Tracer-Adaptive Expert Learning for Generalizable PET Lesion Segmentation

Yiran Shu[1,2], Jingyu Liu[2], Mingyang Yu[1], Xinlu Wang[3], Guiyu Li[4], Weijun Wei[5], Lianghua Li[5], Yanbo Chen[2], Yaozong Gao[2✉], and Dinggang Shen[1,2,6✉]

[1] School of Biomedical Engineering & State Key Laboratory of Advanced Medical Materials and Devices, ShanghaiTech University, Shanghai, China
Dinggang.Shen@gmail.com
[2] Shanghai United Imaging Intelligence Co., Ltd., Shanghai, China
yaozong.gao@uii-ai.com
[3] The First Affiliated Hospital of Guangzhou Medical University, Guangzhou, China
[4] Xijing Hospital, Xi'an, China
[5] Renji Hospital, Shanghai Jiao Tong University School of Medicine, Shanghai, China
[6] Shanghai Clinical Research and Trial Center, Shanghai, China

Abstract. PET/CT imaging is evolving from a single-tracer modality to a multi-tracer paradigm, enabling precise characterization of diverse disease patterns. However, a core diagnostic challenge remains, i.e., distinguishing pathological lesions from physiological uptake, as both appear as high-intensity "hot spots." This distinction is highly dependent on the anatomical location and tracer type, making automated segmentation across tracers particularly difficult. Standard Mixture-of-Experts (MoE) architectures theoretically support multi-task learning but typically rely on feature-driven gating. This leads to uncontrollable routing, because of failing to leverage the anatomical and tracer-specific priors essential for radiological interpretation. In this paper, we propose ProMoE, a Prototype-Driven Mixture-of-Experts framework for context-informed PET lesion segmentation across tracers. Unlike conventional MoE, ProMoE introduces a set of learnable expert prototypes that serve as anchors within a semantic embedding space to represent particular physiological contexts. During inference, each input is first embedded into this space according to anatomical and tracer-specific priors. Then, expert activation is performed driven by the similarity between the embedding and expert prototypes. This mechanism supports physiologically informed expert specialization and tracer-adaptive routing. Validated on a large-scale multi-center dataset of **1,366** PET/CT scans across **3** tracers (^{18}F-FDG, ^{18}F/^{68}Ga-PSMA, and ^{18}F-FAPI), ProMoE significantly outperforms the best state-of-the-art method by 5.97% in average DSC. Furthermore, it yields an 8.82% DSC improvement in prior-guided zero-shot transfer on **30** cases of an unseen tracer (CD70), underscoring its potential as a scalable, biologically grounded framework for prior-guided adaptation to new PET tracers. The code is publicly available at https://github.com/CMYXing/ProMoE.

Keywords: PET/CT · Lesion Segmentation · Mixture-of-Experts · Prior-guided Transfer.

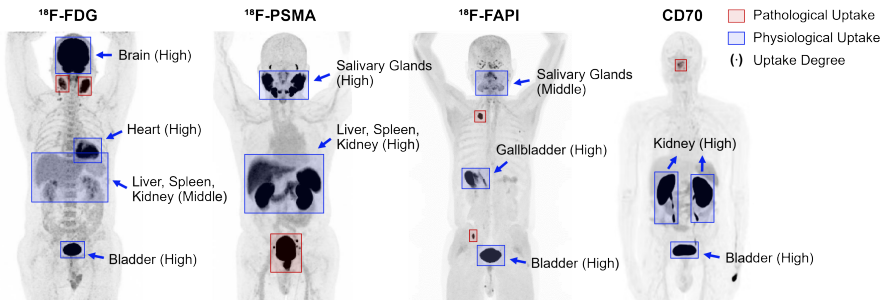

Fig. 1: Pathological lesions and physiological uptake both manifest as high-intensity "hot spots" in PET imaging. Radiologists distinguish them by integrating anatomical context and tracer-specific information.

1 Introduction

Positron Emission Tomography/Computed Tomography (PET/CT) is a cornerstone of precision oncology [11]. While historically dominated by FDG, the field has shifted toward a multi-tracer paradigm, employing targeted radiotracers like PSMA and FAPI to visualize disease processes previously invisible [12]. However, a persistent challenge in PET diagnostics is the ambiguity of signal intensity: Both pathological lesions and physiological uptake appear as high-intensity "hot spots." Their interpretation relies entirely on context. For instance, high FDG uptake is physiological in the brain but pathological in the neck; conversely, high PSMA uptake is normal in the neck (salivary glands) but would indicate pathology if present in the brain (Fig. 1). Consequently, robust segmentation requires a model to look beyond raw intensity to integrate anatomical and tracer-specific priors for interpreting "hot spots" within their correct biological context.

Mixture-of-Experts (MoE) architectures offer a promising solution by allowing experts to specialize in distinct data patterns [13,1,16]. However, standard MoEs typically employ feature-driven gating networks that often lead to context-agnostic routing in multi-tracer datasets. For example, a renal tumor in FDG and a normal kidney in PSMA may appear nearly identical. Appearance-based routers inadvertently assign these conflicting signals to the same expert, forcing it to simultaneously learn to segment and ignore the same visual pattern. This "conflicting gradient" problem results in performance degradation.

We argue that effective multi-tracer learning requires expert selection guided by radiological logic, i.e., determining where the signal is (anatomy) and what the normal metabolic behavior should be (physiological expectation). To this end, we introduce **ProMoE**, a novel **Pro**totype-Driven **M**ixture-**o**f-**E**xperts framework. Departing from black-box gating, ProMoE injects explicit clinical priors into routing via a set of learnable expert prototypes. These prototypes serve as anchors in a semantic embedding space, representing distinct physiological contexts (e.g., "high uptake in the kidney"). During inference, image patches are first embedded into this semantic space based on their anatomical and tracer-

specific priors. Then, expert activation is determined by the similarity between these patch-level embeddings and expert prototypes. This mechanism promotes consistent expert activation for similar physiological patterns across diverse tracers, enabling tracer-adaptive expert learning. Based on this, we devise a dual-stage feature modulation module to sequentially refine PET features via FiLM-based [10] anatomy-aware and prototype-driven physiology-aware modulation. This synergy enables ProMoE to suppress confounding physiological signals in a context-informed manner, thereby supporting generalization in multi-tracer scenarios.

Our contributions are summarized as follows: (i) **ProMoE**, a novel prototype-driven MoE framework, that embeds explicit radiological priors into the routing logic for generalizable PET segmentation across tracers; (ii) a **dual-stage feature modulation** module that sequentially integrates anatomical and physiological cues to refine PET features for highlighting abnormal regions; (iii) an **extensive validation** on 1,366 PET/CT scans across 3 tracers, demonstrating superior performance over state-of-the-art (SOTA) models; and (iv) **prior-guided zero-shot transfer** on 30 cases from an unseen tracer, highlighting ProMoE's scalability and transfer potential for adapting to new PET tracers with curated physiological priors.

2 Methodology

We propose a novel Prototype-Driven MoE framework, called ProMoE, for PET lesion segmentation across tracers. As illustrated in Fig. 2, the framework employs a dual-stream design: a frozen CT backbone provides anatomical grounding, while a PET backbone performs lesion segmentation. We project inputs into a semantic embedding space to govern expert routing (Sec. 2.1). Feature modulation is then driven by specific anatomical and physiological contexts (Sec. 2.2). Our framework is optimized via a joint objective (Sec. 2.3).

2.1 Physiology-Aware Expert Prototype Routing

Physiological Expectation Embedding. To bridge static clinical knowledge with dynamic image context, we propose the Physiological Expectation Embedding (PE^2). First, we establish an offline multi-tracer physiological atlas covering $D = 34$ anatomical structures. The atlas was drafted using a Large Language Model (LLM), reviewed by two senior nuclear medicine experts, and then fixed as a deterministic lookup table of normative uptake levels, $l_d^{(t)} \in \{1, \ldots, 6\}$, for each tracer t and structure d.

During inference, for PET/CT image patches $\{x_{\text{ct}}, x_{\text{pet}}\}$ with tracer t, we retrieve the corresponding physiological priors from the atlas. To identify structures within the current field-of-view, we employ a frozen CT backbone (pre-trained on multi-organ segmentation) to extract anatomical occupancy maps $\mathbf{m} \in \{0, 1\}^D$, where $m_d = 1$ indicates the presence of the d-th structure in patch x_{ct}. The embedding $\mathbf{v} \in \mathbb{R}^D$ is defined element-wise as $v_d = l_d^{(t)} \cdot m_d$, explicitly encoding both the anatomical structure and expected physiological pattern.

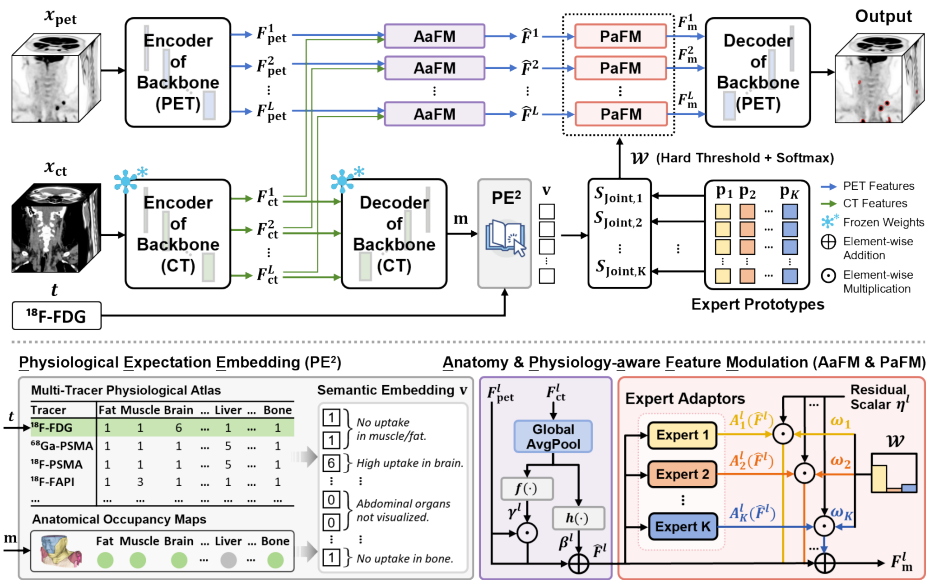

Fig. 2: Overview of the ProMoE framework. The Physiological Expectation Embedding (PE^2) module maps PET/CT inputs into a semantic embedding space grounded in anatomical and tracer-specific priors. Routing weights $\mathcal{W}$ are computed via similarity matching between these embeddings and a set of learnable expert prototypes. Finally, PET features are refined through sequential Anatomy-aware (AaFM) and Physiology-aware (PaFM) modulations to achieve lesion segmentation across tracers.

Prototype-Driven Expert Activation. We govern the experts via a set of K expert prototypes $\{\mathbf{p}_k\}_{k=1}^{K}$. Each prototype $\mathbf{p}_k \in \mathbb{R}^D$ resides in the same semantic space as $\mathbf{v}$ and acts as a learnable anchor for a specific physiological context. The relevance of the k-th expert is determined by the similarity between the input embedding $\mathbf{v}$ and prototype $\mathbf{p}_k$. We compute a joint similarity score $S_{\text{Joint},k}$ comprising two components:

1. Anatomical Alignment (S_{Ana}): The Dice coefficient between binarized non-zero supports $\mathbf{v}_b = \mathbb{1}(\mathbf{v})$ and $\mathbf{p}_{k,b} = \mathbb{1}(\mathbf{p}_k)$, quantifying anatomical overlap:

$$S_{\text{Ana},k} = \text{Dice}(\mathbf{v}_b, \mathbf{p}_{k,b}) = 2(\mathbf{v}_b \cdot \mathbf{p}_{k,b})/(\|\mathbf{v}_b\|_1 + \|\mathbf{p}_{k,b}\|_1), \tag{1}$$

2. Physiological Consistency (S_{Phy}): The Cosine similarity within the mutual support $\Omega_k = \{i \mid v_i \neq 0 \wedge p_{k,i} \neq 0\}$, with $\mathbf{v}_{\Omega_k}$ and $\mathbf{p}_{k,\Omega_k}$ denoting signals restricted to Ω_k. This aligns local expected physiological patterns:

$$S_{\text{Phy},k} = \text{CosSim}(\mathbf{v}_{\Omega_k}, \mathbf{p}_{k,\Omega_k}) = (\mathbf{v}_{\Omega_k} \cdot \mathbf{p}_{k,\Omega_k})/(\|\mathbf{v}_{\Omega_k}\|_2 \|\mathbf{p}_{k,\Omega_k}\|_2). \tag{2}$$

Final routing weights $\mathcal{W}$ are derived by applying Softmax to the joint scores $S_{\text{Joint},k} = S_{\text{Ana},k} \cdot (S_{\text{Phy},k})^2$. To enhance computational efficiency while suppress-

ing noise, we enforce sparsity via a hard threshold $\theta = 1/K$: if $S_{\text{Joint},k} < \theta$, the weight w_k is set to $-\infty$ (pre-softmax), effectively pruning irrelevant experts.

Expert Prototypes Initialization. To ensure stable convergence, we initialize expert prototypes via a data-driven approach. We construct a feature bank by extracting semantic embeddings from 5,000 randomly-sampled training patches. K-Medoids clustering [9] is applied using our joint similarity metric. The resulting cluster medoids are used to initialize the expert prototypes.

2.2 Dual-Stage Feature Modulation

We employ a dual-stage modulation process at each resolution level l of the PET backbone, integrating both anatomical constraints and physiological guidance.

Anatomy-aware Feature Modulation (AaFM). We leverage the anatomical context of CT to condition PET features. Using Feature-wise Linear Modulation (FiLM) [10], the CT feature map F^l_{ct} generates affine parameters γ^l and β^l to scale and shift the PET features F^l_{pet}. This spatially aligns the PET signal with anatomical structures, yielding intermediate features $\hat{F}^l$.

Physiology-aware Feature Modulation (PaFM). Next, the prototype-driven MoE module refines $\hat{F}^l$ by deploying K expert adaptors $\{A^l_k\}_{k=1}^K$ (lightweight $1 \times 1 \times 1$ convolutions). These adaptors operate in one-to-one correspondence with the K expert prototypes. By aggregating the expert responses using the routing weights $\mathcal{W}$ derived in Sec. 2.1, this module supports context-informed modulation of physiological uptake patterns while preserving lesion-discriminative responses. Using a learnable residual scalar η^l to control the expert ensemble's contribution, the modulated feature F^l_{m} is computed as:

$$F^l_{\text{m}} = \hat{F}^l + \eta^l \sum\nolimits_k w_k \cdot A^l_k(\hat{F}^l). \tag{3}$$

2.3 Optimization Objectives

The framework is trained end-to-end using a composite loss function $\mathcal{L}$. The primary segmentation objective $\mathcal{L}_{\text{Seg}}$ combines Dice loss and Focal loss. To regulate the prototype-driven mechanism, we introduce two auxiliary terms:

$$\mathcal{L} = \mathcal{L}_{\text{Seg}} + \lambda_s \mathcal{L}_{\text{Sparse}} + \lambda_p \mathcal{L}_{\text{Proto}}. \tag{4}$$

Here, $\mathcal{L}_{\text{Sparse}}$ minimizes the entropy of the routing weights w_k to encourage crisp expert specialization. The prototype alignment loss, defined as $\mathcal{L}_{\text{Proto}} = \sum_k w_k(1 - S_{\text{Joint},k})$, pulls the activated expert prototypes closer to the embeddings of the current samples. This ensures the learnable prototypes evolve to represent the true physiological expectation clusters present in the training set.

Table 1: Quantitative comparison of lesion segmentation performance across different tracers. The **best** and second-best performance are highlighted among multi-tracer models. Tracer marked with * represents unseen domain. Statistical significance compared to **Ours** is indicated by † ($p < 0.1$) and ‡ ($p < 0.05$) based on Wilcoxon signed-rank tests.

Methods	Params	DSC (%)					NSD (%)				
		FDG	PSMA	FAPI	**Avg.**	CD70*	FDG	PSMA	FAPI	**Avg.**	CD70*
One-for-All Paradigm											
nnU-Net [5]	384M	68.74	51.61‡	54.57†	58.31	27.85‡	77.99	58.77‡	64.68†	67.15	42.51‡
SegResNet [8]	5M	65.30‡	31.24‡	49.92‡	48.82	40.95‡	74.17‡	35.88‡	59.01‡	56.35	62.11
UNETR [3]	105M	48.09‡	31.79‡	42.14‡	40.67	38.22‡	56.69‡	36.63‡	52.94‡	48.75	54.24‡
SwinU [2]	63M	66.42‡	59.38‡	56.70	60.83	42.57‡	76.67	68.68‡	67.31	70.89	**62.87**
MSAM3D [14]	101M	40.52‡	31.01‡	38.09‡	36.54	33.95‡	52.17‡	43.11‡	48.39‡	47.89	55.88‡
BAMF [7]	528M	69.01	52.32‡	53.12†	58.15	40.68‡	**82.34**‡	59.93‡	66.35	69.54	61.60
Ours	6M	**69.04**	**68.31**	**63.05**	**66.80**	**51.39**	78.18	**77.76**	**73.17**	**76.37**	62.79
One-for-One Paradigm											
nnU-Net [5]	384M*3	69.63	71.75	61.12	67.50	–	79.77	81.52	72.69	77.99	–

3 Experiments

3.1 Dataset and Experimental Setup

We curated a dataset of 1,366 PET/CT scans involving three tracer categories: ^{18}F-FDG (846 cases), ^{18}F/^{68}Ga-PSMA (284/66 cases), and ^{18}F-FAPI (170 cases). Data were sourced from three medical centers (in-house) and the public AutoPET challenge [4], covering diverse malignancies (e.g., lung cancer, lymphoma, melanoma, prostate cancer). For prior-guided zero-shot transfer, an additional 30 cases of an unseen CD70-targeted tracer were reserved for evaluation, where the model was deployed using only a fixed physiological prior to assess its transferability to a new tracer. Two senior nuclear medicine radiologists manually annotated all in-house cases based on clinical reports and verified the public labels to ensure consistency. Performance was assessed using a tracer-stratified split (90% training/cross-validation, 10% testing) via Dice Similarity Coefficient (DSC) and Normalized Surface Distance (NSD).

Our framework uses a V-Net backbone [6]. The CT branch utilizes frozen weights pre-trained for 500 epochs on TotalSegmentator [15] for robust anatomical priors. The PET branch integrates the dual-stage modulation module (Sec. 2.2) after each downsampling block. All models were trained for 3,000 epochs on a single NVIDIA L40 GPU (24GB) using $96 \times 96 \times 96$ patches with a batch size of 16. Optimization used Adam, with initial learning rate 1×10^{-3}, weight decay 1×10^{-4}, and a decay factor of 0.1 every 1,000 epochs.

3.2 Comparison study

We compared our framework against diverse SOTA paradigms: nnU-Net [5] (3D fullres) and SegResNet [8] (CNN-based); UNETR [3] and SwinUNETR

Table 2: Ablation study of proposed components. We evaluate modulation strategies in AaFM (Cat (Concatenation) vs. FiLM) and MoE routing schemes in PaFM (MLP, T-Part. (Tracer Partition), A-Part. (Anatomy Partition) vs. ProMoE), including the impact of expert count K. The **best** and second-best performance are highlighted. Tracer marked with * represents unseen domain.

Components			**DSC** (%)					**NSD** (%)				
AaFM	**PaFM**	(K)	FDG	PSMA	FAPI	**Avg.**	CD70*	FDG	PSMA	FAPI	**Avg.**	CD70*
Cat	w/o MoE	(1)	65.60	45.94	51.91	54.48	42.84	75.35	52.17	62.75	63.42	63.53
FiLM	w/o MoE	(1)	66.88	46.73	53.01	55.54	40.82	76.98	53.21	64.70	64.96	63.82
FiLM	MLP	(12)	65.13	39.55	55.78	53.49	41.71	74.00	44.68	63.81	60.83	53.87
	T-Part.	(3)	66.21	48.42	60.95	58.53	–	75.31	54.42	69.87	66.53	–
	A-Part.	(12)	64.54	51.12	57.77	57.81	43.85	72.99	57.56	66.10	65.55	56.56
FiLM	ProMoE	(6)	66.28	58.83	58.58	61.23	50.36	75.03	65.88	68.79	69.90	60.70
	ProMoE	(12)	69.04	**68.31**	**63.05**	**66.80**	51.39	**78.18**	**77.76**	**73.17**	**76.37**	62.79
	ProMoE	(24)	**69.13**	61.65	62.08	64.29	**51.75**	78.03	69.66	72.63	73.44	**65.95**
	ProMoE	(36)	61.92	55.87	59.33	59.04	38.34	70.10	63.68	68.92	67.57	50.50

(SwinU) [2] (Transformer-based); and a fine-tuned SAM-Med3D (MSAM3D) [14]. Additionally, we included the AutoPET III Challenge winner, BAMF Health [7].

As shown in Table 1, ProMoE achieves superior segmentation performance with only 6M parameters (≈1.5% of nnU-Net). In the multi-tracer setting, our model significantly outperforms the second-best SwinUNETR, particularly in the distinct PSMA domain (+8.93% DSC), indicating improved handling of tracer-specific physiological distributions. Notably, in the data-scarce FAPI domain, ProMoE even surpasses the single-tracer specialist, suggesting that our unified framework effectively leverages cross-tracer knowledge to mitigate data scarcity.

Furthermore, in the prior-guided CD70 transfer setting, ProMoE improves DSC by 8.82% over the best competing method on 30 held-out cases from the unseen tracer. Fig. 3B (right) suggests that our prototype-based routing leverages shared physiological patterns to activate the same expert. This mechanism supports prior-guided generalization to an unseen tracer without additional training data. Qualitative results in Fig. 3A further indicate that ProMoE reduces physiological false positives while maintaining high recall for pathological lesions.

3.3 Ablation Study

We investigated the impact of core components and hyper-parameters in Table 2. Replacing simple concatenation with our FiLM-based strategy improves average performance, indicating more effective refinement of PET features by CT structural priors. Regarding the routing mechanism within PaFM, we observe that the standard MLP-based MoE yields the lowest performance. This indicates that purely data-driven routing does not reliably separate pathological interference from underlying physiological variations without explicit prior guidance. Such instability is further evidenced in Fig. 3B (left), where its expert activation maps appear significantly more stochastic compared to the proposed prototype-driven

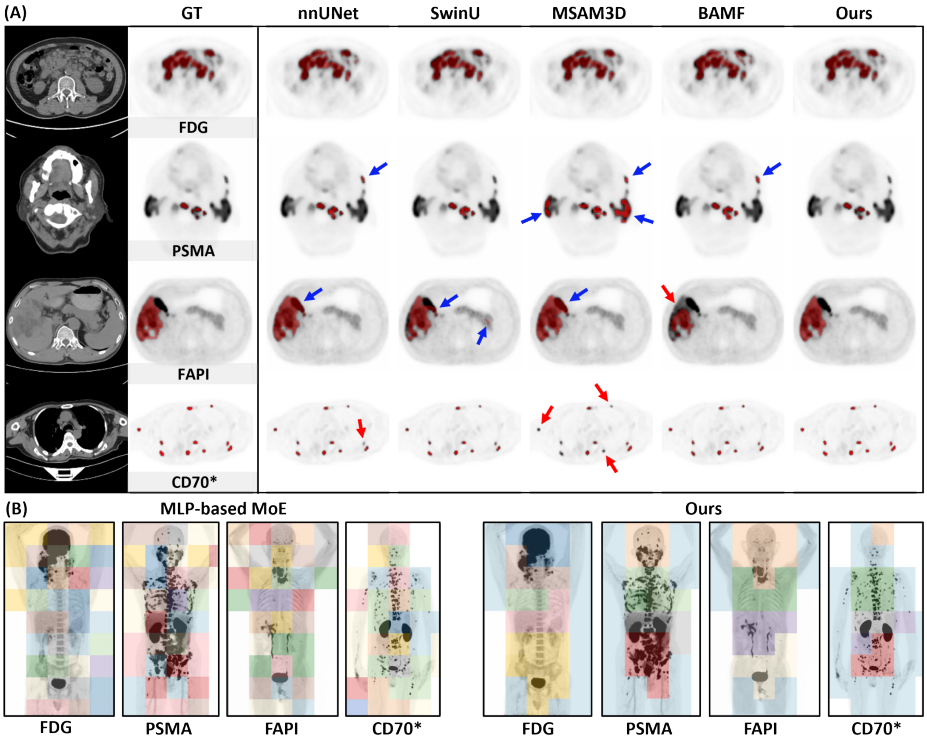

Fig. 3: (A) Representative qualitative results across tracers. Red and blue arrows indicate pathological FNs and physiological FPs of models, respectively; (B) Comparative visualization of the Top-1 expert activation maps ($K = 12$) for two routing schemes, using a unified color-to-expert mapping to highlight consistent expert assignments across various tracers.

approach. While hard-partitioning strategies based on tracer or anatomy improve stability, they exhibit limited generalizability and lack fine-grained physiological granularity. In contrast, ProMoE consistently achieves superior performance across all metrics. Finally, we analyzed expert capacity by varying the number of experts. Increasing experts from 6 to 12 significantly boosts capacity; however, increasing K to 36 degrades performance, suggesting that excessive sparsity hinders the shared representation learning required for effective transfer.

4 Conclusion

In this work, we have presented ProMoE, a prototype-driven Mixture-of-Experts framework to integrate clinical reasoning into multi-tracer PET lesion segmentation. By introducing learnable expert prototypes informed by anatomical and physiological priors, ProMoE supports interpretable and adaptive expert routing. Our dual-stage modulation module refines PET features by sequentially

incorporating FiLM-based anatomical constraints and prototype-driven physiological guidance. This significantly enhances lesion saliency while suppressing confounding physiological noise from healthy organs. Validated on a large-scale multi-tracer dataset, ProMoE achieves SOTA performance, and the prior-guided CD70 transfer experiment provides evidence of its potential to adapt to unseen tracers when curated physiological priors are available. The CD70 transfer evaluation was limited to 30 cases without center-disjoint validation, and tracers with substantially different physiological uptake patterns may require extending the physiological atlas or prototype space. Overall, ProMoE provides a scalable and interpretable framework for multi-tracer PET lesion segmentation, offering a practical path toward prior-guided generalization across tracers.

Acknowledgments. This work was supported in part by National Natural Science Foundation of China (grant numbers 82441023, U23A20295, 62131015, 82394432), the China Ministry of Science and Technology (S20240085, STI2030-Major Projects-2022ZD0209000, STI2030-Major Projects-2022ZD0213100), Shanghai Municipal Central Guided Local Science and Technology Development Fund (No. YDZX20233100001001), The Key R&D Program of Guangdong Province, China (grant number 2023B0303040001), HPC Platform of ShanghaiTech University, and China Medical Health Development Foundation. We also acknowledge The First Affiliated Hospital of Guangzhou Medical University, Xijing Hospital, and Renji Hospital for providing validation data.

Disclosure of Interests. All authors declare no conflicts of interest.

References

1. Chopra, S., Sanchez-Rodriguez, G., Mao, L., Feola, A.J., Li, J., Kira, Z.: MedMoE: Modality-Specialized Mixture of Experts for Medical Vision-Language Understanding. arXiv preprint arXiv:2506.08356 (2025)
2. Hatamizadeh, A., Nath, V., Tang, Y., Yang, D., Roth, H.R., Xu, D.: Swin UNETR: Swin Transformers for Semantic Segmentation of Brain Tumors in MRI Images. In: International MICCAI brainlesion workshop. pp. 272–284. Springer (2021)
3. Hatamizadeh, A., Tang, Y., Nath, V., Yang, D., Myronenko, A., Landman, B., Roth, H.R., Xu, D.: UNETR: Transformers for 3D Medical Image Segmentation. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 574–584 (2022)
4. Ingrisch, M., Dexl, J., Jeblick, K., Cyran, C., Gatidis, S., Kuestner, T.: Automated Lesion Segmentation in Whole-Body PET/CT-Multitracer Multicenter generalization (2024)
5. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. Nature methods **18**(2), 203–211 (2021)
6. Milletari, F., Navab, N., Ahmadi, S.A.: V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation (2016), `https://arxiv.org/abs/1606.04797`
7. Murugesan, G.K., McCrumb, D., Kumar, J., Soni, R., Grigorash, V., Moore, S., Oss, J.V., Health, B.: Multitracer Lesion Segmentation in FDG-18 and PSMA Whole-Body PET/CT Scans Using an Augmented Labeling Approach: AutoPET III Challenge. `https://github.com/bamf-health/autopet2024` (2024)

8. Myronenko, A.: 3D MRI Brain Tumor Segmentation Using Autoencoder Regularization. In: International MICCAI brainlesion workshop. pp. 311–320. Springer (2018)
9. Park, H.S., Jun, C.H.: A simple and fast algorithm for K-medoids clustering. Expert systems with applications **36**(2), 3336–3341 (2009)
10. Perez, E., Strub, F., De Vries, H., Dumoulin, V., Courville, A.: FiLM: Visual Reasoning with a General Conditioning Layer. In: Proceedings of the AAAI conference on artificial intelligence. vol. 32 (2018)
11. Townsend, D.W., Carney, J.P., Yap, J.T., Hall, N.C.: PET/CT today and tomorrow. Journal of Nuclear Medicine **45**(1 suppl), 4S–14S (2004)
12. Trotter, J., Pantel, A.R., Teo, B.K.K., Escorcia, F.E., Li, T., Pryma, D.A., Taunk, N.K.: Positron Emission Tomography (PET)/Computed Tomography (CT) Imaging in Radiation Therapy Treatment Planning: A Review of PET Imaging Tracers and Methods to Incorporate PET/CT. Advances in radiation oncology **8**(5), 101212 (2023)
13. Wang, G., Ye, J., Cheng, J., Li, T., Chen, Z., Cai, J., He, J., Zhuang, B.: SAM-Med3D-MoE: Towards a Non-Forgetting Segment Anything Model via Mixture of Experts for 3D Medical Image Segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 552–561. Springer (2024)
14. Wang, H., Guo, S., Ye, J., Deng, Z., Cheng, J., Li, T., Chen, J., Su, Y., Huang, Z., Shen, Y., et al.: SAM-Med3D: A Vision Foundation Model for General-Purpose Segmentation on Volumetric Medical Images. IEEE Transactions on Neural Networks and Learning Systems (2025)
15. Wasserthal, J., Breit, H.C., Meyer, M.T., Pradella, M., Hinck, D., Sauter, A.W., Heye, T., Boll, D.T., Cyriac, J., Yang, S., et al.: TotalSegmentator: Robust Segmentation of 104 Anatomic Structures in CT Images. Radiology: Artificial Intelligence **5**(5) (2023)
16. Zhong, A., Huang, H., Wang, J., Shen, Z., Song, H., Wu, J., Zhu, Y., Liu, Y., Zuo, C., Wang, Q.: Multi-tracer Uptake Correction for PET-MR via Aligned-Feature Guidance and Multi-scale Pixel-Adaptive Routing. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 413–422. Springer (2025)