To capture inter-pixel correlations, we extend our analysis to utilize a full covariance matrix, calculated as $\Sigma_{full} = \frac{1}{N-1} \sum_{i=1}^N (\mathbf{x}_i^{rec} - \boldsymbol{\mu})(\mathbf{x}_i^{rec} - \boldsymbol{\mu})^\top$ with dimension $\mathbb{R}^{H \cdot W \times H \cdot W}$, leading to

$$MHD_{full}(\mathbf{x}) = \sqrt{(\mathbf{x} - \boldsymbol{\mu})^\top \Sigma_{full}^{-1} (\mathbf{x} - \boldsymbol{\mu})}. \quad (3)$$

After reshaping the MHD map to the input image shape, the final anomaly map is obtained by a per-pixel multiplication of the MHD map with the initial anomaly map for $S_{MHD} = S_{mean} \cdot MHD_{diag}$, and $S_{sMHD} = S_{mean} \cdot MHD_{full}$, respectively.

2.3 Data

Following the principle of reconstruction-based UAD, we utilize data sets without pathologies for training while evaluating data sets that contain annotated pathologies.

For training, we utilize the IXI data set [9], consisting of MRI scans in both T1- and T2-weighting. We split the training set into a healthy test set (N=160) and partition the remaining samples into 5 training sets (N=358) and 5 validation sets (N=44) for cross-validation.

For evaluation, we utilize four different data sets, namely the BRATS21 [1] (N=1152), MSLUB [17] (N=30), ATLAS [18] (N=655) and WMH [14] (N=60) data sets that exhibit tumors, multiple sclerosis, Stroke and white-matter lesions as pathologies, respectively. Note that while we train on both weightings separately, we evaluate BRATS and MSLUB on T2-weightings and ATLAS and WMH on T1-weightings. Pre-processing of the data includes resampling to a voxel dimension of $1 \times 1 \times 1$ mm, skull-stripping, registration to the SRI ATLAS and N4 bias-correction. Furthermore, we crop 15 top and bottom slices and reduce the dimension by a factor of 2, leading to a resolution of $192 \times 192 \times 50$

voxels. During training, we process the volumes slice-wise, with slices sampled with replacement. During evaluation, we iteratively reconstruct all slices to obtain the full volume.

2.4 Implementation Details

In this work, we build upon cDDPMs proposed in [5] as a probabilistic GM. Compared to DDPMs, cDDPMs utilize an additional feature representation of the input image to guide the denoising process. We follow the architectural design of [5] with a 3-layer Unet with channel dimensions [128, 128, 256] as a denoising network. We calculate the SSIM anomaly score with a Gaussian kernel with a standard deviation of 1. When calculating the MHD, we add a small regularization term ($1e-5$) to the diagonal entries of Σ_{full} To ensure numerical stability during inversion. Additionally, we apply a Gaussian filter to the MHD map with a standard deviation of 1. We compare established state-of-the-art baselines for UAD in brain MRI, including AE [2], VAE [2], DAE [13], DDPM [32], pDDPM [4] and cDDPM [5] as reconstruction-based approaches. Moreover, we compare RD [11] and FAE [22] as feature-based methods and the self-supervised approaches PII [30] and DRAEM [33]. Finally, we evaluate the covariance model (CM) of [28], where the MHD is calculated with the healthy training set as a reference distribution. For AEs and VAEs, we set the latent dimension to 128. For VAEs, $\beta_{KLD} = 0.001$ is chosen. We train for 1600 epochs, using the ADAM optimizer, a learning rate of $1e-4$ and a batch size of 32. For all DDPM-based models, we utilize simplex noise as introduced on [32]. We uniformly sample noise levels $t \in [1, T]$ with $T = 1000$ during training and set the noise to $t_{test} = \frac{T}{2} = 500$ during evaluation. All models are implemented in PyTorch v1.10 and trained on an NVIDIA A6000 graphics card³. For evaluation, we utilize the best possible Dice-Coefficient ($\lceil \text{Dice} \rceil$) and the Area under the precision-recall curve (AUPRC). Additionally, we employ the permutation test from the MLXtend library [26]. This test involves 10,000 rounds of permutations and a significance level set at $\alpha = 5\%$ to assess statistical differences.

3 Results

We compare the segmentation performance of different variants of our approach to established state-of-the-art baselines. We average the metrics across the five folds and report the mean $\pm$ standard deviation. We initially tested different values for the number of reconstructions N in the range $N = [5, 10, \dots, 30]$ and observed a moderate improvement in performance up to $N = 10$, after which performance plateaued. Therefore, to balance performance and inference time, we selected $N = 10$ reconstructions for each input image.

³ Code available at github.com/FinnBehrendt/Mahalanobis-Unsupervised-Anomaly-Detection

Table 1. Segmentation performance regarding [Dice] and AUPRC. The highest values are shown in **bold**, where underlines denote statistical significance ($p < 0.05$). S_{mean} denotes the averaging of multiple reconstructions to derive the anomaly map. S_{MHD} and S_{sMHD} denote the use of the MHD either with a diagonal covariance matrix or with a full covariance matrix, respectively.

Model	BRATS		ATLAS		MSLUB		WMH	
	[DICE]	AUPRC	[DICE]	AUPRC	[DICE]	AUPRC	[DICE]	AUPRC
CM [28]	20.47 $\pm$ 0.22	14.03 $\pm$ 0.29	12.52 $\pm$ 0.47	9.31 $\pm$ 0.69	5.24 $\pm$ 0.27	2.59 $\pm$ 0.17	5.59 $\pm$ 0.09	2.70 $\pm$ 0.08
AE [2]	36.69 $\pm$ 0.20	33.58 $\pm$ 0.29	14.03 $\pm$ 0.27	11.68 $\pm$ 0.36	6.22 $\pm$ 0.05	3.55 $\pm$ 0.05	9.44 $\pm$ 0.26	5.60 $\pm$ 0.21
VAE [2]	36.04 $\pm$ 0.91	32.84 $\pm$ 1.07	14.48 $\pm$ 0.38	12.09 $\pm$ 0.41	6.33 $\pm$ 0.14	3.67 $\pm$ 0.11	9.52 $\pm$ 0.23	5.71 $\pm$ 0.23
FAE [22]	44.60 $\pm$ 2.17	43.75 $\pm$ 0.46	17.76 $\pm$ 0.16	13.97 $\pm$ 0.10	6.85 $\pm$ 0.65	4.02 $\pm$ 0.10	8.81 $\pm$ 0.38	4.97 $\pm$ 0.22
RD [11]	32.57 $\pm$ 0.15	27.13 $\pm$ 0.16	19.69 $\pm$ 0.26	15.65 $\pm$ 0.20	6.48 $\pm$ 0.20	3.66 $\pm$ 0.18	7.48 $\pm$ 0.10	4.22 $\pm$ 0.09
DAE [13]	62.93 $\pm$ 0.55	64.76 $\pm$ 0.79	19.42 $\pm$ 0.87	17.73 $\pm$ 0.88	8.35 $\pm$ 0.45	5.64 $\pm$ 0.37	11.14 $\pm$ 0.47	7.92 $\pm$ 0.55
DRAEM [33]	32.75 $\pm$ 3.63	26.38 $\pm$ 4.43	12.80 $\pm$ 1.94	9.63 $\pm$ 1.77	5.78 $\pm$ 2.29	2.66 $\pm$ 1.14	6.25 $\pm$ 1.89	3.23 $\pm$ 1.11
PII [30]	40.83 $\pm$ 2.18	36.49 $\pm$ 2.63	9.73 $\pm$ 1.89	7.26 $\pm$ 1.59	9.46 $\pm$ 0.43	5.21 $\pm$ 0.33	6.59 $\pm$ 1.87	3.49 $\pm$ 1.02
DDPM [32]	49.46 $\pm$ 1.56	47.57 $\pm$ 1.89	15.09 $\pm$ 0.64	11.85 $\pm$ 0.47	9.97 $\pm$ 0.64	6.03 $\pm$ 0.37	13.91 $\pm$ 0.37	9.15 $\pm$ 0.44
pDDPM [4]	54.26 $\pm$ 0.54	53.39 $\pm$ 0.70	18.83 $\pm$ 0.38	15.92 $\pm$ 0.44	10.37 $\pm$ 0.67	6.40 $\pm$ 0.51	15.31 $\pm$ 0.29	10.70 $\pm$ 0.21
cDDPM [5]	54.39 $\pm$ 0.70	54.31 $\pm$ 0.83	19.85 $\pm$ 0.90	16.99 $\pm$ 0.74	11.58 $\pm$ 0.35	7.76 $\pm$ 0.30	16.03 $\pm$ 0.88	12.15 $\pm$ 0.91
cDDPM S_{mean}	58.53 $\pm$ 0.48	59.14 $\pm$ 0.57	21.06 $\pm$ 1.09	18.17 $\pm$ 0.93	11.75 $\pm$ 0.44	7.75 $\pm$ 0.49	17.09 $\pm$ 1.24	13.15 $\pm$ 1.25
cDDPM S_{MHD}	58.47 $\pm$ 0.59	61.28 $\pm$ 0.63	20.34 $\pm$ 1.26	17.51 $\pm$ 1.23	12.25 $\pm$ 0.62	7.99 $\pm$ 0.69	16.82 $\pm$ 1.68	13.34 $\pm$ 1.90
cDDPM S_{sMHD}	64.72 $\pm$ 0.52	68.55 $\pm$ 0.63	26.67 $\pm$ 1.61	24.61 $\pm$ 1.57	15.44 $\pm$ 0.85	11.47 $\pm$ 0.79	16.65 $\pm$ 1.45	13.77 $\pm$ 1.57

The results are shown in Table 1 and Fig. 1. Comparing the baseline models, DAEs exhibit strong segmentation performance for the BRATS data set but are surpassed by cDDPMs for other pathologies in terms of Dice scores. Similarly, feature-based approaches (FAE and RD) perform well on individual data sets but struggle with generalization across all pathologies. Self-supervised approaches (PII and DRAEM) demonstrate poor performance across most data sets. Additionally, the CM method is consistently outperformed across all data sets. Overall, cDDPMs perform robustly across the evaluated data sets while enabling probabilistic sampling of multiple reconstructions. Hence, we consider cDDPMs to generate the pseudo-healthy distributions required for the MHD calculation. Our preliminary experiments indicate that other DDPM variants, such as the baseline DDPMs and pDDPMs, can also be utilized.

We find that averaging multiple reconstructions in cDDPMs (S_{mean}) enhances segmentation performance across most data sets compared to using a single reconstruction. In contrast to leveraging the MHD with a diagonal covariance matrix (S_{MHD}), utilizing the MHD with a full covariance matrix (S_{sMHD}) consistently demonstrates improved or competitive performance across all data sets. Notably, compared to the baseline cDDPMs, sampling multiple reconstructions and calculating the sMHD increases the processing time from 0.4 s to 4.9 s per volume. As illustrated in Fig. 1 (a), refining the anomaly map of cDDPMs by the sMHD leads to focused anomaly maps. Considering Fig. 1 (b), we observe non-zero correlations across the entire brain. Specifically, there exists a symmetric pattern regarding the tumor region with negative correlations in the left hemisphere and positive correlations in the right hemisphere. Exemplary anomaly maps for different models are provided in the supplementary material.




Fig. 1. (a): **Top row:** input, reconstruction, S_{mean} (SSIM), MHD, sMHD and the final anomaly map are shown for an exemplary image taken from the BRATS data set. **Bottom row:** the ground truth (green) and binarized segmentation maps (white) are shown. Note that the threshold for the segmentation maps is derived by optimizing the Dice score, based on the ground truth. (b): The correlation of one pixel (green arrow) with all other pixels, derived from Σ_{full} is visualized as a heatmap.

4 Discussion and Conclusion

A notable challenge of reconstruction-based UAD methods is their high sensitivity to imperfect reconstructions, often resulting in false positives that impede segmentation accuracy. To address this challenge, we propose to refine anomaly scoring by employing the MHD in the pixel space and identifying anomalies as outliers from pseudo-healthy distributions generated by cDDPMs.

Our results (as shown in Table 1) show that averaging multiple reconstructions from pseudo-healthy distributions (S_{mean}) can already improve segmentation performance. This improvement could be attributed to the variability of the noise structure added before reconstruction during the forward process of the cDDPM, resulting in regions with varying levels of complementary information available for denoising. Notably, applying the MHD with a diagonal covariance matrix (S_{MHD}) results in performance comparable to that of averaged reconstructions (S_{mean}). In contrast, using the spatial MHD (S_{sMHD}) substantially improves the segmentation performance. Fig. 1 (a) illustrates the differences between MHD and sMHD. It can be observed that the MHD is less sensitive to the edges of pathologies compared to sMHD. This indicates that the reconstructions exhibit higher variance in these regions, leading to a smaller weight in the anomaly map. In contrast, the anomaly map derived by sMHD shows improved pathology coverage. Fig. 1 (b) indicates the presence of inter-pixel correlations across the entire image, ranging from local neighborhoods to global dependencies exhibiting symmetry. However, the MHD with a diagonal covariance matrix does not capture these correlations. Consequently, the improved performance of sMHD compared to the MHD highlights the importance of considering these dependencies to identify abnormal pixels as outliers. Furthermore, our results indicate that considering the training data as a reference distribution for the MHD, as done in the case of CM [28], is ineffective for segmentation. This finding underscores the importance of constructing a pseudo-healthy reference distribution tailored to each individual test case, which is a key aspect of our approach.

In summary, leveraging the sMHD based on generated pseudo-healthy distributions for refining anomaly scoring can enhance the segmentation performance of DDPMs in the context of UAD in brain MRI. While we demonstrate this improved performance for cDDPMs, the baseline DDPMs and pDDPMs can also benefit from the sMHD, underscoring our method’s versatility and potential impact in enhancing anomaly detection performance. A general limitation of UAD is its restriction to binary segmentation and the overall low performance for subtler anomalies, such as those found in WMH or MSLUB data sets. While our approach increases overall performance, it is important to acknowledge the increased computational demand due to the requirement of multiple reconstructions and matrix inversion. Future work could explore efficient approximations or decompositions to enhance the computational efficiency of MHD calculations.

Acknowledgments. This work was partially funded by grant number KK5208102HV3 and ZF4026303TS9 (Zentrales Innovationsprogramm Mittelstand) and by the Free and Hanseatic City of Hamburg (Interdisciplinary Graduate School).

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

References

1. Baid, U., Ghodasara, S., Mohan, S., Bilello, M., Calabrese, E., Colak, E., Farahani, K., Kalpathy-Cramer, J., Kitamura, F.C., Pati, S., et al.: The rsna-asnr-miccai brats 2021 benchmark on brain tumor segmentation and radiogenomic classification. arXiv preprint arXiv:2107.02314 (2021)
2. Baur, C., Stefan Denner, Benedikt Wiestler, Nassir Navab, Shadi Albarqouni: Autoencoders for unsupervised anomaly segmentation in brain mr images: A comparative study. *Med. Image Anal.* **69**, 101952 (2021). <https://doi.org/10.1016/j.media.2020.101952>
3. Baur, C., Wiestler, B., Albarqouni, S., Navab, N.: Bayesian skip-autoencoders for unsupervised hyperintense anomaly detection in high resolution brain mri. In: ISBI. pp. 1905–1909 (2020). <https://doi.org/10.1109/ISBI45749.2020.9098686>
4. Behrendt, F., Bhattacharya, D., Krüger, J., Opfer, R., Schlaefer, A.: Patched diffusion models for unsupervised anomaly detection in brain mri. In: MIDL (2023)
5. Behrendt, F., Bhattacharya, D., Mieling, R., Maack, L., Krüger, J., Opfer, R., Schlaefer, A.: Guided reconstruction with conditioned diffusion models for unsupervised anomaly detection in brain mris. arXiv preprint arXiv:2312.04215 (2023)
6. Bengs, M., Behrendt, F., Krüger, J., Opfer, R., Schlaefer, A.: Three-dimensional deep learning with spatial erasing for unsupervised anomaly segmentation in brain mri. *IJCARS* **16**(9), 1413–1423 (2021). <https://doi.org/10.1007/s11548-021-02451-9>
7. Bercea, C., Benedikt Wiestler, Daniel Rueckert, Julia A Schnabel: Generalizing unsupervised anomaly detection: Towards unbiased pathology screening. MIDL (2023)
8. Bercea, C.I., Wiestler, B., Rueckert, D., Schnabel, J.A.: Reversing the abnormal: Pseudo-healthy generative networks for anomaly detection. *MICCAI* **14224**, 293–303 (2023)

9. Biomedical Image Analysis Group: Ixi dataset – brain development, <https://brain-development.org/ixi-dataset/>
10. Cai, Y., Chen, H., Yang, X., Zhou, Y., Cheng, K.T.: Dual-distribution discrepancy with self-supervised refinement for anomaly detection in medical images. *Med. Image Anal.* **86**, 102794 (2023). <https://doi.org/10.1016/j.media.2023.102794>
11. Deng, H., Li, X.: Anomaly detection via reverse distillation from one-class embedding. In: CVPR. pp. 9737–9746 (June 2022)
12. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *NeurIPS* **33**, 6840–6851 (2020)
13. Kascenas, A., Pugeault, N., O’Neil, A.Q.: Denoising autoencoders for unsupervised anomaly detection in brain mri. In: MIDL (2022)
14. Kuijff, H.J., Biesbroek, J.M., De Bresser, J., Heinen, R., Andermatt, S., Bento, M., Berseth, M., Belyaev, M., Cardoso, M.J., Casamitjana, A., et al.: Standardized assessment of automatic segmentation of white matter hyperintensities and results of the wmh segmentation challenge. *IEEE TMI* **38**(11), 2556–2568 (2019)
15. Lagogiannis, I., Meissen, F., Kaissis, G., Rueckert, D.: Unsupervised pathology detection: A deep dive into the state of the art. *IEEE TMI* **PP** (2023). <https://doi.org/10.1109/TMI.2023.3298093>
16. Lee, K., Lee, K., Lee, H., Shin, J.: A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *NeurIPS* **31** (2018)
17. Lesjak, Ž., Galimzianova, A., Koren, A., Lukin, M., Pernuš, F., Likar, B., Špiclin, Ž.: A novel public mr image dataset of multiple sclerosis patients with lesion segmentations based on multi-rater consensus. *Neuroinformatics* **16**(1), 51–63 (2018)
18. Liew, S.L., Lo, B.P., Donnelly, M.R., Zavaliangos-Petropulu, A., Jeong, J.N., Barisano, G., Hutton, A., Simon, J.P., Juliano, J.M., Suri, A., et al.: A large, curated, open-source stroke neuroimaging dataset to improve lesion segmentation algorithms. *Scientific data* **9**(1), 320 (2022)
19. Lundervold, A.S., Lundervold, A.: An overview of deep learning in medical imaging focusing on mri. *Zeitschrift fur medizinische Physik* **29**(2), 102–127 (2019). <https://doi.org/10.1016/j.zemedi.2018.11.002>
20. Mahalanobis, P.: On the generalised distance in statistics. In: Proceedings of the National Institute of Science of India. vol. 12, pp. 49–55 (1936)
21. Mao, Y., Xue, F.F., Wang, R., Zhang, J., Zheng, W.S., Liu, H.: Abnormality detection in chest x-ray images using uncertainty prediction autoencoders. In: MICCAI. pp. 529–538. Springer (2020)
22. Meissen, F., Paetzold, J., Kaissis, G., Rueckert, D.: Unsupervised anomaly localization with structural feature-autoencoders. *arXiv preprint arXiv:2208.10992* (2022)
23. Meissen, F., Wiestler, B., Kaissis, G., Rueckert, D.: On the pitfalls of using the residual error as anomaly score. In: MIDL (2022)
24. Pinaya, W.H.L., Graham, M.S., Gray, R., Da Costa, P.F., Tudosiu, P.D., Wright, P., Mah, Y.H., MacKinnon, A.D., Teo, J.T., Jager, R., et al.: Fast unsupervised brain anomaly detection and segmentation with diffusion models. In: MICCAI (2022)
25. Pinaya, W.H.L., Tudosiu, P.D., Gray, R., Rees, G., Nachev, P., Ourselin, S., Cardoso, M.J.: Unsupervised brain imaging 3d anomaly detection and segmentation with transformers. *Med. Image Anal.* **79**, 102475 (2022). <https://doi.org/10.1016/j.media.2022.102475>
26. Raschka, S.: Mlxtend: Providing machine learning and data science utilities and extensions to python’s scientific computing stack. *The Journal of Open Source Software* **3**(24) (2018)

27. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI. pp. 234–241 (2015)
28. Saase, V., Wenz, H., Ganslandt, T., Groden, C., Maros, M.E.: Simple statistical methods for unsupervised brain anomaly detection on mri are competitive to deep learning methods. arXiv preprint arXiv:2011.12735 (2020)
29. Sato, K., Hama, K., Matsubara, T., Uehara, K.: Predictable uncertainty-aware unsupervised deep anomaly segmentation. In: IJCNN. pp. 1–7. IEEE, Piscataway, NJ (2019). <https://doi.org/10.1109/IJCNN.2019.8852144>
30. Tan, J., Hou, B., Day, T., Simpson, J., Rueckert, D., Kainz, B.: Detecting outliers with poisson image interpolation. In: MICCAI (2021). https://doi.org/10.1007/978-3-030-87240-3_56
31. Vasiliuk, A., Frolova, D., Belyaev, M., Shirokikh, B.: Limitations of out-of-distribution detection in 3d medical image segmentation. JMI **9**(9) (2023). <https://doi.org/10.3390/jimaging9090191>
32. Wyatt, J., Leach, A., Schmon, S.M., Willcocks, C.G.: Anoddpm: Anomaly detection with denoising diffusion probabilistic models using simplex noise. In: CVPR. pp. 650–656
33. Zavrtanik, V., Kristan, M., Skočaj, D.: Draem-a discriminatively trained reconstruction embedding for surface anomaly detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8330–8339 (2021)
34. Zimmerer, D., Kohl, S., Petersen, J., Isensee, F., Maier-Hein, K.: Context-encoding variational autoencoder for unsupervised anomaly detection. In: MIDL (2019)