Sparsely Labeled fMRI Data Denoising with Meta-Learning-Based Semi-Supervised Domain Adaptation

Keun-Soo Heo[1], Ji-Wung Han[1], Soyeon Bak[1], Minjoo Lim[1], Bogyeong Kang[1], Sang-Jun Park[1], Weili Lin[2], Han Zhang[3], Dinggang Shen[3], and Tae-Eui Kam[1*]

[1] Department of Artificial Intelligence, Korea University, Republic of Korea
kamte@korea.ac.kr
[2] Department of Radiology, University of North Carolina at Chapel Hill, USA
[3] School of Biomedical Engineering, ShanghaiTech University, China

Abstract. Functional magnetic resonance imaging (fMRI) denoising is a crucial preprocessing step in neuroimaging studies, as noise degrades the reliability of downstream analyses. Previous approaches for fMRI denoising either rely on predefined noise patterns or train dataset-specific models, restricting their reliability across various datasets due to inter-dataset variations in scanner types, scanning protocols, and preprocessing pipelines. Additionally, applying previous approaches to new datasets requires extensive expert signal/noise annotations. To mitigate this reliance, leveraging existing datasets to train sparsely labeled datasets is a practical solution, but inconsistencies in labeling criteria hinder effective adaptation. To address these challenges, we propose a meta-learning-based semi-supervised domain adaptation framework, enabling the learning of dataset-irrelevant features from sparsely labeled datasets by leveraging existing labeled datasets with two key components: (1) a dataset-irrelevant feature extractor trained by meta-learning to capture noise patterns across multiple datasets, and (2) dataset-specific classifiers optimized by decoupled training to handle inconsistencies in labeling criteria. Our proposed approach shows outstanding performance on four fMRI datasets in both fully labeled and sparsely labeled conditions.

Keywords: fMRI denoising · Meta-learning · Semi-supervised domain adaptation · Decoupled training · Sparsely labeled data

1 Introduction

Functional magnetic resonance imaging (fMRI) is one of the most widely used neuroimaging techniques for studying functional connectomes in the brain [2,3,13,15]. However, fMRI data are easily affected by noise originating from experimental environments, scanner artifacts, and non-neuronal physiological fluctuations [7,20,21]. This noise degrades the signal-to-noise ratio (SNR) and reduces the statistical power of fMRI studies in interpreting brain functions [17]. Thus, fMRI denoising, which separates neuronal signals from noise, has become an essential preprocessing step for fMRI-based neuroimaging studies [7,8,17].

Previous approaches for fMRI denoising leverage component-based methods, such as principal component analysis (PCA) and independent component analysis (ICA) [11,14,19,22,24]. Particularly, spatial ICA is widely adopted, as it decomposes fMRI data into independent components (ICs), which consist of pairs of a 3D spatial map and its corresponding 1D time series, enabling the separation of noise from fMRI data by classifying ICs as either signal or noise [1,22]. Based on spatial ICA, FIX [24] classifies signal/noise ICs using machine-learning-based classifiers trained on hand-crafted spatial and temporal features. Recently, deep learning-based methods [11,14,19] employ convolutional neural networks (CNNs) [10] for automatic feature extraction and noise classification.

Nevertheless, they either rely on predefined noise patterns [22] or train dataset-specific models for each dataset [19]. This reliance limits their capacity to learn noise patterns across diverse datasets due to inter-dataset variations in scanner types, scanning protocols, and preprocessing [28,31]. Thus, reducing inter-dataset variations is crucial for effectively capturing a wide range of noise patterns across diverse datasets. Furthermore, applying previous approaches to new datasets requires extensive expert annotations for signal/noise ICs, which demand significant time and cost. To alleviate this burden, leveraging existing labeled datasets is effective for training new datasets with minimal annotations, i.e., sparsely labeled datasets [23]. However, labeling criteria for signal/noise annotations vary depending on research objectives and expert judgment [7,11,27], leading to inconsistencies across datasets. Therefore, addressing these inconsistencies is crucial for ensuring effective application to sparsely labeled datasets.

To address these challenges, we propose a novel meta-learning-based semi-supervised domain adaptation framework for fMRI denoising, enabling the learning of dataset-irrelevant features from a sparsely labeled dataset by leveraging existing labeled datasets. To reduce inter-dataset variations, we adopt meta-learning for the **dataset-irrelevant feature extractor** that effectively learns a wide range of noise patterns across multiple well-labeled datasets and a sparsely labeled dataset. To mitigate inconsistencies in labeling criteria, we introduce decoupled training which optimizes the **dataset-specific classifiers** separately from the dataset-irrelevant feature extractor, allowing each dataset-specific classifier to adapt to its respective dataset's labeling criteria while preserving dataset-irrelevant feature extraction. Our proposed framework iteratively enhances both the feature extractor and the classifiers, outperforming previous approaches on four fMRI datasets under both sparsely labeled and fully labeled settings.

The contributions of this study are summarized as follows:

- We propose a meta-learning-based semi-supervised domain adaptation framework for fMRI denoising, enabling the learning of dataset-irrelevant features from a sparsely labeled dataset by leveraging existing labeled datasets.
- We introduce decoupled training that optimizes dataset-specific classifiers separately from the feature extractor to address differences in labeling criteria across datasets while preserving dataset-irrelevant features.
- Our proposed framework achieves state-of-the-art performance across four fMRI datasets in both fully labeled and sparsely labeled conditions.

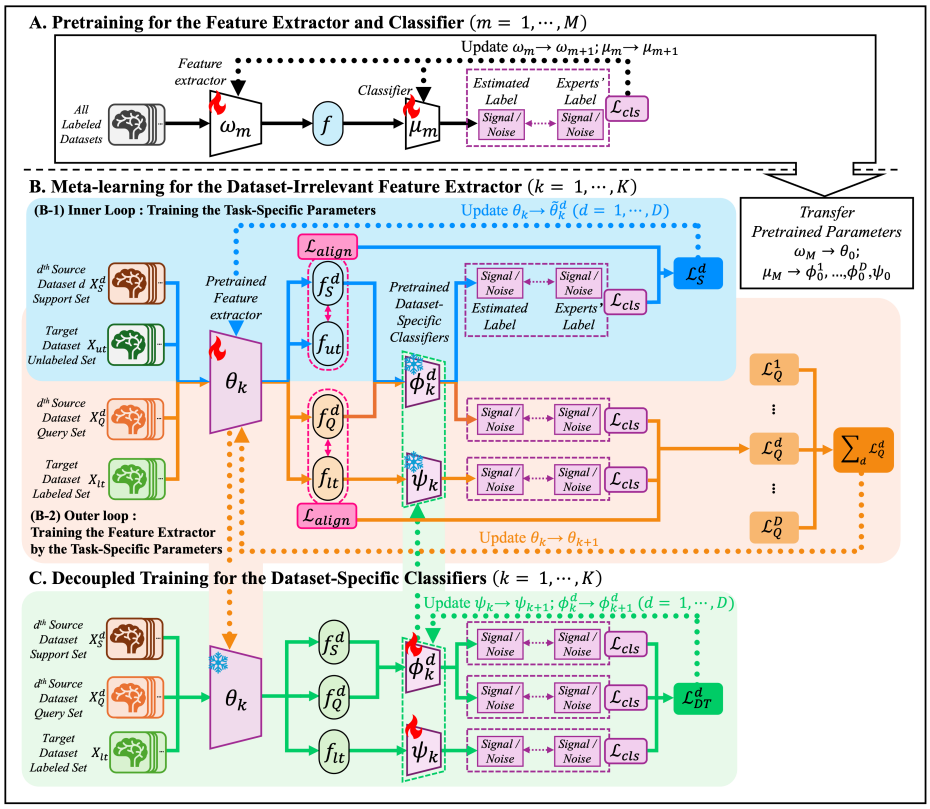

Fig. 1. Overview of our proposed framework for fMRI denoising, which learns dataset-irrelevant features from a sparsely labeled dataset (a target domain) by leveraging existing labeled datasets (source domains) with three stages: (A) pretraining for the feature extractor and classifier, (B) meta-learning for the dataset-irrelevant feature extractor, and (C) decoupled training for the dataset-specific classifiers.

2 Method

We propose a meta-learning-based semi-supervised domain adaptation framework, which learns dataset-irrelevant features from a sparsely labeled dataset (a target domain) with existing labeled datasets (source domains). As shown in Fig. 1, our proposed framework consists of three stages: (A) pretraining for the feature extractor and classifier, (B) meta-learning for the dataset-irrelevant feature extractor, and (C) decoupled training for the dataset-specific classifiers.

2.1 Pretraining for the Feature Extractor and Classifier

We train a feature extractor and a classifier using the fully labeled source domains and a labeled set from the target domain to initialize parameters for the

subsequent meta-learning and decoupled training. We construct the feature extractor with parameters ω to encode each IC into a feature representation, based on the CNN-based architecture proposed in [19]. The classifier μ is implemented as a single-layer fully connected network and predicts whether the feature of the input IC corresponds to a signal or noise component. We adopt cross-entropy loss as the classification loss $\mathcal{L}_{cls}$ to train both models. After M training steps, the feature extractor's parameters ω_M are set as θ_0. Similarly, the classifier's parameters μ_M are used to initialize the parameters of the dataset-specific classifiers $\phi_0^1, \ldots, \phi_0^D, \psi_0$, which consist of $D+1$ individual classifiers, where each classifier learns a specific labeling criterion for its respective dataset. Incorporating labeled data at this stage reduces ambiguity in signal/noise annotations and offers a stable initialization for subsequent processing stages.

2.2 Meta-learning for the Dataset-Irrelevant Feature Extractor

To effectively extract dataset-irrelevant features, we employ a meta-learning approach based on Model-Agnostic Meta-Learning (MAML) [5], pairing each source domain with target domain, resulting in a total of D pairs. The meta-learning process consists of an inner loop and an outer loop, iterating through meta-learning steps ($k = 1, \ldots, K$). For the d^{th} pair ($d = 1, \ldots, D$), we train task-specific parameters $\tilde{\theta}_k^d$ by minimizing both the classification loss and the alignment loss using the support set of the d^{th} source domain and the unlabeled set of the target domain. We adopt cross-entropy loss as the classification loss $\mathcal{L}_{cls}$ and contrastive-based loss [16] as the alignment loss $\mathcal{L}_{align}$. Once the task-specific parameters $(\tilde{\theta}_k^1, \ldots, \tilde{\theta}_k^D)$ are obtained in the inner loop for all pairs, the outer loop updates the parameters of the feature extractor θ_k by validating all task-specific parameters using the query set of the source domain and the labeled set of the target domain, while minimizing both the classification loss and the alignment loss. The meta-learning approach enables robust adaptation to sparsely labeled data under limited supervision [6,9].

2.3 Decoupled Training for the Dataset-Specific Classifiers

After training the dataset-irrelevant feature extractor, we optimize dataset-specific classifiers to address labeling inconsistencies across datasets. We introduce decoupled training for the dataset-specific classifiers with parameters $\phi^1, \ldots, \phi^D, \psi$, where ϕ^d corresponds to the classifier for the d^{th} source domain, and ψ is the classifier for the target domain. Unlike previous approaches [11,19] that jointly update feature extraction and classification, we train the dataset-specific classifiers independently from the feature extractor to accommodate dataset-specific labeling criteria. During this stage, we freeze the feature extractor to prevent interference with dataset-irrelevant feature extraction [29]. Using the source domains and the labeled set of the target domain, we update the parameters of the dataset-specific classifiers by minimizing the classification loss $\mathcal{L}_{cls}$ separately for each dataset.

Table 1. The detail of four fMRI datasets from different studies: the Human Connectome Project (HCP) [26], the Baby Connectome Project (BCP) [12], the Whitehall II Multi-band (WHII-MB6) [4], and the Whitehall II Standard (WHII-STD) [4].

Dataset	Scanner	TR(s)	SR(mm^3)	# Subjects	Age	M/F	# Samples	# ICs	# Signal ICs	# Noise ICs
HCP [26]	3T Siemens Connectome Skyra	0.7	2×2×2	25	22–35 years	8/17	100	22,877	2,326	20,551
BCP [12]	3T Siemens Prisma	0.8	2×2×2	32	379 ± 186 days	-	99	14,850	3,406	11,444
WHII-MB6 [4]	3T Siemens Magnetom Verio	1.3	2×2×2	25	63–75 years	33% F	39	5,143	795	4,348
WHII-STD [4]	3T Siemens Magnetom Verio	3.0	3×3×3	45	63–75 years	33% F	40	2,585	422	2,783

By alternating between meta-learning for the dataset-irrelevant feature extractor and decoupled training for the dataset-specific classifiers over K iterations, our proposed framework progressively improves classification performance while enhancing adaptation from fully labeled source domains to a sparsely labeled target domain, without compromising the dataset-irrelevant feature extractor.

3 Experiment and Discussion

3.1 Datasets and Experiment Settings

Datasets. We utilize four fMRI datasets: Human Connectome Project (HCP) [26], Baby Connectome Project (BCP) [12], and two Whitehall II imaging study datasets, multi-band (WHII-MB6) and standard (WHII-STD) [4], following previous approaches [11,19]. Table 1 shows the detail of the datasets, highlighting inter-dataset variations, such as scanner types, temporal resolution, spatial resolution, subject age, and scan counts.

Experiment Settings. We employ 5-fold cross-validation by uniformly partitioning samples into five folds. Performance metrics, including accuracy (ACC), sensitivity (SEN), specificity (SPEC), F1-score (F1), and G-measure (GM), are evaluated per sample [18]. Sensitivity measures the detection of signal ICs, while specificity measures the detection of noise ICs. Since signal ICs are significantly outnumbered by noise ICs, evaluating performance using the F1-score and G-measure, which provide balanced evaluations of signal/noise IC classification, is essential to ensure robust classification in imbalanced datasets.

Implementation Details. We set the batch size to 12, the number of iterations to $M = 10$, and the feature dimension to 96. The training process alternates between meta-learning for the dataset-irrelevant feature extractor and decoupled training for the dataset-specific classifiers, iterating for a total of $K = 300$ cycles to progressively optimize both components, with a learning rate of 0.01. Our implementation is available on: https://github.com/KeunsooHeo/metaclean

Table 2. Performance comparison for signal/noise IC classification with the baseline model [19]. The upper table presents results in the sparsely labeled (10%) setting, while the lower table provides results for the fully labeled (100%) setting across four fMRI datasets. We present the mean ± standard deviation for accuracy (ACC), sensitivity (SEN), specificity (SPEC), F1-score (F1), and G-measure (GM) on the target dataset. We evaluate three settings by adding key components: (A) pretraining for the feature extractor and classifier by jointly learning with the other datasets (PT), (B) meta-learning for dataset-irrelevant feature extractor (ML), and (C) decoupled training for dataset-specific classifiers (DT). A "✓" indicates the inclusion of the component, while "-" denotes its absence. The highest mean performance in each dataset-label condition is highlighted in bold. Statistical significance compared to Setting (C) is determined using the Wilcoxon signed-rank test (*$p < 0.05$, **$p < 0.001$) [30].

Label Ratio	Target Dataset	Method	Components			Evaluation Metrics				
			PT	ML	DT	ACC (%)	SEN (%)	SPEC (%)	F1 (%)	GM (%)
10%	HCP	Baseline [19]	-	-	-	98.48±0.94**	91.61±8.38**	**99.26**±0.70*	92.36±4.94**	95.25±4.49**
		Setting (A)	✓	-	-	98.58±1.18**	95.55±4.99**	98.92±1.31	93.37±5.03**	97.18±2.54**
		Setting (B)	✓	✓	-	98.75±0.95*	95.55±6.15**	99.11±1.02	94.16±3.80**	97.26±3.13*
		Setting (C)	✓	✓	✓	**98.93**±0.79	**97.82**±3.34	99.06±0.78	**94.99**±3.50	**98.42**±1.75
	BCP	Baseline [19]	-	-	-	96.13±2.02	86.21±14.33**	**98.33**±1.95	88.74±10.88*	91.63±8.57**
		Setting (A)	✓	-	-	96.44±1.78	89.21±10.96	97.95±1.84	90.39±8.04	93.25±6.16
		Setting (B)	✓	✓	-	96.45±2.22	**91.08**±12.43	97.44±2.65	89.96±11.21	**93.90**±7.17
		Setting (C)	✓	✓	✓	**96.60**±1.78	90.01±11.91	97.92±1.99	**90.55**±9.47	93.61±6.75
	WHII-MB6	Baseline [19]	-	-	-	93.89±5.16**	75.70±22.13**	95.82±8.53*	78.01±17.86**	82.85±17.45**
		Setting (A)	✓	-	-	98.71±1.12	95.04±4.66*	**99.40**±1.05	**96.16**±3.00	97.16±2.35*
		Setting (B)	✓	✓	-	98.04±1.70**	94.42±5.74*	98.64±2.09	93.67±5.18*	96.45±2.96**
		Setting (C)	✓	✓	✓	**98.83**±1.02	**97.19**±3.79	99.13±1.14	95.60±6.52	**98.13**±1.93
	WHII-STD	Baseline [19]	-	-	-	86.03±5.68**	0.00 ±0.00**	**100.00**±0.00**	0.00 ±0.00**	0.00 ±0.00**
		Setting (A)	✓	-	-	95.03±3.12**	**93.95**±8.36*	95.30±3.63**	82.31±13.63**	94.49±4.22
		Setting (B)	✓	✓	-	96.88±2.88	81.63±17.67**	99.39±1.02	86.75±13.54	89.46±10.69*
		Setting (C)	✓	✓	✓	**97.67**±1.82	91.12±9.25	98.90±1.69	**91.07**±6.69	**94.79**±4.82

Label Ratio	Target Dataset	Method	Components			Evaluation Metrics				
			PT	ML	DT	ACC (%)	SEN (%)	SPEC (%)	F1 (%)	GM (%)
100%	HCP	Baseline [19]	-	-	-	**99.31**±0.53	97.07±3.92*	**99.57**±0.52*	96.71±2.47	98.29±1.96*
		Setting (A)	✓	-	-	99.12±0.64*	97.09±3.66*	99.34±0.64	95.87±2.59**	98.19±1.83**
		Setting (B)	✓	✓	-	99.21±0.67*	98.09±3.34	99.33±0.76	96.28±2.83*	98.69±1.66
		Setting (C)	✓	✓	✓	**99.31**±0.58	**98.10**±2.76	99.47±0.57	**96.85**±2.47	**98.77**±1.40
	BCP	Baseline [19]	-	-	-	96.02±2.20**	91.67±10.73*	96.78±2.59**	89.48±9.74*	93.96±6.09
		Setting (A)	✓	-	-	96.46±1.90	89.60±12.19	97.82±2.27	90.28±9.48	93.32±6.95*
		Setting (B)	✓	✓	-	96.55±1.96	**92.58**±9.55**	97.11±2.98**	**91.04**±8.29	**94.64**±4.95
		Setting (C)	✓	✓	✓	**96.69**±1.79	90.18±12.42	**98.00**±1.79	90.56±9.92	93.71±7.24
	WHII-MB6	Baseline [19]	-	-	-	98.66±1.26*	94.40±6.23*	99.31±1.19	95.18±4.24*	96.77±3.27*
		Setting (A)	✓	-	-	98.68±1.33*	96.30±5.21	99.01±1.57*	95.07±7.01*	97.60±2.71
		Setting (B)	✓	✓	-	98.72±1.38	95.79±4.52	99.38±1.32	95.17±7.94	97.54±2.44
		Setting (C)	✓	✓	✓	**99.07**±1.03	**96.92** ±4.12	**99.60**±0.64	**97.37**±2.75	**98.23**±2.17
	WHII-STD	Baseline [19]	-	-	-	97.81±2.26*	87.35±13.81**	**99.65**±0.74*	91.27±9.06*	92.98±7.76**
		Setting (A)	✓	-	-	98.10±1.53*	93.29±8.62	99.04±1.17	92.45±6.37*	96.00±4.54
		Setting (B)	✓	✓	-	97.73±2.58*	89.70±12.26*	99.35±1.45	91.95±8.33	94.15±6.84*
		Setting (C)	✓	✓	✓	**98.52**±1.38	**94.14**±7.79	99.32±1.10	**94.20**±5.55	**96.60**±4.01

3.2 Performance Comparison

Table 2 presents a performance comparison of signal/noise IC classification between the baseline model [19] and our proposed framework under both sparsely labeled (10%, top) and fully labeled (100%, bottom) conditions across four fMRI datasets. The evaluation includes accuracy (ACC), sensitivity (SEN), specificity (SPEC), F1-score (F1), and G-measure (GM). We assess the impact of three key

components by progressively adding them. Setting (A) applies pretraining for the feature extractor and classifier by jointly learning with the other datasets (PT), Setting (B) extends this by introducing meta-learning for the dataset-irrelevant feature extractor (ML), and Setting (C) further includes decoupled training for the dataset-specific classifiers (DT). Statistical significance compared to Setting (C) is determined using the Wilcoxon signed-rank test [30].

In the sparsely labeled (10%) condition, the baseline model demonstrates strong specificity but suffers from low sensitivity, particularly in the WHII-STD dataset, where it completely fails to detect positive samples (0%) due to the lack of labeled data. Setting (A) partially mitigates this issue, improving sensitivity (SEN = 93.95%) in the WHII-STD dataset. Setting (B) further enhances robustness, achieving the highest sensitivity (SEN = 91.08%) in the BCP dataset. Setting (C) shows the outstanding performance, achieving the highest accuracy across datasets. Notably, in the HCP dataset, our proposed framework achieves 98.93% accuracy, 97.82% sensitivity, and 98.42% G-measure, surpassing all other settings.

In the fully labeled (100%) condition, the baseline model achieves high specificity. However, its sensitivity remains relatively low. Setting (A) results in slight improvements in sensitivity, but jointly training with the other datasets does not always enhance performance due to variations in noise characteristics and inconsistencies in labeling criteria. For instance, in the HCP dataset, Setting (A) slightly reduces accuracy (99.12%) compared to the baseline (99.31%). Setting (B) further improves sensitivity and F1-score across most datasets, particularly in the BCP dataset, where sensitivity increases to 92.58% and F1-score increases to 91.04%. Setting (C) achieves the highest accuracy, sensitivity, F1-score, and G-measure across most datasets, with the notable improvements observed in the WHII-STD dataset (SEN = 94.14%, GM = 96.60%).

In several cases, our proposed framework in the 10% labeled condition achieves performance exceeding that of the baseline in the 100% labeled condition. In the WHII-MB6 dataset, Setting (C) in the 10% labeled condition achieves 95.60% F1-score and 98.13% G-measure, outperforming 95.18% F1-score and 96.77% G-measure of the baseline in the 100% labeled condition. These results demonstrate that our proposed framework achieves outstanding performance compared to previous approaches and effectively adapts to sparsely labeled datasets.

3.3 Visual Evaluation

To further validate the effectiveness of our proposed framework in fMRI denoising, we conduct a qualitative comparison between our proposed framework and the baseline model [19] on a sample from the HCP dataset under the 10% labeled condition, as shown in Fig. 2. The visualization includes three input modalities: the IC spatial map, its corresponding IC time series, and the wavelet-transformed image, which is generated from the IC time series using a learnable wavelet kernel [19]. Additionally, we generate explanation maps using Grad-CAM [25] to highlight the regions that contribute most to classification decisions.

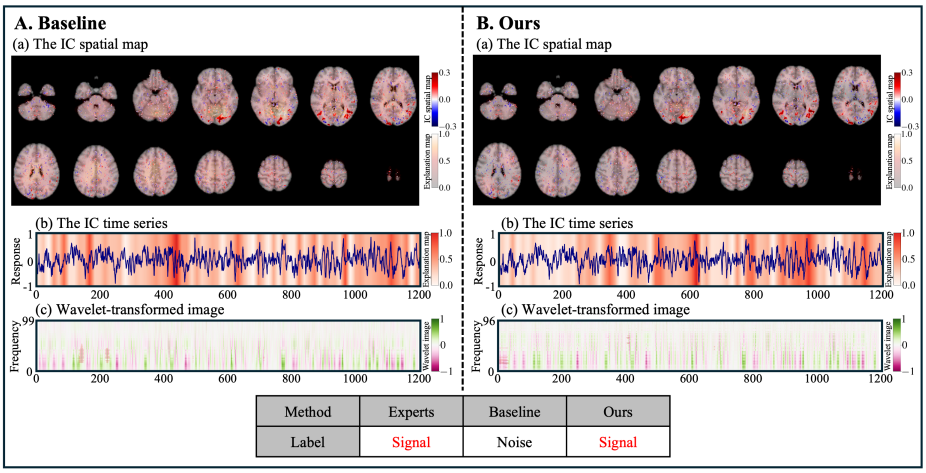

Fig. 2. Visual evaluation comparing the baseline model [19] and our proposed framework on a sample from the HCP dataset under the 10% labeled condition. The visualization includes three modalities: (a) IC spatial map, (b) IC time series, and (c) wavelet-transformed image. The bottom table presents classification results alongside expert annotations. Additionally, explanation maps are generated using Grad-CAM [25] to highlight the regions most influential in the model's classification decisions.

The IC spatial maps for both the baseline model and our proposed framework (ours) capture similar positive clusters in the primary visual cortex of the brain caused by neuronal activity, confirming that both models can detect significant fMRI signal patterns. However, the baseline model incorrectly classifies a signal IC as a noise IC, whereas our proposed framework correctly identifies it as a signal IC. This suggests that our proposed framework effectively adapts to a sparsely labeled dataset.

4 Conclusion

In this study, we propose a meta-learning-based semi-supervised domain adaptation framework for fMRI denoising, enabling the learning of dataset-irrelevant features from sparsely labeled datasets by leveraging existing labeled datasets. To address inter-dataset variations and inconsistencies in labeling criteria, we adopt meta-learning to extract dataset-irrelevant features and employ a decoupled training strategy to optimize dataset-specific classifiers. Experimental results on four fMRI datasets demonstrate that our proposed framework achieves outstanding performance in both sparsely labeled and fully labeled conditions.

Acknowledgments. This work was supported by grants from the Institute of Information & Communications Technology Planning & Evaluation (IITP)—specifically,

the Artificial Intelligence Graduate School Program at Korea University (No. RS-2019-II190079), the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2023-00212498, RS-2024-00415812), the MSIT under the ITRC (Information Technology Research Center) support program (IITP-2025-RS-2024-00436857) supervised by the IITP, and IITP grant funded by MSIT (No. RS-2024-00457882, Artificial Intelligence Research Hub Project).

Disclosure of Interests. The authors declare no competing interests.

References

1. Baxter, L., Fitzgibbon, S., Moultrie, F., Goksan, S., Jenkinson, M., Smith, S., Andersson, J., Duff, E., Slater, R.: Optimising neonatal fMRI data analysis: Design and validation of an extended dHCP preprocessing pipeline to characterise noxious-evoked brain activity in infants. NeuroImage **186**, 286–300 (2019)
2. Bianciardi, M., Fukunaga, M., van Gelderen, P., Horovitz, S.G., de Zwart, J.A., Shmueli, K., Duyn, J.H.: Sources of functional magnetic resonance imaging signal fluctuations in the human brain at rest: a 7T study. Magnetic Resonance Imaging **27**(8), 1019–1029 (2009)
3. Caballero-Gaudes, C., Reynolds, R.C.: Methods for cleaning the BOLD fMRI signal. NeuroImage **154**, 128–149 (2017)
4. Filippini, N., Zsoldos, E., Haapakoski, R., Sexton, C.E., Mahmood, A., Allan, C.L., Topiwala, A., Valkanova, V., Brunner, E.J., Shipley, M.J., et al.: Study protocol: the Whitehall II imaging sub-study. BMC Psychiatry **14**(1), 1–16 (2014)
5. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: International Conference on Machine Learning. pp. 1126–1135. PMLR (2017)
6. Grefenstette, E., Amos, B., Yarats, D., Htut, P.M., Molchanov, A., Meier, F., Kiela, D., Cho, K., Chintala, S.: Generalized inner loop meta-learning. arXiv preprint arXiv:1910.01727 (2019)
7. Griffanti, L., Douaud, G., Bijsterbosch, J., Evangelisti, S., Alfaro-Almagro, F., Glasser, M.F., Duff, E.P., Fitzgibbon, S., Westphal, R., Carone, D., et al.: Hand classification of fMRI ICA noise components. NeuroImage **154**, 188–205 (2017)
8. Griffanti, L., Salimi-Khorshidi, G., Beckmann, C.F., Auerbach, E.J., Douaud, G., Sexton, C.E., Zsoldos, E., Ebmeier, K.P., Filippini, N., Mackay, C.E., et al.: ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging. NeuroImage **95**, 232–247 (2014)
9. Han, J.W., Bak, S., Kim, J.M., Choi, W., Shin, D.H., Son, Y.H., Kam, T.E.: Meta-eeg: Meta-learning-based class-relevant eeg representation learning for zero-calibration brain–computer interfaces. Expert Systems with Applications **238**, 121986 (2024)
10. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Conference on Computer Vision and Pattern Recognition. pp. 770–778 (2016)
11. Heo, K.S., Shin, D.H., Hung, S.C., Lin, W., Shen, D., Zhang, H., Kam, T.E.: Deep attentive spatio-temporal feature learning for automatic fMRI denoising. NeuroImage (2022)
12. Howell, B.R., Styner, M.A., Gao, W., Yap, P.T., Wang, L., Baluyot, K., Yacoub, E., Chen, G., Potts, T., Salzwedel, A., et al.: The UNC/UMN Baby Connectome Project (BCP): An overview of the study design and protocol development. NeuroImage **185**, 891–905 (2019)

13. Jeon, E., Kang, E., Lee, J., Lee, J., Kam, T.E., Suk, H.I.: Enriched representation learning in resting-state fMRI for early MCI diagnosis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 397–406. Springer (2020)
14. Kam, T.E., Wen, X., Jin, B., Jiao, Z., Hsu, L.M., Zhou, Z., Liu, Y., Yamashita, K., Hung, S.C., Lin, W., Zhang, H., Shen, D.: A deep learning framework for noise component detection from resting-state functional MRI. In: Medical Image Computing and Computer Assisted Intervention. pp. 754–762 (2019)
15. Kam, T.E., Zhang, H., Jiao, Z., Shen, D.: Deep learning of static and dynamic brain functional networks for early MCI detection. IEEE Transactions on Medical Imaging **39**(2), 478–487 (2020)
16. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. Advances in Neural Information Processing Systems **33**, 18661–18673 (2020)
17. Kundu, P., Voon, V., Balchandani, P., Lombardo, M.V., Poser, B.A., Bandettini, P.A.: Multi-echo fMRI: a review of applications in fMRI denoising and analysis of BOLD signals. NeuroImage **154**, 59–80 (2017)
18. Lalkhen, A.G., McCluskey, A.: Clinical tests: sensitivity and specificity. Continuing Education in Anaesthesia Critical Care & Pain **8**(6), 221–223 (2008)
19. Lim, M., Heo, K.S., Kim, J.M., Kang, B., Lin, W., Zhang, H., Shen, D., Kam, T.E.: A unified multi-modality fusion framework for deep spatio-spectral-temporal feature learning in resting-state fMRI denoising. IEEE Journal of Biomedical and Health Informatics **28**(4), 2067–2078 (2024)
20. Murphy, K., Birn, R.M., Bandettini, P.A.: Resting-state fMRI confounds and cleanup. NeuroImage **80**, 349–359 (2013)
21. Power, J.D., Mitra, A., Laumann, T.O., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E.: Methods to detect, characterize, and remove motion artifact in resting state fMRI. NeuroImage **84**, 320–341 (2014)
22. Pruim, R.H., Mennes, M., van Rooij, D., Llera, A., Buitelaar, J.K., Beckmann, C.F.: ICA-AROMA: A robust ICA-based strategy for removing motion artifacts from fMRI data. NeuroImage **112**, 267–277 (2015)
23. Saito, K., Kim, D., Sclaroff, S., Darrell, T., Saenko, K.: Semi-supervised domain adaptation via minimax entropy. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8050–8058 (2019)
24. Salimi-Khorshidi, G., Douaud, G., Beckmann, C.F., Glasser, M.F., Griffanti, L., Smith, S.M.: Automatic denoising of functional MRI data: combining independent component analysis and hierarchical fusion of classifiers. NeuroImage **90**, 449–468 (2014)
25. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-CAM: Visual explanations from deep networks via gradient-based localization. In: International Conference on Computer Vision. pp. 618–626 (2017)
26. Smith, S.M., Beckmann, C.F., Andersson, J., Auerbach, E.J., Bijsterbosch, J., Douaud, G., Duff, E., Feinberg, D.A., Griffanti, L., Harms, M.P., et al.: Resting-state fMRI in the human connectome project. NeuroImage **80**, 144–168 (2013)
27. Wachinger, C., Rieckmann, A., Pölsterl, S., Initiative, A.D.N., et al.: Detect and correct bias in multi-site neuroimaging datasets. Medical Image Analysis **67**, 101879 (2021)
28. Wang, M.L., Zhang, D.Q., Huang, J.S., Yap, P.T., Shen, D.G., Liu, M.X.: Identifying autism spectrum disorder with multi-site fMRI via low-rank domain adaptation. IEEE Transactions on Medical Imaging **39**(3), 644–655 (2020)

29. Wang, X., Pan, J., Guo, X., Liu, D., Jiang, J.: Decoupled training: Return of frustratingly easy multi-domain learning. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 15644–15652 (2024)
30. Wilcoxon, F.: Individual comparisons by ranking methods. Biometrics Bulletin **1**(6), 80–83 (1945)
31. Yu, M., Linn, K.A., Cook, P.A., Phillips, M.L., McInnis, M., Fava, M., Trivedi, M.H., Weissman, M.M., Shinohara, R.T., Sheline, Y.I.: Statistical harmonization corrects site effects in functional connectivity measurements from multi-site fMRI data. Human Brain Mapping **39**(11), 4213–4227 (2018)