TracerAD: Training-Free Few-Shot 3D Anomaly Detection for Novel PET Tracers

Haolin Huang[1†], Junlei Wu[1†], Jiaying Lu[2], Zhenrong Shen[1], Xinyu Wang[1], Chuantao Zuo[2], and Qian Wang[1,3](✉)

[1] School of Biomedical Engineering & State Key Laboratory of Advanced Medical Materials and Devices, ShanghaiTech University, Shanghai, China
qianwang@shanghaitech.edu.cn
[2] Department of Nuclear Medicine/PET center, Huashan Hospital, Fudan University, Shanghai, China
[3] Shanghai Clinical Research and Trial Center, Shanghai, China

Abstract. The deployment of PET tracers faces a cold-start dilemma in two common scenarios: when a novel tracer enters early clinical use with scarce healthy controls (HCs) and no diagnostic labels, and when an established tracer is introduced at a new site or resource-limited region where normative references have yet to be accumulated. Existing anomaly detection methods require large training cohorts or operate only on two-dimensional slices, failing to address the joint demands of extreme data scarcity and volumetric reasoning. We propose **TracerAD**, a training-free few-shot framework for three-dimensional brain PET anomaly detection that requires only a handful of HC scans and no model training. TracerAD extracts frozen vision foundation model features from registered brain volumes and detects pathological deviations at three spatial scales: Spatially-Anchored Focal Detection (SAFD), Region-Constrained Semantic Matching (RCSM), and Topology-Aware Metabolic Reasoning (TAMR). This multi-scale approach produces a clinically interpretable 3D anomaly heatmap. Evaluated on ADNI amyloid PET under 1- to 10-shot settings, TracerAD achieves 84.13% AUC at 10-shot, with heatmaps aligning to known amyloid deposition patterns—offering radiologists an immediately deployable tool for cold-start scenarios. Our code will be released at GitHub.

Keywords: PET anomaly detection · training-free · few-shot learning · vision foundation models · neurodegeneration

1 Introduction

The clinical translation of novel positron emission tomography (PET) radiotracers faces a universal bottleneck: the cold-start diagnostic dilemma. Due to ethical and dosimetric constraints, initial deployment cohorts are small, with healthy controls (HCs) being particularly scarce compared to patients. This dilemma

† These authors contributed equally to this work.

also affects established tracers introduced in resource-limited regions where normative databases are unavailable [1, 2]. Lacking statistical references, nuclear medicine physicians must interpret scans by mentally comparing them against a handful of normal references. This reliance on subjective assessment leads to high inter-reader variability and reduced diagnostic confidence, especially when differentiating physiological variation from early pathological deposition [3].

To bridge this gap, automated anomaly detection (AD) systems are urgently needed. However, most AD methods operate on 2D images [4–7], discarding the 3D context essential for characterizing neurodegenerative disorders that originate in deep, small structures such as the substantia nigra and hippocampus [3]. While recent metabolism-aware approaches explore 3D PET analysis [8,9], they require substantial training data and fail under cold-start settings. What is needed is a system that operates under few-shot constraints, reasons in native 3D anatomical space, and produces clinically interpretable anomaly maps.

Recent advances in self-supervised vision foundation models (VFMs) [10, 11] offer a promising direction. These models acquire rich representations from vast unlabeled datasets and have demonstrated strong transferability to medical imaging without task-specific fine-tuning [12–14]. We explore whether this adaptability extends to PET tracer interpretation under cold-start conditions, where only a handful of healthy scans are available.

We propose **TracerAD**, the first training-free few-shot anomaly detection framework that leverages frozen VFM features to detect anomalies directly in 3D anatomical space. Guided by a standard brain atlas, **TracerAD** detects pathological deviations across three spatial scales. **Spatially-Anchored Focal Detection** (SAFD) identifies focal variations via coordinate-level matching of Montreal Neurological Institute (MNI)-registered volumes. **Region-Constrained Semantic Matching** (RCSM) complements this by evaluating dense feature patterns within anatomical regions to account for normal variability. Finally, **Topology-Aware Metabolic Reasoning** (TAMR) models the brain as a sparse dynamic graph to capture circuit-level disruptions in inter-regional connectivity. Using as few as a single reference scan, **TracerAD** generates precise 3D anomaly heatmaps.

Our contributions are: (1) Formalizing the PET cold-start problem as a few-shot anomaly detection task and presenting **TracerAD**, the first training-free framework addressing this constraint in 3D. (2) Proposing a multi-branch architecture that synergizes local anatomical matching with global topological reasoning. (3) Validating the method on ADNI amyloid PET data, demonstrating competitive performance under 1- to 10-shot settings without model training.

2 Method

2.1 Overview

Given a three-dimensional PET volume $\mathbf{X} \in \mathbb{R}^{H \times W \times D}$ and a reference set of N healthy controls $\{\mathbf{X}_n^{\mathrm{HC}}\}_{n=1}^{N}$, TracerAD produces a voxel-wise anomaly map

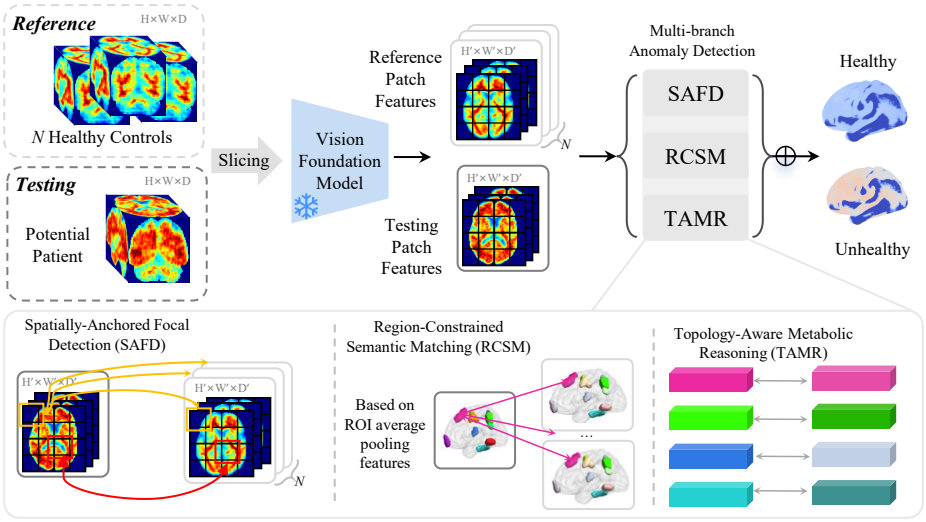

Fig. 1. Overview of proposed TracerAD.

$\mathbf{A} \in [0,1]^{H\times W\times D}$ without any model training or fine-tuning. As illustrated in Fig. 1, all volumes are first registered to MNI152 space and parcellated into N_{ROI} anatomical regions via the Automated Anatomical Labeling (AAL3) atlas [15]. A frozen VFM then extracts a dense feature map $\mathbf{F} \in \mathbb{R}^{H'\times W'\times D'\times C}$ from axial slices of each volume. Three complementary branches operate on these features at progressively broader spatial scales: SAFD compares features at identical coordinates across subjects to detect focal micro-lesions; RCSM performs dense voxel-wise matching within each atlas-defined region to capture intra-regional metabolic deviations; and TAMR constructs a sparse dynamic graph over region-level embeddings to identify circuit-level disruptions in inter-regional co-variation. Each branch yields an ROI-level anomaly score, and the three scores are fused via equal weighting into a single three-dimensional heatmap that is overlaid on the MNI152 template for clinical interpretation.

2.2 Spatially-Anchored Focal Detection (SAFD)

SAFD detects focal anomalies at the finest voxel-level granularity, addressing the clinical need to identify early pathological changes that may involve only a small portion of a region or span the junction between two adjacent regions. Standard global patch matching [4], as employed in industrial anomaly detection [5], compares a query patch against all reference patches regardless of spatial location, which risks matching anatomically unrelated regions and producing false positives. To address this, we exploit the strict coordinate-level alignment that MNI152 registration provides across subjects [6].

Specifically, for each spatial position (x, y, z), a micro memory bank $\mathcal{M}_{xyz}$ is constructed by collecting the features of N HCs at that location. The query

voxel feature is then compared exclusively against these co-located HC features. The focal anomaly score is defined as

$$S_{\mathrm{SAFD}}^{(x,y,z)} = 1 - \max_{n \in \{1,\ldots,N\}} \cos\Big(\mathbf{f}_{\mathrm{test}}^{(x,y,z)},\, \mathbf{f}_{n}^{(x,y,z)}\Big), \tag{1}$$

where $\cos(\cdot,\cdot)$ denotes cosine similarity. Voxel-level scores are aggregated within each ROI via mean pooling to obtain an ROI-level focal score.

2.3 Region-Constrained Semantic Matching (RCSM)

RCSM evaluates whether the overall metabolic pattern within a specific anatomical region conforms to the healthy reference. Different brain regions exhibit inherently different baseline metabolic rates [16]; unconstrained global matching would confuse normal physiological heterogeneity across regions with true pathological change [17]. Moreover, strict coordinate-level matching is sensitive to minor registration errors and normal anatomical variability across subjects, motivating a more flexible matching strategy within anatomical boundaries. To address these issues, we adopt an atlas-constrained matching strategy inspired by component-aware patch matching [5].

For ROI r, all voxel features from the query volume are collected as $\mathbf{Q}_r \in \mathbb{R}^{N_q \times C}$, and the corresponding features from N HCs are concatenated into $\mathbf{R}_r \in \mathbb{R}^{N_{\mathrm{ref}} \times C}$. Each query voxel is matched against all reference voxels within the same ROI, allowing flexible "sliding matching" within the anatomical boundary that tolerates registration imprecision. We adopt dense voxel-level matching rather than average pooling within each ROI, because average pooling discards intra-ROI heterogeneity and would dilute focal anomalies. The anomaly score of each query voxel v is computed as

$$a_v = 1 - \max_{j \in \mathcal{R}_r} \cos\big(\mathbf{f}_v^{q},\, \mathbf{f}_j^{\mathrm{ref}}\big). \tag{2}$$

The ROI-level score is obtained by averaging a_v over all voxels in r and is propagated to every voxel in that ROI via the atlas mask.

2.4 Topology-Aware Metabolic Reasoning (TAMR)

Neurodegenerative pathology can disrupt the metabolic co-variation between functionally connected brain regions, even when each region's absolute metabolic level remains within the normal range [3]. Detecting such network-level disruptions requires modeling inter-regional dependencies rather than evaluating each region in isolation. Inspired by the graph-based component interaction module of UniVAD [5], TAMR models the brain as a dynamic metabolic graph to capture these dependencies.

The voxel features within each ROI are average-pooled to produce a node embedding $\mathbf{h}_r \in \mathbb{R}^C$ for each of the N_{ROI} atlas-defined regions (166 in AAL3), yielding a node feature matrix $\mathbf{F}_{\mathrm{node}} \in \mathbb{R}^{N_{\mathrm{ROI}} \times C}$. A pairwise cosine similarity

matrix is computed as the initial adjacency. However, a fully connected graph forces every ROI to aggregate information from all other ROIs, which introduces cross-region noise and causes oversmoothing of anomaly signals. We therefore retain only the top-κ fraction of edges for each node, setting all other entries to zero:

$$\hat{A}_{ij} = \begin{cases} \dfrac{\cos(\mathbf{h}_i,\, \mathbf{h}_j)}{\sum_{l \in \mathcal{T}_i} \cos(\mathbf{h}_i,\, \mathbf{h}_l)}, & \text{if } j \in \mathcal{T}_i, \\ 0, & \text{otherwise}, \end{cases} \tag{3}$$

where $\mathcal{T}_i$ denotes the set of top-κ most similar nodes to node i and κ is set to 10% by default. A single-layer graph convolution aggregates metabolic context from the sparse neighbourhood: $\hat{\mathbf{F}} = \hat{\mathbf{A}} \cdot \mathbf{F}_{\text{node}}$, so that each node embedding encodes the topology of its core metabolic sub-network rather than the entire graph. The TAMR score for ROI r is computed by comparing the enhanced query embedding against the N enhanced HC embeddings:

$$S_{\text{TAMR}}(r) = 1 - \max_{n \in \{1,\ldots,N\}} \cos\left(\hat{\mathbf{f}}_q^r,\, \hat{\mathbf{f}}_n^r\right). \tag{4}$$

ROI scores are broadcast to voxel space via the atlas mask, as in RCSM.

2.5 Multi-Scale Fusion

The final anomaly score at each voxel is a weighted combination of the three branches:

$$S_{\text{final}} = w_s\, S_{\text{SAFD}} + w_r\, S_{\text{RCSM}} + w_t\, S_{\text{TAMR}}, \tag{5}$$

where w_s, w_r, and w_t are all set to $1/3$ by default. The resulting three-dimensional heatmap is overlaid on the MNI152 template to provide physicians with a spatially interpretable summary of detected anomalies.

3 Experiments

3.1 Dataset and Implementation Details

We evaluated TracerAD on an amyloid PET dataset from the Alzheimer's Disease Neuroimaging Initiative (ADNI) [18], comprising 330 subjects (160 amyloid-positive patients and 170 HCs) scanned with ^{18}F-florbetapir (AV-45) [19]. All volumes were spatially normalized to MNI152 space; RCSM and TAMR operate on anatomical regions defined by the AAL3 atlas. To simulate the cold-start scenario in which only a small number of normal references are available, we adopted an N-shot evaluation protocol ($N \in \{1, 2, 4, 8, 10\}$): for each trial, N HCs were randomly sampled as the reference set, and the remaining subjects were used for testing. This procedure was repeated with 5 random seeds, and we report mean$_{\pm\text{std}}$ over all trials. Three metrics were computed: area under the ROC curve (AUC), weighted F1-score (wF1), and sensitivity (Sen).

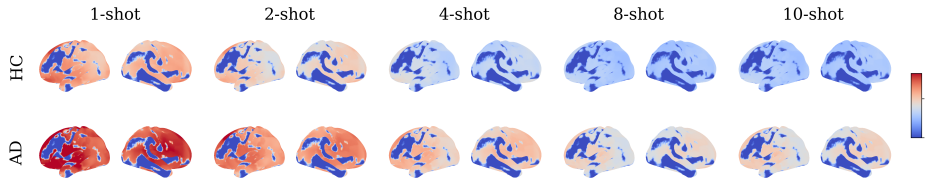

Fig. 2. Visualization of 3D anomaly heatmaps produced by TracerAD under 1- to 10-shot settings, averaged across all test subjects. The top row shows the average heatmap for HCs, and the bottom row for AD patients. Color intensity encodes the mean anomaly score, with warmer hues indicating higher values.

3.2 Comparison with Supervised Methods

We compared TracerAD against the supervised few-shot classification method of Liu *et al.* [14] in both CLS-token and Patch-token KNN variants. A key protocol distinction separates the two approaches: Liu *et al.* require both HC and patient scans in the few-shot reference set ($2N$ labeled samples for an N-shot setting), whereas TracerAD uses only N unlabeled HCs. In a cold-start scenario, only a small number of normal scans are available and patient labels do not yet exist, making the unsupervised protocol the only feasible option. Table 1 reports the results under 1- to 10-shot settings with the MedDINOv3 backbone. Despite operating under this stricter unsupervised protocol, TracerAD outperformed the supervised baselines on the majority of settings. At 10-shot, TracerAD achieved 84.13% AUC, 78.83% wF1, and 81.60% Sen, compared with 82.73% AUC, 76.59% wF1, and 77.47% Sen for the best baseline (Liu *et al.* CLS). TracerAD obtained the highest wF1 and Sen across all shot counts, and the highest AUC at all settings except 2-shot. The overall best sensitivity of 86.60% was reached by TracerAD at 2-shot, which is clinically relevant because in amyloid PET screening, missing a positive case is costlier than a false alarm.

Table 1. Performance comparison using MedDINOv3 backbone under 1-, 2-, 4-, 8-, and 10-shot settings with 5 random repetitions. Results are reported as mean$_{\pm\text{std}}$ (%). Underline indicates best per shot; **bold underline** indicates overall best. AUC: area under the ROC curve; wF1: weighted F1-score; Sen: sensitivity.

	Method	Metric	1-shot	2-shot	4-shot	8-shot	10-shot
Supervised	Liu et al. [14]$_{\text{CLS}}$	AUC	61.60$_{\pm 13.16}$	75.47$_{\pm 11.88}$	78.43$_{\pm 9.50}$	81.49$_{\pm 3.15}$	82.73$_{\pm 5.12}$
		wF1	59.87$_{\pm 14.47}$	63.67$_{\pm 10.43}$	66.14$_{\pm 10.88}$	74.91$_{\pm 4.07}$	76.59$_{\pm 6.41}$
		Sen	60.96$_{\pm 24.06}$	70.11$_{\pm 23.37}$	67.56$_{\pm 21.79}$	75.74$_{\pm 8.93}$	77.47$_{\pm 8.75}$
	Liu et al. [14]$_{\text{Patch}}$	AUC	53.71$_{\pm 6.98}$	62.92$_{\pm 7.37}$	64.19$_{\pm 8.10}$	67.40$_{\pm 6.29}$	69.06$_{\pm 5.42}$
		wF1	51.44$_{\pm 7.88}$	54.81$_{\pm 7.10}$	57.23$_{\pm 7.25}$	60.46$_{\pm 6.41}$	61.97$_{\pm 4.97}$
		Sen	51.69$_{\pm 22.54}$	49.00$_{\pm 21.30}$	54.40$_{\pm 17.18}$	53.54$_{\pm 12.68}$	58.59$_{\pm 13.24}$
Unsup.	TracerAD (Ours)	AUC	64.08$_{\pm 19.02}$	70.59$_{\pm 16.04}$	78.82$_{\pm 5.17}$	82.35$_{\pm 4.35}$	**84.13**$_{\pm 1.02}$
		wF1	61.92$_{\pm 13.88}$	66.93$_{\pm 14.21}$	73.84$_{\pm 4.44}$	76.63$_{\pm 3.42}$	**78.83**$_{\pm 1.90}$
		Sen	62.20$_{\pm 29.90}$	**86.60**$_{\pm 4.50}$	79.00$_{\pm 12.55}$	83.60$_{\pm 4.63}$	81.60$_{\pm 5.04}$

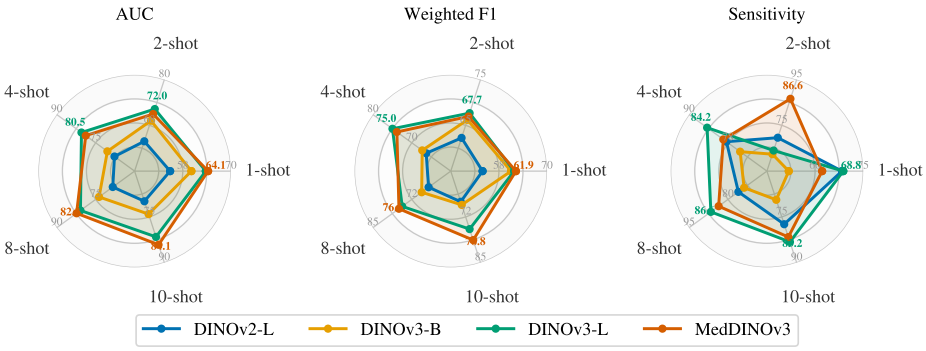

Fig. 3. Radar charts comparing four frozen backbones across 1- to 10-shot settings on AUC, Weighted F1, and Sensitivity.

TracerAD also exhibited lower variance across random seeds (e.g., AUC std of 1.02 vs. 5.12 at 10-shot), indicating that its performance is less sensitive to which specific HCs compose the reference set. Furthermore, TracerAD produces voxel-wise 3D anomaly heatmaps that localize pathological regions, whereas the baselines output only a binary subject-level label, enabling physicians to verify the anatomical plausibility of each detection.

3.3 Backbone Analysis

To investigate the impact of the visual backbone, we evaluated TracerAD with four frozen backbones: DINOv2 ViT-L/14 [10], DINOv3 ViT-B/16, DINOv3 ViT-L/16 [11], and MedDINOv3 [12]. Fig. 3 summarizes the 10-shot AUC for each backbone. MedDINOv3 achieved the highest AUC of 84.13%, followed by DINOv3-L at 81.55%, DINOv3-B at 74.09%, and DINOv2-L at 69.88%. MedDINOv3 outperformed all general-purpose backbones, indicating that medical-domain pretraining produces features better aligned with PET metabolic patterns even without fine-tuning [12]. Among the general-purpose models, DINOv3-L outperformed DINOv3-B by 7.46 AUC points, confirming that larger backbones extract richer representations for this task. DINOv2-L was the weakest despite being a strong general-purpose model [10], suggesting that architecture generation (v3 vs. v2) matters more than model size alone for medical imaging. Even with this weakest backbone, TracerAD still achieved 69.88% AUC, demonstrating robustness to backbone choice.

3.4 Ablation Study

Table 2 reports the additive ablation results using the MedDINOv3 backbone under the same N-shot protocol. SAFD alone achieved 80.07% AUC at 10-shot, confirming that coordinate-level matching already captures focal anomalies. Adding RCSM yielded the largest gain, raising AUC to 83.32% (+3.25

Table 2. Ablation study of the key components using MedDINOv3 backbone: (a) SAFD, (b) RCSM, (c) TAMR with full connection and (d) Top-10%.

Configuration			AUC				
SAFD	RCSM	TAMR	1-shot	2-shot	4-shot	8-shot	10-shot
✓	×	×	**65.01$_{\pm 16.59}$**	69.12$_{\pm 17.83}$	76.33$_{\pm 6.16}$	81.50$_{\pm 4.59}$	80.07$_{\pm 2.47}$
✓	✓	×	63.78$_{\pm 19.44}$	70.56$_{\pm 16.74}$	78.44$_{\pm 5.28}$	81.70$_{\pm 4.51}$	83.32$_{\pm 1.26}$
✓	✓	Full	63.87$_{\pm 19.04}$	70.23$_{\pm 16.34}$	78.49$_{\pm 5.40}$	**82.42$_{\pm 4.78}$**	84.05$_{\pm 1.33}$
✓	✓	Top-10%	64.08$_{\pm 19.02}$	**70.59$_{\pm 16.04}$**	**78.82$_{\pm 5.17}$**	82.35$_{\pm 4.35}$	**84.13$_{\pm 1.02}$**

points), because atlas-constrained dense matching captures regional metabolic deviations that fixed-coordinate comparison misses, such as slight spatial shifts within an ROI. Incorporating TAMR with a full inter-regional graph further improved AUC to 84.05% (+0.73 points), indicating that circuit-level metabolic disruptions provide complementary signal beyond local features. Replacing the full graph with a top-10% sparse graph yielded 84.13% AUC with lower standard deviation (1.02 vs. 1.33), suggesting that sparse connectivity filters noisy edges and focuses on the strongest inter-regional correlations. These results confirm that the three branches address distinct aspects of pathological change and that their combination is complementary.

4 Conclusion

We presented TracerAD, the first training-free few-shot framework that addresses the cold-start diagnostic dilemma in PET tracer deployment. By combining frozen VFM features with atlas-guided multi-scale scoring across three complementary branches, TracerAD produces 3D anomaly heatmaps from as few as a single healthy reference scan without any model training. Experiments on an amyloid PET dataset from ADNI demonstrated that TracerAD achieved 84.13% AUC at 10-shot with only HC references, matching or exceeding a supervised baseline that requires patient labels. These results suggest that TracerAD can serve as a practical computer-aided interpretation tool to accelerate the clinical translation of PET tracers across diverse deployment settings. The current study is limited to a single tracer; future work will extend validation to emerging tracers and multi-site cohorts.

Acknowledgments. This work was partially supported by National Natural Science Foundation of China (62131015, 82394432), Shanghai Municipal Central Guided Local Science and Technology Development Fund (YDZX20233100001001), and AI4S Initiative and HPC Platform of ShanghaiTech University.

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

References

1. Brink, A., Cerci, J.J., Giammarile, F., Yurystovskyi, D., Korde, A., Estrada-Lobato, E., Mikhail-Lette, M., Pellet, O., Macapinlac, H.A., Paez, D.: Global challenges in the provision of theranostic services the international atomic energy agency (iaea) perspective. World Journal of Nuclear Medicine **24**(03), 238–243 (2025)
2. Arshad, F., Alladi, S.: Disease-modifying drugs for alzheimer's disease: Bending the curve. Annals of Indian Academy of Neurology **28**(5), 640–645 (2025)
3. He, Y., Qin, L., Ye, Q., Xie, F.: Illuminating synucleinopathies: advances in α-synuclein pet tracer development for in vivo neuroimaging. Genomic Psychiatry **1**(4), 34–39 (2025)
4. Roth, K., Pemula, L., Zepeda, J., Schölkopf, B., Brox, T., Gehler, P.: Towards total recall in industrial anomaly detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 14318–14328 (2022)
5. Gu, Z., Zhu, B., Zhu, G., Chen, Y., Tang, M., Wang, J.: Univad: A training-free unified model for few-shot visual anomaly detection. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 15194–15203 (2025)
6. Zhang, X., Li, G., Henderson, P., Aragon-Camarasa, G., Long, Z.: Is training necessary for anomaly detection? arXiv preprint arXiv:2601.22763 (2026)
7. Cai, Y., Zhang, W., Chen, H., Cheng, K.T.: Medianomaly: A comparative study of anomaly detection in medical images. Medical Image Analysis **102**, 103500 (2025)
8. Zhao, L., Dong, H., Wu, P., Lu, J., Lu, L., Zhou, J., Liu, T., Zhang, L., Zhang, L., Tang, Y., et al.: Metavit: Metabolism-aware vision transformer for differential diagnosis of parkinsonism with 18 f-fdg pet. In: International Conference on Information Processing in Medical Imaging. pp. 132–144. Springer (2023)
9. Huang, H., Shen, Z., Wang, J., Wang, X., Lu, J., Lin, H., Ge, J., Zuo, C., Wang, Q.: Metaad: metabolism-aware anomaly detection for parkinson's disease in 3d 18 f-fdg pet. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 291–301. Springer (2024)
10. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193 (2023)
11. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., et al.: Dinov3. arXiv preprint arXiv:2508.10104 (2025)
12. Li, Y., Wu, Y., Lai, Y., Hu, M., Yang, X.: Meddinov3: How to adapt vision foundation models for medical image segmentation? arXiv preprint arXiv:2509.02379 (2025)
13. Song, X., Xu, X., Yan, P.: Dino-reg: General purpose image encoder for training-free multi-modal deformable medical image registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 608–617. Springer (2024)
14. Liu, C., Chen, Y., Shi, H., Lu, J., Jian, B., Pan, J., Cai, L., Wang, J., Yu, J., Gao, Z., et al.: Does dinov3 set a new medical vision standard? benchmarking 2d and 3d classification, segmentation, and registration. arXiv preprint arXiv:2509.06467 (2025)
15. Rolls, E.T., Huang, C.C., Lin, C.P., Feng, J., Joliot, M.: Automated anatomical labelling atlas 3. Neuroimage **206**, 116189 (2020)

16. Mazziotta, J.C., Phelps, M.E., Miller, J., Kuhl, D.E.: Tomographic mapping of human cerebral metabolism: normal unstimulated state. Neurology **31**(5), 503–503 (1981)
17. Vallini, G., Baron, G., Silvestri, E., Volpi, T., Vlassenko, A., Goyal, M., Chiuso, A., Cecchin, D., Corbetta, M., Bertoldo, A.: Brain metabolic-functional (de) coupling from health to glioma dysfunction. Communications Biology (2025)
18. Petersen, R.C., Aisen, P.S., Beckett, L.A., Donohue, M.C., Gamst, A.C., Harvey, D.J., Jack, C.R., Jagust, W.J., Shaw, L.M., Toga, A.W., et al.: Alzheimer's Disease Neuroimaging Initiative (ADNI): clinical characterization. Neurology **74**(3), 201–209 (2010)
19. Clark, C.M., Schneider, J.A., Bedell, B.J., Beach, T.G., Bilker, W.B., Mintun, M.A., Pontecorvo, M.J., Hefti, F., Carpenter, A.P., Flitter, M.L., et al.: Use of florbetapir-PET for imaging β-amyloid pathology. JAMA **305**(3), 275–283 (2011)