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.

Selective Test-Time Adaptation for Unsupervised Anomaly Detection using Neural Implicit Representations

Sameer Ambekar[1,2], Julia Schnabel*[1,2,3], Cosmin Bereca*[1,2]

[1] School of Computation, Information and Technology, Technical University of Munich, Germany
[2] Institute of Machine Learning in Biomedical Imaging, Helmholtz Munich, Germany
[3] School of Biomedical Engineering and Imaging Sciences, King's College London, UK

Abstract. Deep learning models in medical imaging often encounter challenges when adapting to new clinical settings unseen during training. Test-time adaptation offers a promising approach to optimize models for these unseen domains, yet its application in anomaly detection (AD) remains largely unexplored. AD aims to efficiently identify deviations from normative distributions; however, full adaptation, including pathological shifts, may inadvertently learn the anomalies it intends to detect. We introduce a novel concept of *selective* test-time adaptation that utilizes the inherent characteristics of deep pre-trained features to adapt *selectively* in a zero-shot manner to any test image from an unseen domain. This approach employs a model-agnostic, lightweight multi-layer perceptron for neural implicit representations, enabling the adaptation of outputs from any reconstruction-based AD method without altering the source-trained model. Rigorous validation in brain AD demonstrated that our strategy substantially enhances detection accuracy for multiple conditions and different target distributions. Specifically, our method improves the detection rates by up to 78% for enlarged ventricles and 24% for edemas. Our code is available: https://github.com/compai-lab/2024-miccai-adsmi-ambekar.

Keywords: model generalization · zero-shot learning

1 Introduction

Unsupervised anomaly detection (UAD) commonly starts by training models on large datasets comprising control patients presumed to be healthy. These models are subsequently deployed in various clinical environments to detect deviations from expected normative anatomical structures [4,6,24,34].

However, these models tend to underperform when deployed in new, unseen clinical environments due to various distribution changes, such as those caused by different imaging equipment, often referred to as scanner shifts [5,23]. Even though it is common practice to train on datasets different from those in deployment

*Shared last-authorship

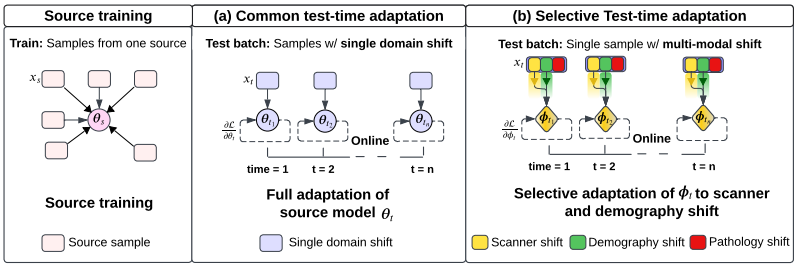

Fig. 1: *Selective* test-time adaptation. (a) Common test-time adaptation use the source model $\boldsymbol{\theta}_s$ to obtain $\boldsymbol{\theta}_t$ by fully adapting to $\mathbf{x}_t$ through iterative updates. (b) In contrast, we propose a new concept of *selective* test-time adaptation. We learn $\boldsymbol{\phi}_t$ and adapt *selectively* to $\mathbf{x}_t$, implicitly excluding pathology shifts.

settings, the resulting domain shifts are largely unaddressed. A common yet flawed approach to mitigate this issue involves using 2D slices from 3D pathological scans labeled for their absence of pathology [31]. This approach, however, presents two critical challenges: the slices may inadvertently contain pathological traces, and there is a risk of data leakage when slices from the same patient are used to represent both healthy and pathological conditions. An alternative that seems straightforward but is complex in implementation is to fine-tune the model on the target data distribution [18,30]. While effective, this strategy demands extensive data on control patients, which many specialized clinics do not possess [27,2].

To address this issue, a recent approach known as "test-time adaptation" has emerged [3,16,28,30], which is applied to adapt to target distributions affected by domain-shifts [19,30], or in generating target datasets from a few samples [8,22,32]. In this setting, the model is trained only on data from the source domain and then adapted to make predictions on unseen target data. However, a significant challenge in AD lies in adjusting for unknown variations caused by different scanners and patient demographics while preserving the sensitivity needed to detect subtle pathological changes. This necessitates a *selective* adaptation approach, where the algorithm adjusts to recognize and compensate for non-pathological variations without compromising its ability to identify anomalies. Current adaptation methods [18,16,30] often struggle in these scenarios because they either fail to address multi-modal distribution shifts, or they fully adapt to target data [16,18,22,30,32] as shown in Figure 1(a).

In this paper, we introduce a novel *selective* test-time adaptation framework for anomaly detection *(STA-AD)*, see Figure 1(b). Our framework is designed to enhance existing AD methods by adapting them to new target distributions on pathological scans in a zero-shot manner. In summary, our contributions are:

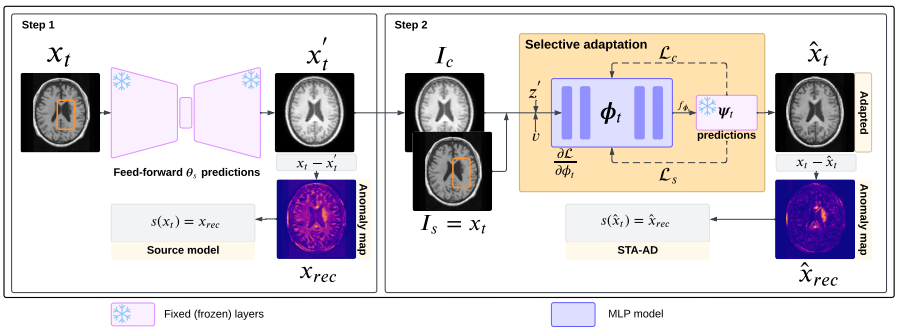

Fig. 2: **STA-AD at test-time.** We leverage source model predictions $\mathbf{x}_t^{'}$ as "content" and target images $\mathbf{x}_t$ as "style" to train neural implicit representations $\boldsymbol{\phi}_t$ to adapt images $\hat{\mathbf{x}}_t$. The initial false positives due to domain shifts seen in $\mathbf{x}_{rec}$ are removed after adaptation ($\hat{\mathbf{x}}_{rec}$), while the detected lesion is highlighted.

- We propose the novel concept of *selective* test-time adaptation for AD. To the best of our knowledge, this innovation is the first to enable AD models to adjust for domain shifts directly on pathological target data.
- Our method leverages an efficient neural implicit learning model that bypasses the need for extensive retraining or large curated datasets, thereby considerably simplifying deployment in clinical settings.
- We demonstrate the efficacy of our *STA-AD* framework through extensive validation, resulting in considerable improvements in detection accuracy across various conditions and target distributions.

2 Materials and methods

Let $\mathcal{X}$ be the space of all images $\mathbf{x} \in \mathbb{R}^n$, where n could be 2 or 3-dimensional. AD networks are trained on a source domain $\mathcal{D}_s = \{(\mathbf{x}_s)^i\}_{i=1}^{N_s}$ containing a large number of control samples (N_s). The method can then be deployed to multiple target distributions $\mathcal{D}_t = \{(\mathbf{x}_t)^i\}_{i=1}^{N_t}$. A generative source model indicated as $\boldsymbol{\theta}_s$ is trained on source domains $\mathcal{D}_s$, by minimizing the common reconstruction losses ($\mathcal{L}$). This process is formalized, where $\mathcal{L}$ consists of generative loss functions [6,9] aimed at obtaining the oracle model, represented as $\boldsymbol{\theta}_s = \arg\min_{\boldsymbol{\theta}} \mathbb{E}_{(\mathbf{x}_s)\in\mathcal{D}_s}[\mathcal{L}(\mathbf{x}_s; \boldsymbol{\theta})]$.

If pathological shifts were the only discrepancies between the domains $\mathcal{D}_s$ and $\mathcal{D}_t$, the methods would accurately identify these deviations as anomalies. However, in reality, $\mathcal{D}_t$ contains multiple scanner and demographic shifts, as shown in Figure 1. These affect the performance of AD (see $\mathbf{x}_{rec}$ in Figure 2). Furthermore, the target domain $\mathcal{D}_t$ may consist solely of pathological samples, such as in specialized clinics for tumor detection. Therefore, *selective* adaptation is crucial. By adjusting to non-pathological distributional shifts while disregarding pathologies, we ensure robust anomaly detection across diverse domains.

Neural implicit representation learning allows the parameterization of signals utilizing neural networks [10,12,20]. Their compactness, resolution independence and inherent differentiability make them particularly advantageous for test-time adaptation. In our work, we utilize them to model the adapted image function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$, where n represents the dimensionality of the adapted image's flattened vector, $\hat{x}_t$. Here, we aim to learn a neural network $F_{\boldsymbol{\phi}} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ parameterized by weights $\boldsymbol{\phi}$, such that $F_{\boldsymbol{\phi}}$ approximates f as closely as possible.

***Selective* test-time adaptation.** We utilize [17] for the test-time training. We propose to *selectively* adapt to specific distribution shifts, e.g., scanner shifts with a single sample from target data. In this framework, a single target sample is treated as the style image (I_s), and its corresponding prediction from a source-trained model as the content image (I_c). By utilizing the content from the source model, we mitigate adapting to anomalies from the target domain. Using I_c and I_s as inputs, we train $\boldsymbol{\phi}_t$ at test-time to generate the target-adapted image ($\hat{\mathbf{x}}_t$), as shown in Figure 2. Initially, we generate z_c and z_s vectors with random normal distribution, along with a random flattened vector v_r. These latent vectors are then interpolated to produce a new latent vector $z' = \text{interpolate}(z_c, z_s)$. This vector z', combined with v_r serves as input for $\boldsymbol{\phi}_t$:

$$I_t = \boldsymbol{\phi}_t(z', (v)). \tag{1}$$

The $\boldsymbol{\phi}_t$ generates the required target-adapted image (I_t), which is then optimized using a pre-trained VGG model (ψ). This optimization process involves minimizing the content and style losses:

$$\mathcal{L}_{\text{cont}}(\psi(\boldsymbol{\phi}_t(z'(i), v; \boldsymbol{\phi}_t)), F_c) + \mathcal{L}_{\text{style}}(\psi(\boldsymbol{\phi}_t(z'(i), v; \boldsymbol{\phi}_t)), F_s) \tag{2}$$

where $\mathcal{L}_{\text{cont}} = \|\psi(\boldsymbol{\phi}_t(z', v; \boldsymbol{\phi}_t)) - F_c\|^2$ is the L2 norm between the adapted image features and content features. Similarly, the style loss $\mathcal{L}_{\text{style}} = \|g(\psi(\boldsymbol{\phi}_t(z', v; \boldsymbol{\phi}))) - g(F_s)\|^2$ is the L2 norm of differences in Gram matrices (g) of the features from the adapted image and the style image at various layers [13] of the ψ network. We optimize $\boldsymbol{\phi}_t$ by updating it to a new state $\boldsymbol{\phi}_t^*$ through the minimization of the content and style losses, which are weighted by β and γ respectively, using learning rate λ:

$$\begin{aligned}\boldsymbol{\phi}_t^* = \boldsymbol{\phi}_t - \lambda \nabla_{\boldsymbol{\phi}_t} \big(&\beta * \mathcal{L}_{\text{cont}}(\psi(\boldsymbol{\phi}_t(z', v; \boldsymbol{\phi}_t)), F_c) \\ &+ \gamma * \mathcal{L}_{\text{style}}(\psi(\boldsymbol{\phi}_t(z', v; \boldsymbol{\phi}_t)), F_s)\big)\end{aligned} \tag{3}$$

Finally, we assign $\hat{\mathbf{x}}_t = I_t$ to obtain the target-adapted image. Note that this process optimizes deep semantic features extracted from specific layers of a pre-trained VGG model [13]. This acts as a surrogate loss [11,21] to prevent adaptation to local/non-smooth deformations, which in our context represent pathological representations. For all experiments, we train all the parameters of the $\boldsymbol{\phi}_t$ model on single test samples using Eqn. 3. Finally, we compute the anomaly map between the input image ($\mathbf{x}_t$) and adapted image ($\hat{\mathbf{x}}_t$) by $s(\hat{\mathbf{x}}_t) = |x_t - \hat{x}_t|$.

3 Experiments and Results

In subsection 3.1, we explore potential distribution shifts between healthy samples from the source and target datasets and apply our approach to mitigate them. In subsection 3.2, assesses whether our *selective* adaptation effectively adjusts to anomalous target samples and evaluates its impact on the downstream AD tasks. Additionally, we introduce a parameter-efficient variant of $\boldsymbol{\phi}_t$, which trains only 1% of the total parameters of $\boldsymbol{\phi}_t$ yet delivers competitive results. Moreover, aligning with the principle that reducing entropy is a core objective of adaptation [25], we demonstrate that our adapted model minimizes entropy post-adaptation.

Datasets. We evaluate our method on public T1-weighted brain Magnetic Resonance Imaging datasets that inherently contain domain shifts due to different scanner vendors and variations in acquisition protocols. The pre-processing of the datasets follows common baselines [6]. We denote IXI [1] as the source dataset (S) containing healthy scans. The target dataset (T_+) is based FastMRI+ [33] containing 13 different anomaly types acquired from various vendors, reminiscent of multi-target adaptation. At test-time, we address these samples from the multiple target datasets with our single sample adaptation capabilities.

Implementation details. Algorithm 1 outlines our approach. We train the $\boldsymbol{\phi}$ network at test time with a single sample per batch in an online manner.
We do not modify the source training process; instead, we utilize the original source model and introduce test-time learning of our MLP network. For our experiments, we use two anomaly detection backbones: RA [6], a generative autoencoder and DDPM [14], a diffusion-based model. We keep their original training protocols and freeze their model weights after the source training phase. The implementation of $\boldsymbol{\phi}$ is an MLP with a 16-layer depth and 128-unit width [17], and ψ is a pre-trained VGG [26] model pre-trained on ImageNet weights. For Table 1, we calculate the evaluation metrics between the input image (x_t) from the test domain and our selectively adapted image ($\hat{x}_t$). For Table 2, we follow [6]. For the additional experiments, we utilize the same MLP with added batch normalization layers and reduce the number of parameters to be trained by only optimizing these layers. Moreover, both $\boldsymbol{\phi}_t$ variants are compute-friendly due to their tiny model capacity and applicable to any existing AD models.

3.1 Test-time adaptation on healthy samples

We evaluate whether non-pathological distributional shifts are present between healthy samples from the source distribution and target distributions T_+ in Table 1. The source model performs well on source data but struggles to generalize to the target domains T_+, as evidenced by nearly halved SSIM scores and higher perceptual errors. Similarly, as shown in Figure 3, the RA and DDPM models exhibit increased false positives in their respective anomaly maps. As a baseline, we consider histogram matching. This method achieves slightly higher MAE and SSIM scores than the source model but achieves lower LPIPS scores, which is

Table 1: **Comparisons for healthy samples from different domains**. The source model naturally performs well on source data. However, on unseen target data, the performance degrades by lower SSIM scores, higher MAE, and higher LPIPS scores. We show the best results in **bold** and underline the runner-up.

Domain	Method	MAE ↓	SSIM ↑	LPIPS ↓
S	RA [6]	0.060 ±0.005	0.702 ±0.028	0.049 ±0.007
	+ STA-AD (ours)	**0.047** ±0.020	0.680 ±0.050	**0.036** ±0.017
	DDPM [14]	0.050 ±0.003	0.753 ±0.024	0.030 ±0.004
	+ STA-AD (ours)	**0.045** ±0.003	**0.791** ±0.01	**0.021** ±0.002
T_+ Multi-target	RA [6]	0.127 ±0.035	0.465 ±0.077	0.084 ±0.028
	+ Histogram matching	0.081 ±0.020	0.535 ±0.069	0.094 ±0.032
	+ STA-AD (ours)	**0.078** ±0.021	**0.600** ±0.070	**0.047** ±0.023
	DDPM [14]	0.080 ±0.008	0.556 ±0.050	0.078 ±0.021
	Histogram matching	**0.058** ±0.008	**0.632** ±0.045	0.095 ±0.026
	+ STA-AD (ours)	0.104 ±0.018	0.473 ±0.070	**0.071** ±0.018

also evident through the generated image artifacts provided in Supplementary. In contrast, our method consistently improves the MAE, SSIM, and LPIPS scores compared to the source model and histogram matching on both healthy source and pathological target datasets. For the DDPM backbone adapted to the target data, our method achieves improved LPIPS scores but lower MAE and SSIM values. This may be because our method intentionally preserves the original content of the source images, which might result in pixel-wise content differences. We aim to keep the source content and align the source style to the target domain.

3.2 Selective Test-time adaptation for anomaly detection

We evaluate the performance of our approach in detecting the various conditions from the T_+ domain in Table 2 and Figure 4. The detection task consists of 13 various conditions with one or more target images per condition. We measure the performance with true positives (TP) and F1 scores. We show the average

Algorithm 1 Selective test-time adaptation.
Input: $\mathcal{T}$: target domain; learned and frozen $\boldsymbol{\theta}_s, \boldsymbol{\phi}_t$: MLP;
Output: Adapted model parameters $\boldsymbol{\phi}_t^*$

Draw a single sample $\mathbf{x}_t$, set as I_s
Get $I_c = f_{\boldsymbol{\theta}}(x_t)$ and initialize $\mathbf{z}_c, \mathbf{z}_s, V_r$
From $\boldsymbol{\phi}_t$ generate $\hat{x}_t$. (Eqn.1)
Update $\boldsymbol{\phi}_t$ with loss from pretrained layers of ψ (Eqn.2)
Optimize $\boldsymbol{\phi}_t^* = \boldsymbol{\phi}_t - \lambda \nabla_{\boldsymbol{\phi}_t} (\beta * \mathcal{L}_{\text{cont}} + (\gamma * \mathcal{L}_{\text{style}})$ (Eqn.3)
Set $\hat{\mathbf{x}}_t = I_t$; compute anomaly map $(\overline{x_t}) = |x_t - \hat{x_t}|$.

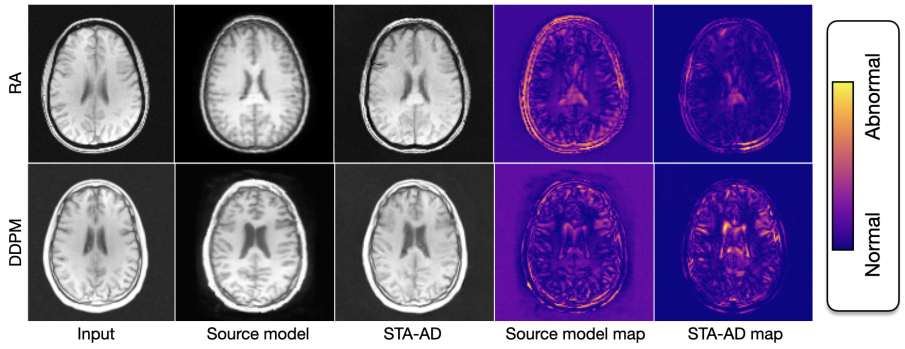

Fig. 3: **Comparison on healthy samples from T_+ dataset.** For both backbones, the source models achieve higher false positives due to inaccurate reconstructions. Our approach, *STA-AD*, through the adapted image, reduces false positive detections in the presented anomaly maps.

Table 2: **Comparisons for anomaly detection on unhealthy samples.** We shot the best results in **bold**. Our method outperforms both of the baselines and achieves up to 87% improvement in average for the diffusion-based model.

Method	Avg.		Edema		E. Ventricles		Encephalomalacia	
	TP ↑	F1 ↑	TP ↑	F1 ↑	TP ↑	F1 ↑	TP ↑	F1 ↑
RA [7]	0.51 ±0.00	0.15 ±0.30	0.72 ±0.00	0.25 ±0.30	0.47 ±0.30	0.32 ±0.01	1.00 ±0.00	0.15 ±0.01
+STA-AD (ours)	**0.70** ±0.30	**0.23** ±0.30	**0.83** ±0.30	**0.40** ±0.30	**0.84** ±0.36	**0.52** ±0.20	1.00 ±0.30	**0.40** ±0.01
DDPM [14]	0.39 ±0.40	0.11 ±0.20	0.67 ±0.40	0.32 ±0.30	0.26 ±0.30	0.15 ±0.00	0.00 ±0.00	0.00 ±0.00
+STA-AD (ours)	**0.73** ±0.40	**0.19** ±0.30	**0.83** ±0.30	**0.33** ±0.20	**0.47** ±0.40	**0.25** ±0.30	**1.00** ±0.00	**0.29** ±0.00

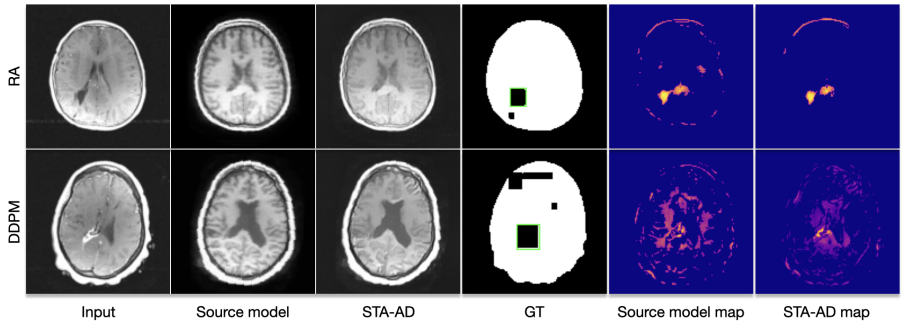

Fig. 4: **Comparisons with different backbones on T_+ unhealthy samples.** The *STA-AD* framework can enhance the performance of AD methods by reducing false positives and accurately detecting anomalies after target adaptation.

performance (Avg.) and three exemplary conditions. The detailed results on all conditions are presented in the Supplementary.

Table 3: **Entropy minimization.**

	Healthy ↓	Unhealthy ↓
RA	13.9	13.9
STA-AD	**6.8**	**6.7**

Table 4: **Efficiency analysis.**

	Parameters	**MAE ↓**	**LPIPS ↓**
w/o BN	600000	0.078 ±0.02	0.04 ±0.02
w/ BN	4096	0.141 ±0.02	0.12 ±0.03

Our method demonstrates an average 87% increase in true positives using the DDPM backbone and an average 37% increase with the RA backbone, alongside improvements of 53% and 72% in overall F1 scores, respectively. Specifically, for the condition of encephalomalacia, the baseline DDPM model notably struggled, missing the anomaly. Our adapted approach successfully detected the true positive with a high accuracy of 1.0, considerably enhancing the F1 score. Similarly, for structural anomalies such as the enlarged ventricles, our method uniformly boosts true positives and F1 scores across both backbones. This pattern of increasing true positives and consistently improving F1 scores extends to other conditions as well, emphasizing our approach's effectiveness in enhancing disease detection rates and specificity through adaptation. Figure 4 contains the visual examples. Our method significantly reduces false positives while maintaining true positive detection, such as detecting lesions close to the left posterior lateral ventricle. However, the second row shows the limitations of our approach. Although we still reduce false positive detections, the DDPM reconstruction differs too much in content from the input, resulting in suboptimal anomaly detection.

Implicit entropy minimization. Following the principle that entropy minimization facilitates adaptation [16,30,29,25], we also investigate the entropy error [25] on the T_+ dataset. We report the entropy errors (lower is better) on healthy and unhealthy variants of T_+. As shown in Table 3, the predictions of source model 'RA' obtain high entropy error, while our proposed approach reduces the entropy error considerably after adaptation. By nearly halving the entropy error, our approach achieves significantly more precise adaptation.

Efficiency analysis. Our approach optimizes the ϕ_t model with a total of 600,000 parameters in about 4 minutes per single sample. We also propose a variant that requires optimizing just 1% (4096 parameters) of the total parameters by adjusting the batch norm affine parameters [15] and still achieving competitive performance on the T_+ dataset, as detailed in Table 4. Both proposed methods are highly efficient, requiring only one GPU for training due to their compact model sizes of 600,000 and 6,000 parameters, respectively. This demonstrates the practicality of our approach, especially in environments with limited computational resources.

4 Conclusion

In this work, we proposed a *selective* adaptation paradigm, *STA-AD*, that allows the adaptation of any reconstruction-based anomaly detection (AD) method directly to target pathological domains, thereby improving their performance. We validated our approach on a multi-target brain MRI dataset containing multiple lesions and demonstrated performance boosts of up to 87%. We believe our work will encourage future research in adapting AD methods to unseen pathological distributions. Further improvements in inference times and reducing reliance on source model content predictions could enhance its clinical benefits, making it a valuable tool for deploying AD methods in various clinical settings.

Acknowledgements

This paper is supported by the DAAD programme Konrad Zuse Schools of Excellence in Artificial Intelligence, sponsored by the Federal Ministry of Education and Research. C.I.B. is funded via the EVUK program ("Next-generation AI for Integrated Diagnostics") of the Free State of Bavaria and partially supported by the Helmholtz Association under the joint research school 'Munich School for Data Science'. We thank Lina Felsner for her assistance with the diagrams.

References

1. IXI Dataset, Information eXtraction from Images (EPSRC GR/S21533/02). https://brain-development.org/ixi-dataset/ (2012), accessed: 2023-01-12
2. Ambekar, S., Ankit, A., van der Mast, D., Alence, M., Tafuro, M.: Skdcgn: Source-free knowledge distillation of counterfactual generative networks using cgans. arXiv preprint arXiv:2208.04226 (2022)
3. Ambekar, S., Xiao, Z., Shen, J., Zhen, X., Snoek, C.G.: Learning variational neighbor labels for test-time domain generalization. arXiv preprint arXiv:2307.04033 (2023)
4. Behrendt, F., Bhattacharya, D., Krüger, J., Opfer, R., Schlaefer, A.: Patched diffusion models for unsupervised anomaly detection in brain mri. International Conference on Medical Imaging with Deep Learning (2023)
5. Bercea, C.I., Puyol-Antón, E., Wiestler, B., Rueckert, D., Schnabel, J.A., King, A.P.: Bias in unsupervised anomaly detection in brain mri. In: Workshop on Clinical Image-Based Procedures. pp. 122–131. Springer (2023)
6. Bercea, C.I., Wiestler, B., Rueckert, D., Schnabel, J.A.: Generalizing unsupervised anomaly detection: Towards unbiased pathology screening. In: Medical Imaging with Deep Learning (2023)
7. Bercea, C.I., Wiestler, B., Rueckert, D., Schnabel, J.A.: Reversing the abnormal: Pseudo-healthy generative networks for anomaly detection. arXiv preprint arXiv:2303.08452 (2023)
8. Chen, X., Li, Y., Yao, L., Adeli, E., Zhang, Y., Wang, X.: Generative adversarial u-net for domain-free few-shot medical diagnosis. Pattern Recognition Letters **157**, 112–118 (2022)

9. Daniel, T., Tamar, A.: Soft-introvae: Analyzing and improving the introspective variational autoencoder. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4391–4400 (2021)
10. De Luigi, L., Cardace, A., Spezialetti, R., Ramirez, P.Z., Salti, S., Di Stefano, L.: Deep learning on implicit neural representations of shapes. arXiv preprint arXiv:2302.05438 (2023)
11. Grabocka, J., Scholz, R., Schmidt-Thieme, L.: Learning surrogate losses. arXiv preprint arXiv:1905.10108 (2019)
12. Grattarola, D., Vandergheynst, P.: Generalised implicit neural representations. Advances in Neural Information Processing Systems **35**, 30446–30458 (2022)
13. Härkönen, E., Hertzmann, A., Lehtinen, J., Paris, S.: Ganspace: Discovering interpretable gan controls. Advances in neural information processing systems **33**, 9841–9850 (2020)
14. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in neural information processing systems **33**, 6840–6851 (2020)
15. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: International conference on machine learning. pp. 448–456. pmlr (2015)
16. Iwasawa, Y., Matsuo, Y.: Test-time classifier adjustment module for model-agnostic domain generalization. In: Advances in Neural Information Processing Systems. vol. 34 (2021)
17. Kim, S., Min, Y., Jung, Y., Kim, S.: Controllable style transfer via test-time training of implicit neural representation. Pattern Recognition **146**, 109988 (2024)
18. Liang, J., Hu, D., Feng, J.: Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In: International Conference on Machine Learning. pp. 6028–6039. PMLR (2020)
19. Liu, Y., Kothari, P., van Delft, B., Bellot-Gurlet, B., Mordan, T., Alahi, A.: Ttt++: When does self-supervised test-time training fail or thrive? In: Advances in Neural Information Processing Systems. vol. 34 (2021)
20. Molaei, A., Aminimehr, A., Tavakoli, A., Kazerouni, A., Azad, B., Azad, R., Merhof, D.: Implicit neural representation in medical imaging: A comparative survey. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2381–2391 (2023)
21. Nguyen, X., Wainwright, M.J., Jordan, M.I.: On surrogate loss functions and f-divergences. The Annals of Statistics **37**(2), 876 – 904 (2009). https://doi.org/10.1214/08-AOS595, https://doi.org/10.1214/08-AOS595
22. Ojha, U., Li, Y., Lu, J., Efros, A.A., Lee, Y.J., Shechtman, E., Zhang, R.: Few-shot image generation via cross-domain correspondence. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10743–10752 (2021)
23. Pandey, P., Kyatham, V., Mishra, D., Dastidar, T.R., et al.: Target-independent domain adaptation for wbc classification using generative latent search. IEEE Transactions on Medical Imaging **39**(12), 3979–3991 (2020)
24. Pinaya, W.H., Tudosiu, P.D., Gray, R., Rees, G., Nachev, P., Ourselin, S., Cardoso, M.J.: Unsupervised brain imaging 3d anomaly detection and segmentation with transformers. Medical Image Analysis **79**, 102475 (2022)
25. Shannon, C.E.: A mathematical theory of communication. The Bell system technical journal **27**(3), 379–423 (1948)
26. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)

27. Song, J., Lee, J., Kweon, I.S., Choi, S.: Ecotta: Memory-efficient continual test-time adaptation via self-distilled regularization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11920–11929 (2023)
28. Sun, Y., Wang, X., Liu, Z., Miller, J., Efros, A., Hardt, M.: Test-time training with self-supervision for generalization under distribution shifts. In: International Conference on Machine Learning. pp. 9229–9248. PMLR (2020)
29. Vu, T.H., Jain, H., Bucher, M., Cord, M., Pérez, P.: Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2517–2526 (2019)
30. Wang, D., Shelhamer, E., Liu, S., Olshausen, B., Darrell, T.: Tent: Fully test-time adaptation by entropy minimization. In: International Conference on Learning Representations (2021)
31. Wolleb, J., Bieder, F., Friedrich, P., Zhang, P., Durrer, A., Cattin, P.C.: Binary noise for binary tasks: Masked bernoulli diffusion for unsupervised anomaly detection. arXiv preprint arXiv:2403.11667 (2024)
32. Xiao, J., Li, L., Wang, C., Zha, Z.J., Huang, Q.: Few shot generative model adaption via relaxed spatial structural alignment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11204–11213 (2022)
33. Zhao, R., Yaman, B., Zhang, Y., Stewart, R., Dixon, A., Knoll, F., Huang, Z., Lui, Y.W., Hansen, M.S., Lungren, M.P.: fastmri+: Clinical pathology annotations for knee and brain fully sampled multi-coil mri data. arXiv preprint arXiv:2109.03812 (2021)
34. 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. pp. 289–297. Springer (2019)