



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Binary Noise for Binary Tasks: Masked Bernoulli Diffusion for Unsupervised Anomaly Detection

Julia Wolleb, Florentin Bieder, Paul Friedrich, Peter Zhang,
Alicia Durrer, and Philippe C. Cattin

Department of Biomedical Engineering, University of Basel, Allschwil, Switzerland
julia.wolleb@unibas.ch

Abstract. The high performance of denoising diffusion models for image generation has paved the way for their application in unsupervised medical anomaly detection. As diffusion-based methods require a lot of GPU memory and have long sampling times, we present a novel and fast unsupervised anomaly detection approach based on latent Bernoulli diffusion models. We first apply an autoencoder to compress the input images into a binary latent representation. Next, a diffusion model that follows a Bernoulli noise schedule is employed to this latent space and trained to restore binary latent representations from perturbed ones. The binary nature of this diffusion model allows us to identify entries in the latent space that have a high probability of flipping their binary code during the denoising process, which indicates out-of-distribution data. We propose a masking algorithm based on these probabilities, which improves the anomaly detection scores. We achieve state-of-the-art performance compared to other diffusion-based unsupervised anomaly detection algorithms while significantly reducing sampling time and memory consumption. The code is available at https://github.com/JuliaWolleb/Anomaly_berdiff.

Keywords: Bernoulli Diffusion · Unsupervised Anomaly Detection

1 Introduction

The recent success of denoising diffusion models [8,10] has paved the way for unsupervised anomaly detection methods, which aim to identify unusual patterns, outliers, or deviations from normal behavior without relying on labeled examples [9,11,24]. In this work, we aim to identify pixel-level anomalous changes in medical images while training on a dataset of healthy subjects only. We thus aim to transform pathological tissue of the input image into healthy tissue, and use the difference map between input and output to highlight abnormal areas [3,13,25,30]. As approaches operating in the image-space suffer from long sampling times, we propose a novel method that relies on binary latent diffusion models [28]. An overview of the model training is presented in Fig. 1. Our approach uses a binarizing autoencoder to compress input images into a binary latent space. A diffusion model that follows a Bernoulli noise schedule is then

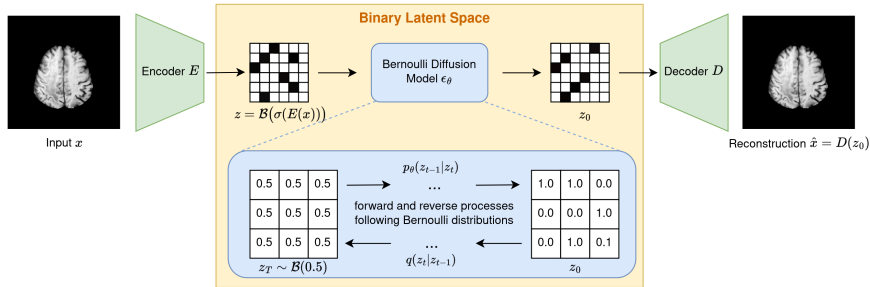


Fig. 1. We train a binarizing autoencoder consisting of an encoder E and a decoder D . In the latent space, a diffusion model is trained to restore the binary latent code of healthy images by reversing a Bernoulli noise driven diffusion process. Both the autoencoder and the diffusion model are exclusively trained on healthy data.

trained to gradually denoise perturbed representations within the latent space. For each entry of the latent representation, the diffusion model directly extracts the probability of flipping its binary code during the denoising process. We argue that entries associated with anomalous changes are more likely to flip, indicating out-of-distribution data. Based on this assumption, we design a novel masking strategy that enables reconstructions close to the input image, thus improving the accuracy of the difference map and enhancing anomaly detection scores. We assess our method on the BRATS2020 dataset for brain tumor detection, and the OCT2017 dataset for drusen detection in retinal OCT images.

Related Work Many unsupervised anomaly detection algorithms rely on variational autoencoders (VAEs) [16] trained on healthy samples only. They extract anomaly scores from the reconstruction error between the pathological input image and its healthy reconstruction [20,31]. Others are based on self-supervised learning using segmentation [21], patch-based regression [5], or denoising autoencoders [13]. Due to their high reconstruction quality, diffusion models [8,10] have been applied to weakly supervised anomaly detection tasks [26,29]. If no labels are available, *AnoDDPM* [30] suggests replacing the Gaussian noise scheme with simplex noise. This is further explored in [14], where a coarse noise schedule is proposed. Various diffusion-based masking mechanisms have been explored [3,12,18] to reconstruct healthy tissue. To further improve the quality of these healthy reconstructions, *AutoDDPM* [4] proposes an iterative masking, stitching, and resampling algorithm. To address concerns of memory consumption and sampling time, latent diffusion models (LDMs) [25] encode input images into a low-dimensional latent space, where a diffusion model is trained to restore healthy latent representations. While classical denoising diffusion models are based on Gaussian noise, the use of Bernoulli noise was introduced as early as 2015 [27] and has been applied for fully supervised lesion segmentation [6]. Furthermore, a binary LDM, using a binarizing autoencoder in combination with a latent Bernoulli diffusion model, was proposed for image synthesis [28].

Contribution We present a novel approach for unsupervised anomaly detection in medical images. Our method uses a binary latent diffusion model based on Bernoulli noise to learn a robust representation from healthy images. Taking advantage of the binary nature of Bernoulli diffusion models, we design a novel masking scheme to preserve the anatomical information of the input images, thereby improving the anomaly detection performance. We achieve state-of-the-art anomaly detection results on two different medical datasets while reducing sampling time and memory footprint.

2 Method

Our approach consists of two neural networks, shown in Fig. 1: An autoencoder (in green) mapping an input image into a binary latent space, as well as a Bernoulli diffusion model (in blue) in the binary latent space. First, we train the autoencoder, consisting of an encoder E and a decoder D .

Autoencoder Following [28], the convolutional encoder E extracts a compressed representation $y = \sigma(E(x)) \in \mathbb{R}^{C \times \frac{h}{k} \times \frac{w}{k}}$ from an input image $x \in \mathbb{R}^{c \times h \times w}$. The channel dimension of the input image is denoted by c , k is the compression rate, σ is the sigmoid activation function, and C is the number of encoded feature channels. By sampling $z \sim \mathcal{B}(y)$, we obtain a binary representation z . The decoder D reconstructs the image $\hat{x} = D(z)$. We train the autoencoder using a combined MSE, perceptual, and adversarial loss between x and \hat{x} [28].

Bernoulli Diffusion Model Following [6], the Bernoulli diffusion model takes the latent representation $z = z_0$ as input. The Bernoulli forward process follows a sequence of noise perturbations

$$q(z_t | z_{t-1}) := \mathcal{B}\left((1 - \beta_t)z_{t-1} + \frac{\beta_t}{2}\right) \text{ for } t \in 1, \dots, T, \quad (1)$$

where \mathcal{B} denotes the Bernoulli distribution, and $\beta_{1:T}$ follow a predefined noise schedule. By defining $\alpha_t := 1 - \beta_t$ and $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$, we can directly compute

$$z_t = z_0 \oplus \epsilon, \quad \text{with } \epsilon \sim \mathcal{B}\left(\frac{1 - \bar{\alpha}_t}{2}\right), \quad (2)$$

where \oplus is the logical XOR operation. The Bernoulli posterior is defined as

$$q(z_{t-1} | z_t, z_0) = \mathcal{B}(\theta_{post}(z_t, z_0)), \quad (3)$$

which is described in detail in the Supplementary Material. For the reverse diffusion process, we train a deep neural network ϵ_θ , following the architecture proposed in [8], to predict the residual $\hat{\epsilon} = \epsilon_\theta(z_t, t)$, i.e., the flipping probability of the binary code z_t [28]. For each time step t we predict $\tilde{z}_0 = |z_t - \epsilon_\theta(z_t, t)|$ and approximate z_0 in Equation 3 with \tilde{z}_0 , i.e.,

$$p_\theta(z_{t-1} | z_t, \tilde{z}_0) := \mathcal{B}(\theta_{post}(z_t, \tilde{z}_0)) = \mathcal{B}(\theta_{post}(z_t, |z_t - \epsilon_\theta(z_t, t)|)). \quad (4)$$

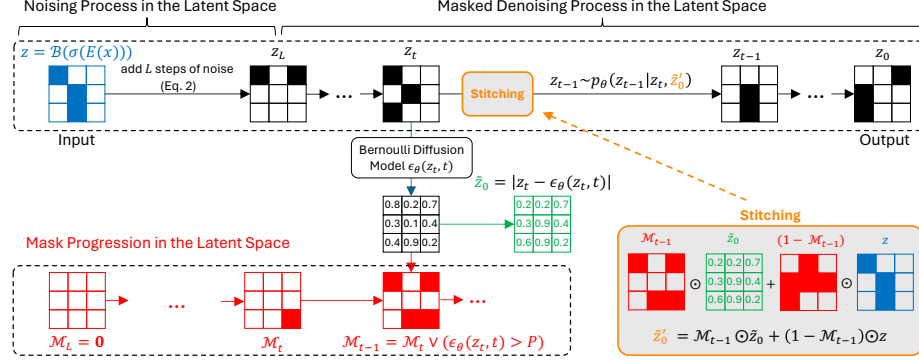


Fig. 2. In our proposed denoising process, we aim to restore a healthy representation z_0 from an initial input representation z . We define a mask \mathcal{M} based on the threshold probability P . By stitching the sampled and original latent representations at each timestep, we ensure the preservation of anatomical information from the input.

To train the model ϵ_θ , we use the binary cross entropy (BCE) as loss function

$$\mathcal{L}(z_t, t) = BCE(\epsilon_\theta(z_t, t), z_t \oplus z_0). \quad (5)$$

To generate synthetic images, we start from random noise $z_T \sim \mathcal{B}(\frac{1}{2})$ and iterate through the denoising process by sampling

$$z_{t-1} \sim \mathcal{B}(\theta_{post}(z_t, \tilde{z}_0)) \text{ for } t = T, \dots, 1. \quad (6)$$

Anomaly Detection during Inference An overview of the pipeline is shown in Fig. 2. To detect anomalies in an image x of a possibly diseased subject, we first pass it through the encoder E to get a binary representation $z = \mathcal{B}(\sigma(E(x)))$. Similar to other diffusion-based reconstruction methods [14,25,29,30], we add L steps of noise to z using Equation 2. The resulting noisy latent representation z_L is then passed through the denoising process described in Equation 6 for $t = L, \dots, 1$. Since the diffusion model was trained on healthy subjects only, it restores a latent representation z_0 following the learned healthy distribution. We finally reconstruct an output image $\hat{x} = D(z_0)$. The desired pixel-wise anomaly map is computed by $a = \sum_c (x - \hat{x})^2$.

Masking As described in [4], restoration-based methods face the *noise paradox*: We have to strike a balance between sufficiently large noise levels L while preserving identity information in healthy tissue. While increasing the noise level L reduces anomalies in the reconstruction, the integrity of healthy tissue is at risk, leading to false positive pixels in the anomaly map a . We propose a novel masking algorithm, presented in Fig. 2, to tackle this challenge. We take advantage of the binary nature of our Bernoulli diffusion model. We argue that in any timestep

t , entries of the latent representation related to anomalous changes are likely to have a high flipping probability $\epsilon_\theta(z_t, t)$, as they do not follow the expected healthy distribution. First, we initialize a mask $\mathcal{M}_L = \mathbf{0} = [0]^{C \times \frac{h}{k} \times \frac{w}{k}}$ in the latent space, and define a threshold probability P . In every step t , we update \mathcal{M}_t by masking out entries that have a flipping probability higher than P , by computing $\mathcal{M}_{t-1} := \mathcal{M}_t \vee (\epsilon_\theta(z_t, t) > P)$. To preserve anatomical information from the original latent representation z , we stitch the predicted $\tilde{z}_0 = |z_t - \epsilon_\theta(z_t, t)|$ and the original z according to

$$\tilde{z}'_0 = \mathcal{M}_{t-1} \odot \tilde{z}_0 + (1 - \mathcal{M}_{t-1}) \odot z, \quad (7)$$

where \odot denotes an element-wise multiplication. Thereby, we only change entries with a high flipping probability, while preserving the rest of the original representation z . Finally, we sample $z_{t-1} \sim \mathcal{B}(\theta_{post}(z_t, \tilde{z}'_0))$ according to Equation 6. The whole inference procedure is presented in Algorithm 1.

Algorithm 1 Anomaly detection using our masked Bernoulli diffusion model

Input: input image x , probability threshold P , noise level L
Output: healthy reconstruction \hat{x} , anomaly map a
Initialize $\mathcal{M}_L = \mathbf{0}$, $z = \mathcal{B}(\sigma(E(x)))$
 $z_L \leftarrow z \oplus \epsilon_L$, where $\epsilon_L \sim \mathcal{B}(\frac{1-\alpha_L}{2})$
for all t from L to 1 **do**
 $\tilde{z}_0 \leftarrow |z_t - \epsilon_\theta(z_t, t)|$ {prediction of z at time 0}
 $\mathcal{M}_{t-1} \leftarrow \mathcal{M}_t \vee (\epsilon_\theta(z_t, t) > P)$ {update mask}
 $\tilde{z}'_0 \leftarrow \mathcal{M}_{t-1} \odot \tilde{z}_0 + (1 - \mathcal{M}_{t-1}) \odot z$ { element-wise stitching of z and \tilde{z}_0 }
 $z_{t-1} \sim \mathcal{B}(\theta_{post}(z_t, \tilde{z}'_0))$ { sample z_{t-1} }
end for
 $\hat{x} = D(z_0)$
 $a \leftarrow \sum_c (x - \hat{x})^2$
return \hat{x} , a

3 Experiments

The binary autoencoder is trained as proposed in [28], with a batch size of 6. We encode images of resolution $c \times 256 \times 256$ to a resolution of $128 \times 32 \times 32$ in the binary latent space, where c is the number of input channels. For the Bernoulli diffusion model, we choose a U-Net architecture as described in [6,8]. The diffusion model has $T = 1000$ timesteps and is trained using the Adam optimizer with a learning rate of 10^{-4} and a batch size of 32. By choosing the number of channels in the first layer of the diffusion model to be 128, and using one attention head at resolution 16, the total number of parameters is 36,034,432 for the Bernoulli diffusion model and 4,256,808 for the binarizing autoencoder. The diffusion model is trained for 2×10^5 iterations and the autoencoder for 1.2×10^4 iterations on an NVIDIA Quadro RTX 6000 GPU using Pytorch 2.1.0.

BRATS2020 This dataset [1,2,23] consists of 3D brain MR images of subjects diagnosed with a brain tumor, together with pixel-wise ground truth labels outlining tumor regions. Our focus on a 2D approach leads us to analyse axial slices. Each slice has four channels corresponding to four different MR sequences, is zero-padded to dimensions of 256×256 and normalized in the range 0 to 1. Given the typical occurrence of tumors in the central regions of the brain [17], we omit the bottom 80 and the top 26 slices. A slice is classified as healthy if the ground truth label mask is zero. We split the dataset at the patient level, resulting in a training set of 5,598 healthy slices. The test set contains 1,038 slices with tumors and 705 without. To post-process the obtained anomaly maps a , similar to [3], we apply median filtering with a kernel-size of 5, apply a threshold at 0.5 to obtain a segmentation [30], and remove components below 10 pixels.

OCT2017 This dataset [15] contains optical coherence tomography (OCT) images of the retina. We select 18,232 healthy subjects for training. The dataset comprises grayscale images resized to a resolution of 256×256 and normalized to values between 0 and 1. We randomly select subjects suffering from drusen for qualitative comparison. Drusen manifest as accumulations of extracellular material between the Bruch’s membrane and the retinal pigment epithelium (RPE), causing elevations of the RPE [7]. There are no ground truth anomaly maps available for the OCT dataset, so quantitative results are not possible for this dataset.

4 Results and Discussion

Hyperparameter Selection To explore the effect of the threshold probability P and the noise level L , we perform a grid search on the BRATS2020 dataset. In Fig. 3 (left), we report the Dice score between the anomaly map a and the ground truth tumor segmentation. Note that a threshold of $P = 0$ corresponds to no mask being applied. Notably, a noise level of $L = 200$ combined with $P = 0.5$ achieves the highest scores. We choose this setting for further comparison with the state-of-the-art. Additional analyses of the effect of the hyperparameters L and P on the restored images \hat{x} can be found in the Supplementary Material.

Image-Level Masking Scores Taking advantage of the binary nature of Bernoulli diffusion models, we can directly extract image-level anomaly scores from the model output. In Fig. 3 (right), we plot the percentage of masked entries in \mathcal{M} for different masking thresholds P and noise levels L . We plot these values separately for the BRATS2020 test set of diseased and healthy slices. The box plots show clear differences between the two groups, indicating that the masked flipping probabilities extracted by our masking schedule are indicative for out-of-distribution data.

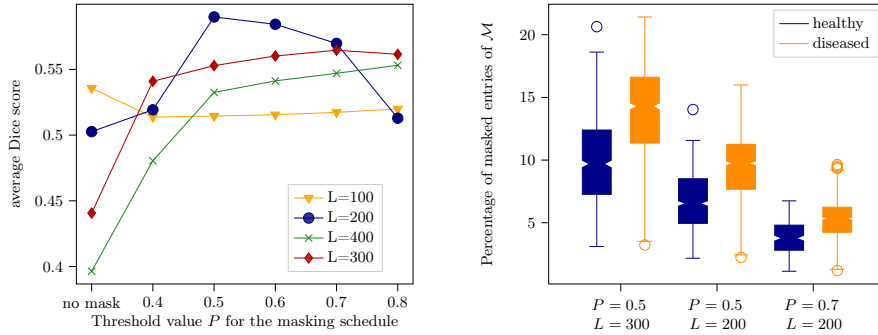


Fig. 3. On the left, we present Dice scores on the BRATS2020 test set for different noise levels L and probability thresholds P . On the right, we show the percentage of the masked entries of \mathcal{M} for the healthy (blue) and diseased (orange) groups for exemplary settings of P and L .

Comparison to the State of the Art We compare our method to *AnoDDPM* [30], *pDDPM* [3], a latent diffusion model *LDM* [25], and *AutoDDPM* [4]. Implementation details can be found at https://github.com/JuliaWolleb/Anomaly_berdiff. Quantitative results on the BRATS2020 test set are summarized in Table 1, including the Dice scores, the area under the precision-recall curve (AUPRC), the peak signal-to-noise ratio (PSNR) between x and \hat{x} , as well as the sampling time and GPU-memory consumption per input image. Our method achieves anomaly Dice scores comparable to the state of the art, while significantly reducing the sampling time to 5s. However, our method still lags behind AnoDDPM in terms of AUPRC. Compared to the *LDM*, which also applies a masking schedule in a compressed latent space, we obtain higher scores, highlighting the effectiveness of our proposed binary masking scheme. We present exemplary results for all methods in Fig. 4 for a qualitative comparison. On the BRATS2020 dataset, all methods produce good reconstruction results. *AutoDDPM* generates high-quality restorations, also reflected in the high PSNR values. Our method obtains a clear anomaly map with slight changes in the ventricles. As noted by [22], hyperintense lesions can be effectively identified using a simple thresholding function. However, addressing abnormalities in anatomical

Table 1. We report the mean \pm standard deviation over the BRATS2020 test set of diseased slices for the different comparing methods.

| Method | Noise | Dice \uparrow | AUPRC \uparrow | PSNR \uparrow | Time \downarrow | Mem. \downarrow |
|-----------------|-----------|------------------|------------------|-----------------|-------------------|-------------------|
| AnoDDPM | Simplex | 0.594 ± 0.26 | 0.727 ± 0.28 | 25.45 ± 2.4 | 50s | 3.85 GB |
| pDDPM | Simplex | 0.473 ± 0.20 | 0.471 ± 0.24 | 22.82 ± 1.9 | 103s | 3.44 GB |
| LDM | Gaussian | 0.490 ± 0.24 | 0.576 ± 0.27 | 20.03 ± 2.2 | 8s | 1.87 GB |
| AutoDDPM | Gaussian | 0.568 ± 0.21 | 0.632 ± 0.27 | 26.06 ± 1.3 | 32s | 4.26 GB |
| Ours (unmasked) | Bernoulli | 0.503 ± 0.22 | 0.515 ± 0.26 | 23.31 ± 2.0 | 4s | 1.47 GB |
| Ours (masked) | Bernoulli | 0.590 ± 0.24 | 0.656 ± 0.27 | 25.63 ± 2.1 | 5s | 1.47 GB |




Fig. 4. Qualitative results of all comparing methods for a subject from the BRATS2020 test set, as well as a subject from the OCT2017 dataset suffering from drusen.

structure deformations requires more than just eliminating hyperintense tissue. Thus, we evaluate our method for drusen detection in OCT images. As shown in Fig. 4 on the right, our approach successfully restores a smooth RPE. Additional examples are provided in the Supplementary Material.

5 Conclusion

In this paper, we present a novel unsupervised anomaly detection method based on a latent Bernoulli diffusion model. By exploiting the binary nature of Bernoulli noise, we can extract anomaly scores directly from the predicted flipping probabilities. Based on this assumption, we design a novel masking procedure to improve the pixel-wise anomaly detection performance in medical images. We evaluate our approach on two different medical datasets and successfully translate images containing anomalies into images without, resulting in competitive anomaly scores. Quantitative comparisons show that our method significantly reduces sampling time and memory consumption while still providing convincing restoration of healthy subjects. Our approach is particularly promising for

3D applications, where memory constraints and sampling times are significant challenges. In addition, this work creates opportunities for other binary downstream tasks in medical image analysis. For future work, an extension to binary convolutional neural networks [19] can be explored.

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

References

1. Bakas, S., Akbari, H., Sotiras, A., Bilello, M., Rozycki, M., Kirby, J.S., Freymann, J.B., Farahani, K., Davatzikos, C.: Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features. *Scientific data* 4(1), 1–13 (2017)
2. Bakas, S., Reyes, M., Jakab, A., Bauer, S., Rempfler, M., Crimi, A., Shinohara, R.T., Berger, C., Ha, S.M., Rozycki, M., et al.: Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the BRATS challenge. *arXiv preprint arXiv:1811.02629* (2018)
3. Behrendt, F., Bhattacharya, D., Krüger, J., Opfer, R., Schlaefer, A.: Patched diffusion models for unsupervised anomaly detection in brain mri. In: *Medical Imaging with Deep Learning*. pp. 1019–1032. PMLR (2024)
4. Bercea, C.I., Neumayr, M., Rueckert, D., Schnabel, J.A.: Mask, stitch, and re-sample: Enhancing robustness and generalizability in anomaly detection through automatic diffusion models. *arXiv preprint arXiv:2305.19643* (2023)
5. Bieder, F., Wolleb, J., Sandkühler, R., Cattin, P.C.: Position regression for unsupervised anomaly detection. In: *International Conference on Medical Imaging with Deep Learning*. pp. 160–172. PMLR (2022)
6. Chen, T., Wang, C., Shan, H.: Berdiff: Conditional bernoulli diffusion model for medical image segmentation. *arXiv preprint arXiv:2304.04429* (2023)
7. Despotovic, I.N., Ferrara, D.: 6.1.1 - drusen. In: Goldman, D.R., Waheed, N.K., Duker, J.S. (eds.) *Atlas of Retinal OCT: Optical Coherence Tomography*, pp. 16–23. Elsevier (2018)
8. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021)
9. Gonzalez-Jimenez, A., Lionetti, S., Pouly, M., Navarini, A.A.: Sano: Score-based diffusion model for anomaly localization in dermatology. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2987–2993 (2023)
10. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
11. Huijben, E., Amirrajab, S., Pluim, J.P.: Histogram-and diffusion-based medical out-of-distribution detection. *arXiv preprint arXiv:2310.08654* (2023)
12. Iqbal, H., Khalid, U., Chen, C., Hua, J.: Unsupervised anomaly detection in medical images using masked diffusion model. In: *International Workshop on Machine Learning in Medical Imaging*. pp. 372–381. Springer (2023)
13. Kascenas, A., Pugeault, N., O’Neil, A.Q.: Denoising autoencoders for unsupervised anomaly detection in brain mri. In: *International Conference on Medical Imaging with Deep Learning*. pp. 653–664. PMLR (2022)

14. Kascenas, A., Sanchez, P., Schrenpf, P., Wang, C., Clackett, W., Mikhael, S.S., Voisey, J.P., Goatman, K., Weir, A., Pugeault, N., et al.: The role of noise in denoising models for anomaly detection in medical images. *Medical Image Analysis* **90**, 102963 (2023)
15. Kermany, D., Zhang, K., Goldbaum, M., et al.: Labeled optical coherence tomography (oct) and chest x-ray images for classification. *Mendeley data* **2**(2), 651 (2018)
16. Kingma, D.P., Welling, M.: An introduction to variational autoencoders. *arXiv preprint arXiv:1906.02691* (2019)
17. Larjavaara, S., Mäntylä, R., Salminen, T., Haapasalo, H., Raitanen, J., Jääskeläinen, J., Auvinen, A.: Incidence of gliomas by anatomic location. *Neuro-oncology* **9**(3), 319–325 (2007)
18. Liang, Z., Anthony, H., Wagner, F., Kamnitsas, K.: Modality cycles with masked conditional diffusion for unsupervised anomaly segmentation in mri. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 168–181. Springer (2023)
19. Lin, X., Zhao, C., Pan, W.: Towards accurate binary convolutional neural network. *Advances in neural information processing systems* **30** (2017)
20. Marimont, S.N., Tarroni, G.: Anomaly detection through latent space restoration using vector quantized variational autoencoders. In: *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*. pp. 1764–1767. IEEE (2021)
21. Marimont, S.N., Tarroni, G.: Achieving state-of-the-art performance in the medical out-of-distribution (mood) challenge using plausible synthetic anomalies (2023)
22. Meissen, F., Kaissis, G., Rueckert, D.: Challenging current semi-supervised anomaly segmentation methods for brain mri. In: *International MICCAI brain-lesion workshop*. pp. 63–74. Springer (2021)
23. Menze, B.H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., et al.: The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE transactions on medical imaging* **34**(10), 1993–2024 (2014)
24. Naval Marimont, S., Baugh, M., Siomos, V., Tzelepis, C., Kainz, B., Tarroni, G.: Disyre: Diffusion-inspired synthetic restoration for unsupervised anomaly detection. In: *Proceedings/IEEE International Symposium on Biomedical Imaging: from nano to macro*. IEEE (2024)
25. Pinaya, W.H., Graham, M.S., Gray, R., Da Costa, P.F., Tudosi, 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: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 705–714. Springer (2022)
26. Sanchez, P., Kascenas, A., Liu, X., O’Neil, A.Q., Tsaftaris, S.A.: What is healthy? generative counterfactual diffusion for lesion localization. In: *MICCAI Workshop on Deep Generative Models*. pp. 34–44. Springer (2022)
27. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: *International conference on machine learning*. pp. 2256–2265. PMLR (2015)
28. Wang, Z., Wang, J., Liu, Z., Qiu, Q.: Binary latent diffusion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 22576–22585 (2023)
29. Wolleb, J., Bieder, F., Sandkühler, R., Cattin, P.C.: Diffusion models for medical anomaly detection. In: *International Conference on Medical image computing and computer-assisted intervention*. pp. 35–45. Springer (2022)

30. Wyatt, J., Leach, A., Schmon, S.M., Willcocks, C.G.: Anoddpm: Anomaly detection with denoising diffusion probabilistic models using simplex noise. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 650–656 (2022)
31. Zimmerer, D., Isensee, F., Petersen, J., Kohl, S., Maier-Hein, K.: Unsupervised anomaly localization using variational auto-encoders. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part IV 22. pp. 289–297. Springer (2019)