



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.

PRA-PoE: Robust Multimodal Alzheimer’s Disease Classification under Arbitrary Modality Missingness

Guangqian Yang¹, Ye Du¹, Wenlong Hou¹, Qian Niu^{2✉}, Shujun Wang^{1,3✉}, and
for the Alzheimer’s Disease Neuroimaging Initiative

¹ Department of Biomedical Engineering, The Hong Kong Polytechnic University,
Hong Kong SAR, China

² Department of Technology Management for Innovation, The University of Tokyo,
Japan

³ Department of Data Science and Artificial Intelligence, The Hong Kong
Polytechnic University, Hong Kong SAR, China

qian.niu@weblab.t.u-tokyo.ac.jp, shu-jun.wang@polyu.edu.hk

✉ Corresponding authors.

Abstract. Missing modalities are prevalent in real-world Alzheimer’s disease (AD) assessment and pose a significant challenge to multimodal learning, particularly when the distribution of observed modality subsets differs between training and deployment. Such missingness pattern mismatch induces a conditional representation shift across modality subsets. Existing approaches that rely on implicit imputation or modality synthesis often fail to explicitly model modality availability and uncertainty, leading to overconfident dependence on synthesized features, reduced robustness, and miscalibrated uncertainty estimates. To address these limitations, we propose PRA-PoE, an incomplete multimodal learning framework that is equipped with Prototype-anchored Representation Alignment (PRA) and an Uncertainty-aware Product of Experts (UA-PoE) fusion mechanism. First, PRA uses learnable global prototypes and availability-conditioned tokens to encode modality availability, distinguish observed from missing modalities, re-synthesize features for missing modalities, and adaptively refine observed representations to align latent spaces across modality subsets, with the goal of reducing representation shift under varying missingness patterns. Second, UA-PoE models each modality as a Gaussian expert and performs closed-form Product of Experts fusion, where experts with higher uncertainty are automatically down-weighted via lower precision, improving uncertainty reliability. We evaluate PRA-PoE under a clinically realistic protocol by training with naturally missing data and testing on all non-empty modality combinations. PRA-PoE consistently outperforms the state-of-the-art across datasets, achieving a **5.4%** relative improvement in average accuracy on ADNI and a **10.9%** relative gain in average F1 on OASIS-3 over the strongest baseline across all non-empty modality subsets. Code is available at <https://github.com/SDH-Lab/PRA-PoE>.

Keywords: Missing Modality · Alzheimer’s Disease · Prototype Learning · Product of Experts

1 Introduction

Accurate classification of Alzheimer’s disease (AD) relies on multimodal evidence [8, 5, 20]. Structural Magnetic Resonance Imaging (sMRI), Positron Emission Tomography (PET) (e.g., FDG-PET and Amyloid-PET), and clinical assessments provide complementary cues on anatomy, metabolism/pathology [10, 18], and symptoms, making multimodal learning a key paradigm for AD staging. However, in real-world clinical cohorts, multimodal acquisition is routinely incomplete due to cost, scanner availability, patient burden, contraindications, and protocol heterogeneity [6, 12], as illustrated in Fig. 1. This missingness is largely protocol-driven. Tabular data (e.g., demographic information and cognitive assessments) are routinely available, while MRI/PET can be absent due to contraindications, radiation concerns, and cost [6, 12], yielding non-random and highly imbalanced modality subsets. Consequently, modality availability varies across subjects, and models must remain accurate and reliable under heterogeneous missingness.




Fig. 1. Modality missing patterns and sample counts for each combination in the ADNI dataset [11].

To learn under heterogeneous modality availability, prior work on incomplete multimodal learning mainly follows two lines: (i) missingness-robust deterministic fusion, e.g., modality dropout [3, 24], mixture-of-fusion strategies [21, 4, 22], and knowledge distillation [6]; and (ii) imputation or cross-modal translation that predicts missing modalities or their features from observed ones [23, 9, 1, 14, 13]. Yet clinical AD cohorts exhibit a long-tailed distribution over observed modality subsets (Fig. 1) due to hierarchical acquisition protocols, creating a missingness pattern mismatch between the dominant subsets seen during training and the rare subsets encountered at deployment. This mismatch raises two challenges. First, models tend to co-adapt to frequent modality subsets, so the representation of a modality can vary with its surrounding context (i.e.,

which other modalities are present), leading to missingness-induced representation shift and degraded generalization on rare subsets [25, 2]. Second, although imputation-based approaches can yield uncertain feature estimates, the previous methods [23, 1] often treat synthesized features as if they were fully observed, lacking an explicit mechanism to down-weight unreliable estimates. Such blind fusion can induce negative transfer, where noisy imputations corrupt multimodal aggregation and reduce performance even when strong observed modalities are available.

To address these challenges, we propose Prototype-anchored Representation Alignment with Product-of-Experts fusion (PRA-PoE), a unified framework for robust multimodal learning under missing data. PRA introduces learnable global prototypes together with availability-conditioned tokens to (i) align modality representations into a subset-consistent space, reducing context-dependent representation shift, and (ii) synthesize missing modality features from the observed ones at the feature level. Building on these representations, UA-PoE treats each modality as a Gaussian expert and performs adaptive Product-of-Experts fusion, where experts with higher uncertainty contribute less via lower precision. This explicitly prevents overconfident reliance on synthesized features and mitigates negative transfer under severe missingness. We evaluate PRA-PoE under a clinically realistic protocol that trains with naturally missing modalities and tests on all non-empty modality subsets, including rare long-tailed patterns. Extensive experiments on two large-scale benchmarks demonstrate the effectiveness of PRA-PoE, which surpasses the strongest baseline by relative improvements of 5.4% in average accuracy on ADNI and 10.9% in average F1 on OASIS-3 under arbitrary missing-modality scenarios.

2 Method

2.1 Overview and Problem Formulation

Overview. We study AD stage classification with incomplete multimodal inputs. Training follows the naturally missing modalities, while testing must work with any non-empty observed subset. This subset-dependent input creates a representation shift and can cause negative transfer when missing modalities are completed with overconfident features. We propose **PRA-PoE**: modality encoders map inputs to a shared latent space, PRA performs subset-conditioned prototype-anchored alignment and completion, and UA-PoE fuses modalities by Gaussian product of experts with uncertainty-aware down-weighting, followed by Monte Carlo sampling for prediction.

Problem formulation. Let $\mathcal{D} = \{(X_i, y_i)\}_{i=1}^N$, where $y_i \in \{\text{CN}, \text{MCI}, \text{AD}\}$ and $X_i = \{x_i^m\}_{m=1}^M$ contains $M = 4$ modalities: T1-sMRI, FDG-PET, Amyloid-PET, and tabular clinical or demographic variables. We use an observation mask $\delta_{i,m} \in \{0, 1\}$ and define the observed set $\mathcal{O}_i = \{m \mid \delta_{i,m} = 1\}$. Tabular variables are usually present but can be missing, while imaging modalities are often absent, which yields imbalanced modality subsets.




Fig. 2. Overall Framework of PRA-PoE.

2.2 Unified Feature Encoding

We employ modality-specific encoders $E_m(\cdot)$ to map the heterogeneous inputs into a shared D -dimensional space, defined as $\mathbf{h}_{i,m} = E_m(x_i^m) \in \mathbb{R}^D$. For 3D neuroimaging modalities (T1-sMRI, FDG-PET, Amyloid-PET), $E_m(\cdot)$ is a 3D volume encoder; for tabular data, $E_m(\cdot)$ is an MLP on standardized variables. When modality m is missing, we set $\mathbf{h}_{i,m} = \mathbf{0}$ to maintain tensor-shape consistency. These zero placeholders do not contribute to cross-modal aggregation, as they are excluded from attention via an availability mask in PRA (Sec. 2.3).

2.3 Prototype-Anchored Representation Alignment (PRA)

Direct fusion of incomplete multimodal features leads to *missingness-induced representation shift*, where a modality’s embedding becomes dependent on the specific subset of co-observed modalities, resulting in context-dependent drift and unstable alignment. To mitigate this issue, we propose **PRA**, which enforces subset-invariant semantics prior to fusion. PRA maps each modality into a shared semantic space by anchoring representations to learnable global prototypes and refining them via availability-conditioned cross-modal attention. Applied to all modalities, including observed ones, **PRA** explicitly corrects subset-induced drift by constraining representations toward global anchors while leveraging the available cross-modal context.

Prototype-anchored Query with Availability Embedding. For each modality m , we use a learnable global prototype $\mathbf{P}_m \in \mathbb{R}^D$ and two learnable availability-conditioned tokens $\mathbf{t}_{\text{obs}}, \mathbf{t}_{\text{mis}} \in \mathbb{R}^D$ to construct query for each modality m , as follows:

$$\mathbf{a}_{i,m} = \delta_{i,m} \mathbf{t}_{\text{obs}} + (1 - \delta_{i,m}) \mathbf{t}_{\text{mis}}, \quad \mathbf{q}_{i,m} = \phi_m([\mathbf{P}_m \parallel \mathbf{a}_{i,m}]) \in \mathbb{R}^D, \quad (1)$$

where $\phi_m : \mathbb{R}^{2D} \rightarrow \mathbb{R}^D$ is a linear layer and $[\cdot \parallel \cdot]$ denotes concatenation.

Masked Multi-head Cross-Attention. For sample i and target modality m , we stack the features of the remaining $M-1$ modalities into a context matrix $\mathbf{C}_{i,m} = [\mathbf{h}_{i,m_1}; \dots; \mathbf{h}_{i,m_{M-1}}] \in \mathbb{R}^{(M-1) \times D}$, where $\{m_1, \dots, m_{M-1}\} = \{1, \dots, M\} \setminus \{m\}$. The prototype-anchored query $\mathbf{q}_{i,m}$ serves as the query, and $\mathbf{C}_{i,m}$ serves as both keys and values. To prevent attention to missing modalities, we derive an additive bias $\mathbf{B}_{i,m} \in \mathbb{R}^{1 \times (M-1)}$ directly from the observation mask:

$$(\mathbf{B}_{i,m})_j = \begin{cases} 0, & \text{if } \delta_{i,m_j} = 1, \\ -\infty, & \text{otherwise,} \end{cases} \quad (2)$$

and compute the refined representation via multi-head cross-attention:

$$\hat{\mathbf{h}}_{i,m} = \text{softmax} \left(\frac{\mathbf{q}_{i,m} \mathbf{C}_{i,m}^\top}{\sqrt{d_k}} + \mathbf{B}_{i,m} \right) \mathbf{C}_{i,m}, \quad (3)$$

where d_k is the per-head dimension. This formulation maintains a fixed $(M-1)$ -token context shape regardless of the missing pattern, while ensuring that only observed modalities contribute to cross-modal refinement.

Adaptive Calibration and PRA Output. For observed modalities, we adaptively fuse the original feature $\mathbf{h}_{i,m}$ with the aligned feature $\hat{\mathbf{h}}_{i,m}$:

$$\alpha_{i,m} = \sigma \left(\psi_m([\mathbf{h}_{i,m} \parallel \hat{\mathbf{h}}_{i,m}]) \right), \quad \mathbf{f}_{i,m} = \alpha_{i,m} \mathbf{h}_{i,m} + (1 - \alpha_{i,m}) \hat{\mathbf{h}}_{i,m}, \quad (4)$$

where ψ_m is a linear layer. The final aligned representation is

$$\mathbf{z}_{i,m} = \text{LN} \left(\delta_{i,m} \mathbf{f}_{i,m} + (1 - \delta_{i,m}) \hat{\mathbf{h}}_{i,m} \right). \quad (5)$$

Thus, missing modalities never serve as cross-modal information sources; their $\mathbf{z}_{i,m}$ are inferred from $\mathcal{O}_i$ via $\hat{\mathbf{h}}_{i,m}$, while observed modalities are calibrated to reduce subset-induced drift.

2.4 Uncertainty-Aware Product-of-Experts (UA-PoE)

Given the PRA-aligned features $\{\mathbf{z}_{i,m}\}_{m=1}^M$, we formulate multimodal fusion as Gaussian PoE in a shared latent space [16]. The standard strategy for handling missing modalities removes absent experts from the product, fusing only over $\mathcal{O}_i$. In contrast, we retain all M experts, including those whose features are completed by PRA, and let each expert independently predict its own mean and variance. Because PRA-completed features carry higher epistemic uncertainty, their corresponding experts learn larger variances and thus contribute less precision to the fused posterior, achieving implicit uncertainty-aware down-weighting within the closed-form PoE framework.

Concretely, a lightweight head $G_m(\cdot)$ maps each $\mathbf{z}_{i,m}$ to a diagonal Gaussian expert $q_{i,m}(\mathbf{s} \mid X_{\mathcal{O}_i}) = \mathcal{N}(\mathbf{s}; \boldsymbol{\mu}_{i,m}, \text{diag}(\boldsymbol{\sigma}_{i,m}^2))$. Combined with a standard

normal prior $p(\mathbf{s}) = \mathcal{N}(\mathbf{0}, \mathbf{I})$, the product $p(\mathbf{s}) \prod_{m=1}^M q_{i,m}$ yields a closed-form Gaussian $q_i(\mathbf{s} \mid X_{\mathcal{O}_i}) = \mathcal{N}(\boldsymbol{\mu}_i, \text{diag}(\boldsymbol{\sigma}_i^2))$ via precision aggregation:

$$\boldsymbol{\sigma}_i^{-2} = \mathbf{1} + \sum_{m=1}^M \boldsymbol{\sigma}_{i,m}^{-2}, \quad \boldsymbol{\mu}_i = \boldsymbol{\sigma}_i^2 \odot \sum_{m=1}^M (\boldsymbol{\mu}_{i,m} \odot \boldsymbol{\sigma}_{i,m}^{-2}). \quad (6)$$

Note that all M experts participate regardless of missingness; experts backed by imputed features naturally exhibit larger $\boldsymbol{\sigma}_{i,m}^2$ and thus contribute less precision to the fused result.

Prediction via Monte Carlo (MC) Sampling. We draw L samples from the fused posterior and average the classifier outputs:

$$\mathbf{p}_i = \frac{1}{L} \sum_{k=1}^L \text{Softmax}(W \mathbf{s}_i^{(k)} + b), \quad \mathbf{s}_i^{(k)} \sim q_i(\mathbf{s} \mid X_{\mathcal{O}_i}). \quad (7)$$

2.5 Loss Function

We train the model with the cross-entropy loss. To regularize the fused distribution and avoid over-confident experts, we add a Kullback-Leibler (KL) Divergence penalty to the standard normal prior $p(\mathbf{s}) = \mathcal{N}(\mathbf{0}, \mathbf{I})$:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N \left(\mathcal{L}_{\text{CE}}(\mathbf{p}_i, y_i) + \beta D_{\text{KL}}(q_i(\mathbf{s} \mid X_{\mathcal{O}_i}) \parallel p(\mathbf{s})) \right), \quad (8)$$

where $\mathbf{p}_i$ is the predictive probability (Sec. 2.4) and $q_i(\mathbf{s} \mid X_{\mathcal{O}_i})$ is the UA-PoE fused Gaussian, β is a hyper-parameter.

3 Experimental Results

3.1 Dataset and Implementation Details

Datasets. We evaluate on baseline ADNI [11] and OASIS-3 [7]. ADNI contains T1-sMRI, FDG-PET, Amyloid-PET, and tabular variables (Age, gender, education, marital status, MMSE, APOE4, ADNI.EF, ADNI.MEM); OASIS-3 contains T1-sMRI, Flair-MRI, Amyloid-PET, and tabular variables (Age, gender, education, APOE4, MMSE), which are commonly used in AD multi-modal learning [19, 15, 20]. We follow [19] for 3D image preprocessing and Min-Max normalize tabular features; all volumes are resized to 128^3 . ADNI contains 825/980/429 Cognitively Normal (CN)/Mild Cognitive Impairment (MCI)/AD subjects (N=2,234), and OASIS-3 has 791/197 CN/AD subjects (N=988). We split complete subjects into dev/test (1:1), then split (dev + incomplete) into train/val (6:1) stratified by diagnosis, resulting in 1504/259/471 (ADNI) and 757/127/104 (OASIS-3) for train/val/test. Train sets exhibit realistic missingness (e.g., 37.4% AV45 available), while test sets are complete to simulate arbitrary missing-modality subsets at inference (Fig. 1).

Implementation Details. We train with Adam (LR 2×10^{-4}), 15-epoch warm-up + cosine decay, batch size 32, up to 100 epochs with early stopping (patience 40); $\beta = 0.001$ unless specified. Experiments use two NVIDIA RTX 4090 GPUs.

Table 1. Performance comparison (mean (std), %) on ADNI and OASIS-3 for average metrics on 15 modality combinations, with computational efficiency. The orange background refers to the best one, and the second-best results are underlined.

Methods	CN vs. MCI vs. AD			CN vs. AD			Computational Efficiency		
	ACC	AUC	Macro-F1	ACC	AUC	F1 _{AD}	Params (M)	FLOPs (G)	Speed (ms)
mmFormer [24]	51.7 (2.1)	68.4 (1.1)	44.6 (2.6)	86.8 (2.1)	72.8 (6.2)	40.3 (12.2)	26.4	65.2	31.7
ShaSpec [17]	56.9 (1.9)	73.6 (1.4)	54.1 (1.0)	87.7 (0.0)	88.2 (2.4)	51.0 (13.5)	144.1	482.4	43.9
AnyMod [3]	59.7 (1.1)	76.6 (1.3)	56.8 (3.6)	85.5 (1.6)	82.3 (4.8)	51.5 (4.5)	11.9	92.4	129.8
M3Care [23]	59.7 (1.1)	<u>78.7 (1.5)</u>	59.9 (2.0)	88.5 (1.7)	<u>88.7 (0.7)</u>	52.9 (6.2)	19.2	18.4	11.9
FuseMoE [4]	58.5 (1.2)	77.6 (0.3)	57.0 (2.0)	88.9 (1.7)	86.6 (1.3)	57.7 (5.9)	118.8	<u>18.5</u>	2.5
Flex-MoE [21]	57.1 (1.7)	74.3 (1.4)	48.1 (2.5)	88.7 (2.3)	87.5 (1.0)	57.3 (7.0)	<u>13.8</u>	18.4	28.0
MoE-retriever [22]	60.6 (3.1)	78.3 (1.9)	57.5 (2.6)	89.8 (0.3)	86.4 (1.9)	58.0 (2.8)	15.1	18.4	74.2
PRA-PoE (Ours)	<u>63.9 (0.4)</u>	81.4 (0.3)	61.7 (0.8)	<u>90.3 (0.6)</u>	89.5 (0.5)	<u>64.3 (0.9)</u>	<u>13.8</u>	18.4	<u>8.1</u>

Table 2. ACC comparison under 8 clinical common modality combinations on ADNI dataset.

Modalities	ACC (%)							
Tab T1 FDG Amyloid	mmFormer	ShaSpec	AnyMod	M3Care	FuseMoE	Flex-MoE	MoE-retr	PRA-PoE
✓	56.8 (5.0)	<u>67.9 (3.5)</u>	64.2 (2.4)	60.3 (2.9)	63.6 (2.5)	65.2 (3.3)	61.3 (7.6)	65.3 (0.5)
✓ ✓	56.3 (4.0)	56.3 (3.7)	<u>64.8 (1.8)</u>	63.4 (5.4)	61.1 (4.6)	64.8 (2.3)	63.7 (7.5)	<u>68.7 (1.6)</u>
✓ ✓ ✓	60.6 (2.0)	59.9 (1.7)	67.4 (0.5)	65.6 (1.1)	68.0 (0.9)	<u>69.1 (0.7)</u>	65.0 (2.6)	70.1 (0.5)
✓ ✓ ✓ ✓	53.3 (2.8)	58.5 (3.3)	<u>65.9 (0.6)</u>	61.1 (3.3)	59.7 (3.1)	<u>62.1 (2.3)</u>	62.0 (6.2)	68.4 (0.5)
✓ ✓ ✓ ✓ ✓	60.2 (1.8)	60.7 (1.9)	66.8 (1.0)	66.7 (2.7)	68.3 (1.2)	<u>70.3 (0.4)</u>	67.3 (2.1)	69.5 (1.2)
✓ ✓ ✓ ✓ ✓ ✓	53.8 (3.0)	57.7 (4.6)	<u>66.2 (0.6)</u>	64.1 (4.4)	60.4 (5.4)	66.0 (3.4)	64.0 (5.0)	69.1 (0.8)
✓ ✓ ✓ ✓ ✓ ✓ ✓	59.9 (2.1)	61.7 (3.3)	<u>67.3 (0.8)</u>	65.7 (1.3)	69.4 (1.5)	69.8 (0.8)	69.3 (2.1)	71.9 (2.1)
✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓	61.0 (1.9)	60.7 (3.6)	66.9 (1.4)	65.8 (2.8)	68.8 (1.8)	<u>69.5 (0.1)</u>	69.0 (4.0)	72.1 (0.8)
Average	57.7 (2.7)	60.4 (2.5)	66.2 (1.0)	64.1 (0.7)	64.9 (1.2)	67.1 (1.2)	<u>65.2 (4.1)</u>	69.4 (1.3)

3.2 Comparison with State-of-the-art Methods.

We evaluate all methods under arbitrary missing-modality settings for CN vs. MCI vs. AD (ADNI) and CN vs. AD (OASIS-3). For each test subject, we randomly drop modalities to form different modality subsets and report the average performance over all 15 non-empty modality combinations. We compare PRA-PoE with state-of-the-art approaches for incomplete multimodal learning, including mmFormer [24] and ShaSpec [17] (adapted to classification by replacing decoders with classification heads), M3Care [23], AnyMod [3], and recent flexible Mixture-of-Experts (MoE) frameworks (FuseMoE [4], Flex-MoE [21], and MoE-retriever [22]). All experiments are repeated three times with different random seeds, and results are reported as mean(standard deviation).

The overall quantitative results are summarized in Table 1. For CN vs. MCI vs. AD, PRA-PoE achieves the best average performance across modality com-

binations, with an average ACC of 63.9%, AUC of 81.4%, and Macro-F1 of 61.7%, outperforming MoE-retriever (the best baseline in ACC) by relative improvements of **5.4%** in ACC, **4.0%** in AUC, and **7.8%** in Macro-F1 score on the ADNI dataset. For CN vs. AD, PRA-PoE outperforms other methods with an average ACC of 90.3%, AUC of 89.5%, and $F1_{AD}$ of 64.3% on the OASIS-3 dataset. Due to page limitations, we only present the Accuracy of 8 clinically common modality combinations in ADNI and OASIS-3 in Table 2 and Table 3. In addition, Table 1 reports model complexity: PRA-PoE is parameter-efficient (13.8M) and computationally lightweight (18.4G FLOPs), while maintaining fast inference (8.1 ms/sample), making it suitable for practical low-latency deployment.

Table 3. ACC (%) comparison under 8 clinical common modality combinations on OASIS-3 dataset.

Modalities				ACC (%)							
Tab	T1	Flair	Amyloid	mmFormer	ShaSpec	AnyMod	M3Care	FuseMoE	Flex-MoE	MoE-retr	PRA-PoE
✓				85.6 (1.7)	86.5 (6.7)	90.4 (1.0)	89.7 (2.8)	89.4 (1.0)	88.5 (1.0)	89.1 (2.1)	90.4 (1.9)
✓	✓			85.6 (3.5)	89.4 (1.0)	86.2 (1.5)	87.5 (3.3)	88.8 (1.5)	85.3 (4.5)	90.7 (2.4)	91.7 (1.1)
✓		✓		82.7 (2.5)	84.3 (2.4)	84.6 (2.9)	86.2 (2.4)	85.3 (4.5)	88.1 (2.8)	86.5 (1.9)	89.7 (0.6)
✓			✓	87.5 (4.4)	90.4 (1.0)	90.4 (3.3)	89.1 (3.1)	92.9 (0.6)	91.3 (1.7)	90.4 (1.7)	92.9 (2.2)
✓	✓	✓		86.5 (1.0)	87.8 (1.5)	84.3 (1.5)	87.8 (2.0)	87.8 (2.9)	86.2 (3.6)	92.3 (1.7)	91.3 (1.0)
✓	✓		✓	88.8 (3.1)	90.7 (2.0)	89.7 (3.6)	90.4 (2.5)	92.6 (1.5)	89.7 (1.5)	92.3 (1.0)	92.0 (0.6)
✓		✓	✓	88.5 (1.0)	87.5 (1.9)	87.5 (5.4)	90.7 (1.1)	90.7 (2.4)	92.0 (2.2)	92.3 (0.0)	92.9 (1.1)
✓	✓	✓	✓	88.5 (1.7)	88.5 (1.0)	88.1 (4.9)	90.1 (2.0)	90.1 (2.4)	89.7 (3.6)	93.3 (1.0)	93.6 (0.6)
Average				86.7 (2.2)	88.1 (0.3)	87.7 (2.7)	88.9 (1.7)	89.7 (1.4)	88.9 (2.1)	90.9 (0.8)	91.8 (0.8)




Fig. 3. Ablation study (a) and hyperparameter sensitivity analysis (b).

3.3 Ablation Study and Sensitivity Analysis

Fig. 3 reports ablations and hyperparameter sensitivity on ADNI. In Fig. 3(a), PRA-PoE achieves the best average performance compared to variants without PRA, UA-PoE, or MC sampling across missingness patterns, showing all

components contribute to robustness. Removing PRA yields the largest drop (Average ACC $64.1\% \rightarrow 61.0\%$), indicating its key role in mitigating distribution shift, while removing UA-PoE also degrades performance (Average ACC $64.1\% \rightarrow 62.2\%$), supporting uncertainty-aware fusion. In Fig. 3(b), increasing MC draws L improves performance and saturates at $L \approx 10$. The best KL weight is $\beta = 0.001$; larger β hurts performance, suggesting overly strong regularization limits discriminative representations.

4 Conclusion

In this paper, we propose PRA-PoE, a robust framework for AD classification under arbitrary modality missingness. By combining Prototype-anchored Representation Alignment to mitigate missingness-induced representation shift with uncertainty-aware precision fusion to attenuate unreliable modality estimates, PRA-PoE delivers robust predictions across diverse missing patterns. Experiments on two datasets demonstrate improved performance over strong baselines, indicating its potential for clinical application. Future work will extend the framework to longitudinal modeling and external multi-center validation to further assess generalizability.

Acknowledgments. This work was partially supported by RGC Collaborative Research Fund (No. C5055-24G), the Start-up Fund of The Hong Kong Polytechnic University (No. P0045999), the Seed Fund of the Research Institute for Smart Ageing (No. P0050946), and Tsinghua-PolyU Joint Research Initiative Fund (No. P0056509), and PolyU UGC funding (No. P0053716).

Disclosure of Interests. The authors have no competing interests.

References

1. Chen, Y., Pan, Y., Xia, Y., Yuan, Y.: Disentangle first, then distill: a unified framework for missing modality imputation and Alzheimer’s disease diagnosis. *IEEE Transactions on Medical Imaging* **42**(12), 3566–3578 (2023)
2. Dai, R., Li, C., Yan, Y., Mo, L., Qin, K., He, T.: Unbiased missing-modality multimodal learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 24507–24517 (2025)
3. Feng, Y., Gao, B., Deng, S., Qiu, A., Qin, J.: Unified multi-modal learning for any modality combinations in Alzheimer’s disease diagnosis. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 487–497. Springer (2024)
4. Han, X., Nguyen, H., Harris, C., Ho, N., Saria, S.: FuseMoE: Mixture-of-experts transformers for fleximodal fusion. *Advances in Neural Information Processing Systems* **37**, 67850–67900 (2025)
5. Hou, W., Yang, G., Du, Y., Lau, Y., Liu, L., He, J., Long, L., Wang, S.: Ada-agent: Llm agent for alzheimer’s disease analysis with collaborative coordinator. In: *International Workshop on Agentic AI for Medicine*. pp. 23–32. Springer (2025)

6. Kwak, M.G., Mao, L., Zheng, Z., Su, Y., Lure, F., Li, J.: A cross-modal mutual knowledge distillation framework for Alzheimer’s disease diagnosis: Addressing incomplete modalities. *IEEE Transactions on Automation Science and Engineering* **22**, 14218–14233 (2025)
7. LaMontagne, P.J., Benzinger, T.L., Morris, J.C., Keefe, S., Hornbeck, R., Xiong, C., Grant, E., Hassenstab, J., Moulder, K., Vlassenko, A.G., et al.: OASIS-3: longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and Alzheimer disease. *medrxiv* pp. 2019–12 (2019)
8. Lei, B., Li, Y., Fu, W., Yang, P., Chen, S., Wang, T., Xiao, X., Niu, T., Fu, Y., Wang, S., et al.: Alzheimer’s disease diagnosis from multi-modal data via feature inductive learning and dual multilevel graph neural network. *Medical Image Analysis* **97**, 103213 (2024)
9. Meng, X., Sun, K., Xu, J., He, X., Shen, D.: Multi-modal modality-masked diffusion network for brain mri synthesis with random modality missing. *IEEE Transactions on Medical Imaging* (2024)
10. Minoshima, S., Cross, D., Thientunyakit, T., Foster, N.L., Drzezga, A.: 18f-fdg pet imaging in neurodegenerative dementing disorders: insights into subtype classification, emerging disease categories, and mixed dementia with copathologies. *Journal of nuclear medicine* **63**(Supplement 1), 2S–12S (2022)
11. Mueller, S.G., Weiner, M.W., Thal, L.J., Petersen, R.C., Jack, C., Jagust, W., Trojanowski, J.Q., Toga, A.W., Beckett, L.: The Alzheimer’s disease neuroimaging initiative. *Neuroimaging Clinics* **15**(4), 869–877 (2005)
12. Odusami, M., Damaševičius, R., Milieškaitė-Belousovienė, E., Maskeliūnas, R.: Alzheimer’s disease stage recognition from mri and pet imaging data using pareto-optimal quantum dynamic optimization. *Heliyon* **10**(15) (2024)
13. Pan, Y., Liu, M., Lian, C., Zhou, T., Xia, Y., Shen, D.: Synthesizing missing pet from mri with cycle-consistent generative adversarial networks for alzheimer’s disease diagnosis. In: *International conference on medical image computing and computer-assisted intervention*. pp. 455–463. Springer (2018)
14. Pan, Y., Liu, M., Xia, Y., Shen, D.: Disease-image-specific learning for diagnosis-oriented neuroimage synthesis with incomplete multi-modality data. *IEEE transactions on pattern analysis and machine intelligence* **44**(10), 6839–6853 (2022)
15. Qiu, S., Miller, M.I., Joshi, P.S., Lee, J.C., Xue, C., Ni, Y., Wang, Y., De Anda-Duran, I., Hwang, P.H., Cramer, J.A., et al.: Multimodal deep learning for Alzheimer’s disease dementia assessment. *Nature Communications* **13**(1), 3404 (2022)
16. Shi, Y., Paige, B., Torr, P., et al.: Variational mixture-of-experts autoencoders for multi-modal deep generative models. *Advances in Neural Information Processing Systems* **32** (2019)
17. Wang, H., Chen, Y., Ma, C., Avery, J., Hull, L., Carneiro, G.: Multi-modal learning with missing modality via shared-specific feature modelling. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 15878–15887 (2023)
18. Xie, H., Li, Y., Wu, X., Wang, R., Long, X., Su, M., Chen, Q., Li, L., Tian, R., Jia, Z.: The image quality, amyloid- β detectability, and acquisition time of clinical florbetapir positron emission tomography in alzheimer’s disease and healthy adults. *Quantitative Imaging in Medicine and Surgery* **13**(12), 7765 (2023)
19. Xue, C., Kowshik, S.S., Lteif, D., Puducheri, S., Jasodanand, V.H., Zhou, O.T., Walia, A.S., Guney, O.B., Zhang, J.D., Poésy, S., et al.: AI-based differential diagnosis of dementia etiologies on multimodal data. *Nature Medicine* **30**(10), 2977–2989 (2024)

20. Yang, G., Du, K., Yang, Z., Du, Y., Cheung, E.Y.W., Zheng, Y., Yang, M., Kourtzi, Z., Schonlieb, C.B., Wang, S.: ADFound: A foundation model for diagnosis and prognosis of Alzheimer’s disease. *IEEE Journal of Biomedical and Health Informatics* **29**(11), 8395–8408 (2025)
21. Yun, S., Choi, I., Peng, J., Wu, Y., Bao, J., Zhang, Q., Xin, J., Long, Q., Chen, T.: Flex-MoE: Modeling arbitrary modality combination via the flexible mixture-of-experts. *Advances in Neural Information Processing Systems* **37**, 98782–98805 (2025)
22. Yun, S., Xin, J., Choi, I., Peng, J., Ding, Y., Long, Q., Chen, T.: Generate, then retrieve: Addressing missing modalities in multimodal learning via generative AI and MoE. In: *Workshop on Large Language Models and Generative AI for Health at AAAI 2025* (2025)
23. Zhang, C., Chu, X., Ma, L., Zhu, Y., Wang, Y., Wang, J., Zhao, J.: M3Care: Learning with missing modalities in multimodal healthcare data. In: *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*. pp. 2418–2428 (2022)
24. Zhang, Y., He, N., Yang, J., Li, Y., Wei, D., Huang, Y., Zhang, Y., He, Z., Zheng, Y.: mmFormer: Multimodal medical transformer for incomplete multimodal learning of brain tumor segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 107–117. Springer (2022)
25. Zhao, B., Zhang, W., Zou, Z.: Mce: Towards a general framework for handling missing modalities under imbalanced missing rates. *Pattern Recognition* p. 112591 (2025)