



City Research Online

City, University of London Institutional Repository

Citation: Naval Marimont, S., Baugh, M., Siomos, V., Tzelepis, C., Kainz, B. & Tarroni, G. (2024). DISYRE: Diffusion-Inspired SYNthetic REstoration for Unsupervised Anomaly Detection. Paper presented at the IEEE International Symposium on Biomedical Imaging 2024, 27-30 May 2024, Athens, Greece.

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/32200/>

Link to published version:

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

DISYRE: DIFFUSION-INSPIRED SYNTHETIC RESTORATION FOR UNSUPERVISED ANOMALY DETECTION

Sergio Naval Marimont¹ Matthew Baugh² Vasilis Siomos¹ Christos Tzelepis¹
Bernhard Kainz^{2,3} Giacomo Tarroni^{1,2}

¹ CitAI Research Centre, Department of Computer Science City, University of London

² BioMedIA, Department of Computing, Imperial College London

³ Department of Artificial Intelligence in Biomedical Engineering, FAU Erlangen-Nürnberg

ABSTRACT

Unsupervised Anomaly Detection (UAD) techniques aim to identify and localize anomalies without relying on annotations, only leveraging a model trained on a dataset known to be free of anomalies. Diffusion models learn to modify inputs x to increase the probability of it belonging to a desired distribution, *i.e.*, they model the score function $\nabla_x \log p(x)$. Such a score function is potentially relevant for UAD, since $\nabla_x \log p(x)$ is itself a pixel-wise anomaly score. However, diffusion models are trained to invert a corruption process based on Gaussian noise and the learned score function is unlikely to generalize to medical anomalies. This work addresses the problem of how to learn a score function relevant for UAD and proposes DISYRE: Diffusion-Inspired SYnthetic REstoration. We retain the diffusion-like pipeline but replace the Gaussian noise corruption with a gradual, synthetic anomaly corruption so the learned score function generalizes to medical, naturally occurring anomalies. We evaluate DISYRE on three common Brain MRI UAD benchmarks and substantially outperform other methods in two out of the three tasks.

Index Terms— Unsupervised anomaly detection, out-of-distribution detection, diffusion models, synthetic anomalies

1. INTRODUCTION

Unsupervised Anomaly Detection (UAD) in the field of medical image analysis involves the detection and/or localization of irregularities in medical images without the need for annotations. This approach is particularly valuable since obtaining annotations can be both expensive and challenging [1]. During the prediction phase, UAD methodologies employ comparisons between a test image and the established distribution, often using a parameterized model of this distribution. Anomalies are detected when deviations from this healthy distribution are observed in the test image.

Many UAD methods rely on modelling the expected distribution with generative models such as variational auto-encoders [2] or generative adversarial networks [3]. These

rely on the assumption that generative models trained on in-distribution data will be unable to reconstruct out-of-distribution regions, meaning that the reconstruction error can be used to identify anomalous pixels. Unfortunately, this assumption often fails, as models may not accurately reconstruct normal data and the remaining residual error needs to be calibrated as subject-specific bias [2, 4]. Furthermore, auto-encoders are also known to have the ability to also reconstruct anomalies [4].

As an alternative, restoration-based methods [5] use Maximum-A-Posteriori (MAP) estimation to remove local anomalies from images. Many approaches have since followed this idea, often using a two-stage restoration process; first identifying anomalous areas or latent variables and then resampling them to become in-distribution [6, 4].

More recently, diffusion models have emerged as the state-of-the-art method for image generation. The most common of these is the Denoising Diffusion Probabilistic Model (DDPM) [7], which constructs a mapping from Gaussian noise to a desired image distribution. To do this, DDPMs leverage Langevin dynamics [7] to learn a score function, *i.e.* $\nabla_x \log p(x)$, which is able to iteratively denoise a sample of Gaussian noise to produce a sample from the target image distribution. Such a score function is potentially very relevant for UAD, as it is able to indicate both where and how to edit the image to increase its likelihood, directly acting as a pixel-wise anomaly score. However, as a standard DDPM learns a score function that is specific to images corrupted with Gaussian noise, it does not generalize to correcting images containing other types of anomalies.

Because of this, DDPM’s score function has not been favoured when applying diffusion models to UAD. AnODPM [8] instead replaces Gaussian noise with simplex noise, as the multi-scale nature of simplex noise trains the score function to be able to correct larger anomalous structures. This is consistent with what is reported by DAE [9] authors, who found that denoising models trained on reasonably coarse noise outperformed those using pixel-level noise. An alternative approach [10] involves utilizing classifier guid-

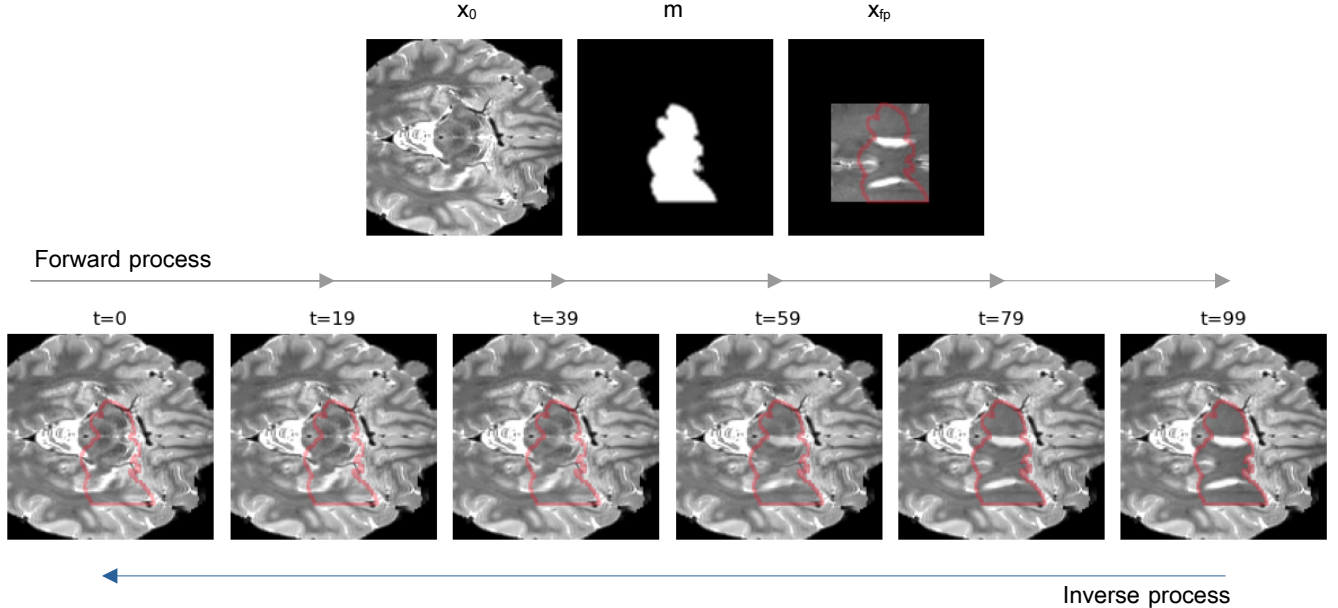


Fig. 1. 1st row: examples of anomaly generation inputs: healthy image x_0 , shape mask m and foreign patch x_{fp} . 2nd row: synthetic corruption forward process for different t values (with $T = 100$).

ance to encourage the model to restore the image to be part of an expected distribution. However, this requires training a discriminator model with sample-level labels.

Separately from generative models, Foreign Patch Interpolation (FPI) [11] proposed directly training anomaly segmentation networks by adding synthetic anomalies to otherwise healthy anatomies. The anomalies are generated by interpolating a square region of the sample with a patch extracted from a separate sample, controlling the interpolation with a factor α . One of the limitations of FPI is that it produces image gradients on the boundaries between synthetic anomalies and original images, a feature which the model could exploit as a shortcut to identify synthetic anomalies. Poisson Image Interpolation (PII) [12] proposes to use Poisson image blending to mitigate this by performing the interpolation on the gradients of the image, seamlessly blending the anomaly with the surrounding area. Naval et al. [13] took a simpler approach, softening the interpolation factor at the boundaries to make a more gradual change, as well as generating anomalies with randomly generated shapes to produce greater variety in the generated anomalies.

Contribution: In this work we introduce DISYRE: Diffusion-Inspired SYnthetic REstoration. By replacing the Gaussian noise corruption in diffusion models with a synthetic anomaly degradation process allows for directly using a diffusion-like model for unsupervised anomaly detection and localisation. It does this by learning a score function that generalizes to naturally-occurring medical anomalies and indicates how to modify a test image to bring it within the healthy distribution. Diffusion-like pipelines without noise have been previously

proposed for image generation [14], but to our knowledge this is the first approach specifically designed for UAD.

2. METHOD

Synthetic anomaly generation: We adopt the synthetic anomaly generation mechanism from [13], where a section of the image is corrupted replacing the original pixel intensities with a linear interpolation of original intensities and a foreign patch extracted randomly from the training set. Local synthetic anomalies are introduced by

$$x_t = (1 - \alpha \cdot m) \cdot x_0 + \alpha \cdot m \cdot x_{fp} \quad (1)$$

where $x_0 \in \mathbb{R}^{H \times W}$ denotes the original (healthy) image, $x_{fp} \in \mathbb{R}^{H \times W}$ denotes the foreign patch extracted from a healthy sample different from x_0 , $\alpha \in [0, 1]$ denotes the interpolation factor that controls the convex combination of x_0 and x_{fp} within the synthetic anomaly, and $m \in \mathbb{R}^{H \times W}$ denotes the mask component that both controls the shape of the synthetic anomaly and softens the interpolation factor towards the edges of the anomalous region. Note that $m_i \in [0, 1]$.

Forward and backward corruption process: In our diffusion-inspired proposed framework, we define a schedule to gradually corrupt the images increasing the α factor defined in Eq. 1 and train a model to revert the modified forward corruption process. We note that we do not use random noise, although synthetic anomalies are randomly generated.

Following the DDPM literature, we define the schedule α_t to increase with time steps t , *i.e.* when $t = 0$, $\alpha_t = 0$ and the image belongs to the expected distribution, while when

$t = T$, $\alpha_t = 1$ and the image is highly anomalous. We use the DDPM [7] schedule, however in our case it defines the interpolation of the original image and foreign patch instead of original image and noise. Note that our α_t is analogous to $1 - \sqrt{\alpha_t}$ in the diffusion literature. Given that we want interpolations x_t to remain in the range $[0, 1]$, we use a convex combination of image and foreign patch, so the weighting of the original image becomes $(1 - \alpha_t)$. In Fig. 1, we include an example of the synthetic corruption process that shows similarities with the standard diffusion process where textures / high frequencies are corrupted at small t values and structures / low frequencies are corrupted at high t values.

We train our model P_θ to restore corrupted images x_t into their healthy counterparts when conditioning on t , $x_0 \approx \hat{x}_0 = P_\theta(x_t, t)$. Network parameters are therefore optimized using the objective $\mathbb{E}_{t \sim [1, T], x_0, x_{fp}} (\|x_0 - P_\theta(x_t, t)\|^2)$.

Anomaly Score: At inference time, we propose to use the trained model P_θ to *restore* test images. This process is similar to the standard sampling of a DDPM [7], but crucially different in that we do not start from noise, but instead from a test image. Additionally, instead of iterating through all T steps, we skip steps based on a *step_size* hyper-parameter.

In order to localize anomalies we propose an Anomaly Score (*AS*) defined as the accumulation of the absolute gradients across all the restoration steps, so it is sensitive to the different degrees of abnormalities present in the corruption process. Algorithm 1 describes our *healing* procedure, with P as the trained model P_θ and Q the process specified in Eq. 1 which corrupts \hat{x}_o , interpolating it with the *unhealed* x_t using t_{next} and $m = 1$. In addition to the *AS*, the algorithm produces a restoration \hat{x}_o , a *healed* version of x_t . In our experiments we also evaluate using directly single-step restorations residuals, *i.e.*, $|P(x_t, t) - x_0|$, as anomaly score conditioning on different t values.

Algorithm 1 Inference: Image restoration

```

1:  $x_t \leftarrow x$ 
2:  $t \leftarrow T$ 
3:  $AS \leftarrow \text{zeros\_like}(x)$ 
4: while true do
5:    $\hat{x}_0 \leftarrow P(x_t, t)$ 
6:    $AS \leftarrow AS + |\hat{x}_0 - x_t|$ 
7:   if  $t_{next} \leq 0$  then
8:     break ▷ Exit the loop
9:    $t_{next} \leftarrow t - \text{step\_size}$ 
10:   $x_t \leftarrow Q(\hat{x}_0, x_t, t_{next}, m = 1)$ 
11:   $t \leftarrow t_{next}$ 
12: end

```

During training we use a patch-based pipeline. At inference time we use sliding window inference to obtain predictions for full axial slices. We hypothesized that the model is more accurate on the patches with a bigger proportion of foreground vs background. Consequently, we propose to

weight patch predictions by the foreground percentage prior to combine patch-predictions and report results both without (DISYRE) and with foreground weighting (DISYRE f.w.).

3. EXPERIMENTS

Experimental setup: We adopt the experimental setup for Brain MRI from [15], which includes the following datasets:

- Cambridge Centre for Ageing and Neuroscience dataset (CamCAN) [16]: T1- and T2-weighted images of 652 healthy, lesion-free, adult subjects.
- Anatomical Tracings of Lesions After Stroke (ATLAS) dataset [17]: T1-weighted scans of stroke patients ($N = 655$). This dataset contains annotated lesions after stroke.
- Multimodal Brain Tumor Segmentation (BraTS) Challenge dataset [18] 2020 edition: We used the T1- and T2-weighted images from the challenge training set. The dataset provides manual segmentations of gliomas.

We use the CamCAN dataset for training ($N = 602$), holding out a set with $N = 50$ subjects to identify training convergence. BraTS and ATLAS are used to evaluate UAD performance. We follow the registration and normalization protocols provided in [15], training our model using patches of 128×128 pixels obtained from non-empty axial slices.

For our training pipeline we use the default Brain MRI configuration of the task from [13]. We also adapted the same α_t schedule and network architecture proposed in DDPM but reduced the steps to $T = 100$. Our UNet has six down-sampling blocks with [32, 64, 96, 128, 256, 256] channels and attention layers from the 3rd block onwards. Networks are trained using the AdamW optimizer and OneCycle learning rate with maximum rate of $1e-4$ for 100,000 steps.

Performance analysis: To quantitatively evaluate our method, we use the most common metrics in the medical UAD literature: Average Precision (AP) and an estimate of the best possible Sørensen-Dice index ([Dice] score). Some qualitative results are reported in Fig. 3.

We evaluated single-step restorations to localize anomalies as a function of t , *i.e.*, $\|x - P(x, t)\|$. We also evaluated our cumulative multi-step restoration strategy using different *step_size* settings. Results are included in Table 1. Additionally, Fig. 2 compares single-step restoration AP with cumulative multi-step restoration performance.

Finally, Table 2 compares our results with baseline UAD methods from [15]. For our method DISYRE we report mean metrics for four training runs with multi-step restoration and *step_size* = 25. In addition to our method we also evaluate the method from [13], included in the table as *MOOD22*. Note that this approach uses the same synthetic anomaly generation process but binary cross-entropy with α as soft target. When implementing this baseline we use the same network architecture without conditioning on t .

	AP	ATLAS	BraTS-T1	BraTS-T2
single-step t	25	0.03 \pm 0.00	0.38 \pm 0.01	0.23 \pm 0.01
	50	0.20 \pm 0.02	0.44 \pm 0.02	0.45 \pm 0.01
	75	0.29 \pm 0.03	0.34 \pm 0.03	0.59 \pm 0.03
	100	0.25 \pm 0.03	0.28 \pm 0.02	0.70 \pm 0.02
multi-step $step_size$	10	0.27 \pm 0.02	0.32 \pm 0.02	0.73 \pm 0.02
	20	0.29 \pm 0.02	0.34 \pm 0.02	0.74 \pm 0.02
	25	0.29 \pm 0.02	0.34 \pm 0.02	0.75 \pm 0.02
	33	0.29 \pm 0.02	0.35 \pm 0.02	0.75 \pm 0.01
	50	0.30 \pm 0.03	0.35 \pm 0.02	0.75 \pm 0.02

Table 1. Upper section: AP for single-step restorations ($AS = \|x - P(x, t)\|$) as a function of the t used to condition the model. Lower section: AP for multi-step restorations, as a function of $step_size$. Mean and std. dev. on 4 seeds.

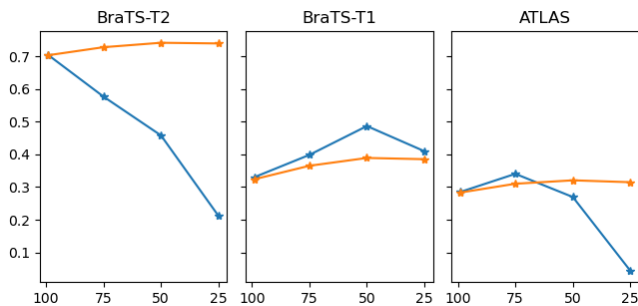


Fig. 2. AP profiles as a function of t . Blue line shows single-step restoration. Orange line shows the cumulative AS as defined in Section 2 up-to t stage ($step_size = 25$). Single seed.

Method	ATLAS		BraTS-T1		BraTS-T2		
	AP	[Dice]	AP	[Dice]	AP	[Dice]	
IR	VAE [2]	0.11	0.20	0.13	0.19	0.28	0.33
	r-VAE [5]	0.09	0.17	0.13	0.19	0.36	0.40
	f-AnoGAN [3]	0.02	0.06	0.06	0.12	0.15	0.21
FM	FAE [19]	0.08	0.18	0.42	0.45	0.51	0.52
	RD [20]	0.11	0.22	0.36	0.42	0.47	0.50
AB	AMCons [21]	0.01	0.03	0.05	0.12	0.35	0.40
S-S	PII [12]	0.03	0.07	0.13	0.22	0.13	0.22
	DAE [9]	0.05	0.13	0.13	0.20	0.47	0.49
	MOOD22 [13]	0.10	0.21	0.24	0.31	0.47	0.48
	DISYRE (Ours)	0.19	0.27	0.22	0.30	0.69	0.67
	DISYRE f.w. (Ours)	0.29	0.37	0.34	0.41	0.75	0.70

Table 2. Localization results of the image-reconstruction (IR), feature-modeling (FM), attention-based (AB), and self-supervised (S-S) methods. Best scores are bold.

4. DISCUSSION & CONCLUSION

Results in Table 2 highlight the effectiveness of our strategy: DISYRE achieves an AP of 0.75 in BraTS-T2 dataset, an increase of 0.24 from the second best method. DISYRE also improves the state of the art of the ATLAS task to 0.29, up from 0.19 while achieving competitive results in BraTS-T1.




Fig. 3. Restoration results on test images drawn from the ATLAS [17], BraTS-T1, and BraTS-T2 [18], after $step_size = 25$, along with anomaly scores AS in grayscale compared with the ground truth outline in red.

When using single-step restorations we found that AP peaks at different t values for different anomalies and modalities (Fig. 2). We hypothesize that these profiles are associated with how the anomalies manifest in each image modality, with high-frequency anomalies being restored at lower t values and more structural corruptions restored as t approaches T . We aim to be sensitive to all classes of anomaly, from structural to high frequency, which is why we integrate multiple steps into the proposed AS . Despite improving results in BraTS-T2 and ATLAS, multi-step AS does not improve single-step AS in BraTS-T1, potentially due to the nature of how gliomas manifest in T1-weighted Brain MRI. In future work we will explore alternative approaches to achieve better coverage of the anomaly spectrum.

As seen in Table 1, DISYRE is robust to the choice of the $step_size$ hyper-parameter. We opted for $step_size = 25$ in our final experiments as it provides a good balance between coverage of the anomaly schedule and inference speed.

In this work we propose DISYRE, a new strategy that adapts diffusion models to UAD. With the motivation of learning a score function $\nabla_x \log p(x)$ which generalizes to naturally occurring medical anomalies, we replace the standard Gaussian noise corruption with gradual, synthetic anomalies. DISYRE opens a new route to leverage diffusion-like models in medical image analysis; we show that forward processes can be designed so they become relevant for downstream tasks. In DISYRE we leverage synthetic anomalies for UAD and substantially outperform other UAD baseline methods in two out of three datasets.

5. COMPLIANCE WITH ETHICAL STANDARDS

This research study was conducted retrospectively using human subject data made available in open access by [16, 17, 18]. Ethical approval was not required as confirmed by the license attached with the open-access data.

6. ACKNOWLEDGMENTS

Bernhard Kainz received support by the ERC - project MIA-NORMAL 101083647, Matthew Baugh by a UKRI DTP award.

7. REFERENCES

- [1] Geert Litjens et al., “A Survey on Deep Learning in Medical Image Analysis,” *Medical image analysis*, vol. 42, pp. 60–88, 2017.
- [2] Christoph Baur et al., “Autoencoders for Unsupervised Anomaly Segmentation in Brain MR Images: A Comparative Study,” *Medical image analysis*, , no. 02 Jan 2021, pp. 69:101952, 2021.
- [3] Thomas Schlegl et al., “f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks,” *Medical image analysis*, vol. 54, pp. 30–44, 2019.
- [4] Cosmin I. Bercea et al., “Reversing the Abnormal: Pseudo-Healthy Generative Networks for Anomaly Detection,” in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*, Cham, 2023, pp. 293–303.
- [5] Xiaoran Chen et al., “Unsupervised lesion detection via image restoration with a normative prior,” *Medical Image Analysis*, vol. 64, pp. 101713, Aug. 2020.
- [6] Sergio Naval Marimont and Giacomo Tarroni, “Anomaly detection through latent space restoration using vector quantized variational autoencoders,” in *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, 2021.
- [7] Jonathan Ho et al., “Denoising diffusion probabilistic models,” *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [8] Julian Wyatt et al., “AnoDDPM: Anomaly Detection with Denoising Diffusion Probabilistic Models using Simplex Noise,” in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, New Orleans, LA, USA, June 2022, pp. 649–655.
- [9] Antanas Kascenas et al., “Denoising autoencoders for unsupervised anomaly detection in brain mri,” in *International Conference on Medical Imaging with Deep Learning*. PMLR, 2022, pp. 653–664.
- [10] Julia Wolleb et al., “Diffusion Models for Medical Anomaly Detection,” *arXiv preprint arXiv:2203.04306*.
- [11] Jeremy Tan et al., “Detecting outliers with foreign patch interpolation,” *Machine Learning for Biomedical Imaging*, vol. 1, no. April 2022 issue, pp. 1–27, 2022.
- [12] Jeremy Tan et al., “Detecting Outliers with Poisson Image Interpolation,” in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*, Cham, 2021.
- [13] Sergio Naval Marimont and Giacomo Tarroni, “Achieving state-of-the-art performance in the Medical Out-of-Distribution (MOOD) challenge using plausible synthetic anomalies,” Nov. 2023, arXiv:2308.01412 [cs].
- [14] Arpit Bansal et al., “Cold Diffusion: Inverting Arbitrary Image Transforms Without Noise,” Aug. 2022, arXiv:2208.09392 [cs].
- [15] Ioannis Lagogiannis et al., “Unsupervised Pathology Detection: A Deep Dive Into the State of the Art,” *IEEE Transactions on Medical Imaging*, pp. 1–1, 2023, arXiv:2303.00609 [cs].
- [16] Jason R Taylor et al., “The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) data repository: Structural and functional MRI, MEG, and cognitive data from a cross-sectional adult lifespan sample,” *neuroimage*, vol. 144, pp. 262–269, 2017.
- [17] Sook-Lei Liew et al., “A large, open source dataset of stroke anatomical brain images and manual lesion segmentations,” *Scientific data*, vol. 5, pp. 1–11, 2018.
- [18] Bjoern H Menze et al., “The multimodal brain tumor image segmentation benchmark (brats),” *IEEE transactions on medical imaging*, vol. 34, no. 10, pp. 1993–2024, 2014.
- [19] Felix Meissen et al., “Unsupervised Anomaly Localization with Structural Feature-Autoencoders,” in *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*, Cham, 2023, pp. 14–24.
- [20] Hanqiu Deng and Xingyu Li, “Anomaly detection via reverse distillation from one-class embedding,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 9737–9746.
- [21] Julio Silva-Rodríguez et al., “Constrained unsupervised anomaly segmentation,” *Medical Image Analysis*, vol. 80, pp. 102526, Aug. 2022, arXiv:2203.01671 [cs, eess].