

Foundation-Model-Boosted Multimodal Learning for fMRI-based Neuropathic Pain Drug Response Prediction

Wenrui Fan^{1,2}[0009-0007-1394-0092], L. M. Riza Rizky^{1,2}[0009-0003-3073-9789], Jiayang Zhang^{1,2}[0009-0007-8446-2462], Chen Chen^{1,2,6}[0000-0002-3525-9755], Haiping Lu^{1,2}[0000-0002-0349-2181], Kevin Teh³[0000-0003-2538-5157], Dinesh Selvarajah^{4,5}[0000-0001-7426-1105], and Shuo Zhou^{1,2}(✉)[0000-0002-8069-2814]

¹ School of Computer Science, University of Sheffield, Sheffield, UK

² Centre for Machine Intelligence, University of Sheffield, Sheffield, UK

³ Academic Unit of Radiology, University of Sheffield, Sheffield, UK

⁴ Diabetes Research Unit, Royal Hallamshire Hospital, Sheffield, UK

⁵ Division of Clinical Medicine, University of Sheffield, Sheffield, UK

⁶ Department of Computing, Imperial College London, London, UK

{wenrui.fan, l.m.rizky, jiayang.zhang, chen.chen2, h.lu, k.teh, d.selvarajah, shuo.zhou(✉)}@sheffield.ac.uk

Abstract. Neuropathic pain, affecting up to 10% of adults, remains difficult to treat due to limited therapeutic efficacy and tolerability. Although resting-state functional MRI (rs-fMRI) is a promising non-invasive measurement of brain biomarkers to predict drug response in therapeutic development, the complexity of fMRI demands machine learning models with substantial capacity. However, *extreme data scarcity* in neuropathic pain research limits the application of high-capacity models. To address the challenge of data scarcity, we propose **FMM_{TC}**, a **Foundation-Model-boosted Multimodal** learning framework for fMRI-based neuropathic pain drug response prediction, which leverages both internal multimodal information in *pain-specific* data and external knowledge from large *pain-agnostic* data. Specifically, to maximize the value of limited pain-specific data, FMM_{TC} integrates complementary information from two rs-fMRI modalities: **T**ime series and functional **C**onnectivity. FMM_{TC} is further boosted by an fMRI foundation model with its external knowledge from extensive pain-agnostic fMRI datasets enriching limited pain-specific information. Evaluations with an in-house dataset and a public dataset from OpenNeuro demonstrate FMM_{TC}'s superior representation ability, generalizability, and cross-dataset adaptability over existing unimodal fMRI models that only consider one of the rs-fMRI modalities. The ablation study validates the effectiveness of multimodal learning and foundation-model-powered external knowledge transfer in FMM_{TC}. An integrated gradient-based interpretation study explains how FMM_{TC}'s cross-dataset dynamic behaviors enhance its adaptability. In conclusion, FMM_{TC} boosts clinical trials in neuropathic pain therapeutic development by accurately predicting drug responses to improve the participant stratification efficiency. The code is released on https://github.com/Shef-AIRE/FMM_TC.

Keywords: Multimodal learning · Foundation model · fMRI · Drug response prediction · Neuropathic pain

1 Introduction

Neuropathic pain, caused by damage or disease in the somatosensory nervous system [14], affects up to 10% of people over 30 years old [1]. Despite advances in pharmacological treatments, nearly half of patients fail to achieve adequate pain relief [8] due to low drug efficacy [6] and high rates of side effects [9]. This underscores the urgent need for novel therapeutic approaches [21].

Although resting-state functional magnetic resonance imaging (rs-fMRI) is a promising non-invasive measurement of brain biomarkers [21,3], its high dimensionality [17], noisiness [5], and cross-individual variability [28] demand machine learning (ML) models with substantial capacity. To capture complex patterns in fMRI data, previous works have proposed advanced deep neural networks with strong representation ability (e.g. transformers [15] and Mamba-based models [31]) that perform well with large-scale training data such as UK Biobank [22].

However, directly applying such cutting-edge machine learning models for neuropathic pain fMRI analysis faces a significant challenge of *extreme data scarcity*, leading to severe overfitting or training collapse problems. In fact, neuropathic pain datasets usually have less than 100 participants [24] due to the challenges in the data acquisition process (e.g. high costs, ethics). Therefore, practical applications of advanced machine learning models with sophisticated network architectures for neuropathic pain fMRI analysis are hindered.

To tackle this challenge, we propose **FMM_{TC}**, a **F**oundation-**M**odel-boosted **M**ultimodal learning framework for fMRI-based drug response prediction in neuropathic pain therapeutic development. As shown in Fig. 1, FMM_{TC} proposes two novel modules to address data scarcity: a) multimodal learning across two modalities of rs-fMRI: **T**ime series (TS) and functional **C**onnectivity (FC), and b) transferring external knowledge from large-scale pain-agnostic fMRI datasets via foundation models (FMs). Our contributions are three-fold:

Firstly, we introduce multimodal learning in FMM_{TC} to exploit internal information in the limited pain-specific data by integrating complementary information from TS and FC modalities. Although previous works [15,16,4] regard FC and TS as one modality due to coming from the same source, FMM_{TC} rethinks them as two distinct modalities, encodes them in two separate streams, and fuses their latent features by a simple yet effective design to enhance predictions. By multimodal learning, FMM_{TC} improves Matthews correlation coefficient (MCC) by at least 2.80% compared to unimodal baselines across all experiments.

Secondly, FMM_{TC} exploits the benefits of external knowledge that FM [4] learns from large-scale pain-agnostic fMRI datasets (e.g. UK Biobank [22] with 50,000+ scans). Enriching limited internal information with rich knowledge from the general population [22] beyond pain patients, this FM-powered external knowledge transfer improves MCC by at least 14.71% compared to the model with the same network architecture but no pre-training or external knowledge.

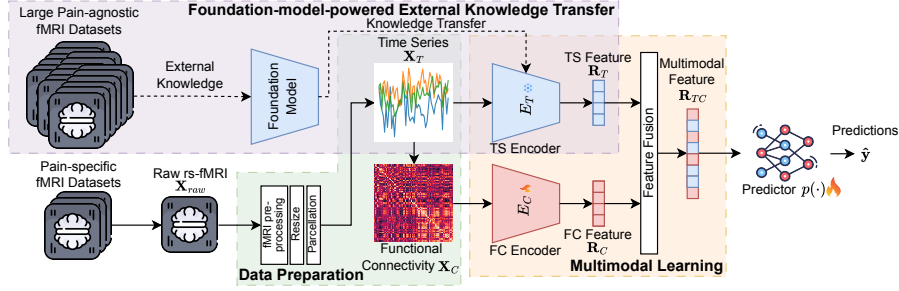


Fig. 1. FMM_{TC} pipeline. The raw rs-fMRI \mathbf{X}_{raw} undergoes three processing steps to obtain time-series (TS) \mathbf{X}_T . Then, functional connectivity (FC) \mathbf{X}_C is computed from \mathbf{X}_T . Next, \mathbf{X}_T and \mathbf{X}_C are processed by a time-series (TS) encoder E_T (frozen fMRI foundation model) and a functional connectivity (FC) encoder E_C (learnable CNN), respectively. Finally, features of two modalities \mathbf{R}_T and \mathbf{R}_C are fused to get the multimodal feature \mathbf{R}_{TC} for the following prediction.

Thirdly, we comprehensively evaluate FMM_{TC} with one public dataset from OpenNeuro [24] and an in-house dataset. Drug-agnostic and within-domain drug-specific response prediction demonstrate FMM_{TC}’s promising representation ability and cross-dataset adaptability (up to MCC=69.64% and AUROC=77.56%). The out-of-domain drug-specific response prediction results show FMM_{TC}’s strong generalizability (up to MCC=86.02%). We conduct an ablation study to demonstrate the effectiveness of multimodal learning and external knowledge transfer, and an interpretation study to understand the reason for its strong adaptability.

2 Methodology

Problem formulation. The drug response prediction task is formulated as a binary classification problem. With a dataset of rs-fMRI images $\{\mathbf{X}_{raw}\}$ from a cohort of participants, the goal is to learn a function f_{θ} parameterized by θ to predict participant’s drug response $\hat{y} = f_{\theta}(\mathbf{X}_{raw})$ with label $y=1$ indicating effective (responder) and 0 indicating no response (non-responder).

Overview of FMM_{TC}. FMM_{TC} has three major modules: a) data preparation, b) multimodal learning, and c) FM-powered external knowledge transfer.

Data preparation. FMM_{TC} requires dual stream and multimodal inputs: TS data \mathbf{X}_T and FC data \mathbf{X}_C . To get TS data \mathbf{X}_T , we first process the raw fMRI \mathbf{X}_{raw} by 3 steps: (1) fMRI pre-processing: \mathbf{X}_{raw} is processed by fMRI pre-processing pipelines (e.g. fMRIPrep [7]) with steps like head motion correction. (2) Resize: The processed fMRI images are resized to match the standard ICBM152 space [10]. (3) Parcellation: we parcellize the resized fMRI data into N regions of interest (ROIs) by an atlas to obtain TS data $\mathbf{X}_T \in \mathbb{R}^{N \times t}$, where t refers to time steps. In this context, an atlas is a neuroscientist-curated annotation that clusters brain voxels based on their functional similarity, and each ROI is a group of neurons that perform similar functions. After getting \mathbf{X}_T , we further

compute FC data \mathbf{X}_C by the correlation [2] between TS data \mathbf{X}_T and itself along the time dimension $\mathbf{X}_C = \text{Cor}(\mathbf{X}_T, \mathbf{X}_T)$, $\mathbf{X}_C \in \mathbb{R}^{N \times N}$.

Multimodal learning. A key contribution to this work is that we rethink TS and FC as two distinct modalities instead of two features from one modality. We use two separate encoders: a time-series (TS) encoder E_T and a functional connectivity (FC) encoder E_C to process the multimodal input, which gives us TS and FC features $\mathbf{R}_T = E_T(\mathbf{X}_T) \in \mathbb{R}^M$ and $\mathbf{R}_C = E_C(\mathbf{X}_C) \in \mathbb{R}^M$, respectively. Here, M is the size of output features. To fuse \mathbf{R}_T and \mathbf{R}_C and obtain multimodal feature \mathbf{R}_{TC} , a feature fusion scheme $\text{fusion}(\cdot)$ is introduced: $\mathbf{R}_{TC} = \text{fusion}(\mathbf{R}_T, \mathbf{R}_C)$.

We regard TS and FC as two distinct modalities for two reasons: (1) Their information is different and complementary. TS focuses on fine-grained, regional features and provides more temporal information [27], while FC constructs a global brain network emphasizing spatial relationships between regions [11]. (2) For ML models, the global patterns in FC might not be directly learnable from TS data, especially with such small datasets [25]. In this situation, FC provides an explicit and complementary global summary of the information in TS data. To exploit this complementarity, FMM_{TC} treats them as two distinct modalities.

FM-powered external knowledge transfer. To enrich the limited information in pain-specific data, we use FM to introduce external knowledge. FM learns rich knowledge by pre-training on large fMRI datasets, with which it can extract robust representations from data and generalize effectively with small, domain-specific pain datasets [4]. FMM_{TC} leverages BrainLM [4] here, which is a foundation model for TS data with rich knowledge obtained from 6,700 hours of fMRI data from 50k+ UK Biobank participants [22].

Detailed design of FMM_{TC} . We now discuss the designs of encoders E_T , E_C and the feature fusion method $\text{fusion}(\cdot)$, and the motivations for design decisions.

BrainLM as TS encoder and rationale. we use BrainLM’s encoder [4] to implement E_T , which is a 4-layer 4-head transformer with an efficient self-attention mechanism to reduce computational complexity while maintaining global context [32]. We use the class token of the output of BrainLM to get \mathbf{R}_T . BrainLM is frozen during training to stabilize training and avoid mode collapse when fine-tuned on small, domain-specific pain datasets. Traditional knowledge distillation [33] is not applied in our study due to a) limited pain-specific training data, and b) restricted access to BrainLM’s original pre-training dataset: UK Biobank [22].

ResNet-18 as FC encoder and rationale. We simply implement E_C with the backbone of ResNet-18 E_{ResNet} [13] alongside a linear projection layer h afterward to align E_{ResNet} ’s output dimensions with \mathbf{R}_T . Thus, $E_C = h \circ E_{ResNet}$ and we train it from scratch. The choice of using ResNet for E_C is based on three reasons: (1) FC matrices capture spatial relationships between brain regions, sharing structural similarities with 2D images. (2) While more advanced models like transformers [15] work well in large-scale fMRI analysis, they deliver unsatisfactory performances with limited pain-specific data. (3) CNN is proven effective for FC data [16,30]. By applying ResNet as E_C , we balance simplicity and performance under severe data scarcity.

Table 1. Simple feature concatenation outperforms advanced feature fusion schemes.

Method	Unidirectional cross-attention [19]	Bidirectional cross-attention [19]	Mixture-of-experts [12]	Element-wise summation	Concatenation (ours)
MCC (%)	0	44.12	0	53.93	69.64

Table 2. Metadata of the in-house dataset and the OpenNeuro [24] dataset.

Dataset	No.	Mean age	Gender	Given drugs	Response labels
In-house	61	58.3 (40-82)	Male(38), Female(23)	Lidocaine(61)	$y = 1(24), y = 0(37)$
OpenNeuro [24]	56	57.9 (44-73)	Male(26), Female(30)	Duloxetine(19), Placebo(37)	$y = 1(26), y = 0(30)$

Multimodal feature fusion via concatenation and rationale. We apply concatenation ($\text{concat}(\cdot)$), a simple yet effective method to get the fused multimodal feature \mathbf{R}_{TC} by $\mathbf{R}_{TC} = \text{concat}(\mathbf{R}_T, \mathbf{R}_C)$, $\mathbf{R}_{TC} \in \mathbb{R}^{2M}$. In our preliminary experiments (see Table 1), we found advanced multimodal feature fusion techniques including attention-based [19] and mixture-of-experts (MoE) [12] delivered unsatisfactory performance, likely due to limited training data. Element-wise summation’s sub-optimal performance could be incompatible element-wise semantics in features from the transformer (E_T) and ResNet (E_C). Concatenation preserves individual semantics in \mathbf{R}_T and \mathbf{R}_C , allowing the following layers to dynamically adjust feature prioritization.

Model training. FMM_{TC} is trained end-to-end. After feature fusion, a single-layer perceptron classifier $p(\cdot)$ is attached for the downstream drug response prediction: $\hat{y} = p(\text{concat}(E_T(\mathbf{X}_T), E_C(\mathbf{X}_C)))$. The training is driven by a binary cross-entropy loss between predictions $\hat{\mathbf{y}}$ and labels \mathbf{y} : $\mathcal{L}_{BCE}(\hat{\mathbf{y}}, \mathbf{y})$.

3 Experiments

3.1 Experiment Setup

Datasets. The metadata of the in-house and OpenNeuro [24] datasets are shown in Table 2. The OpenNeuro dataset is the patient cohort of the DS000208 dataset from the OpenNeuro platform [24]. The in-house data are collected from the local hospital and de-identified. Written informed consent for the in-house data was obtained before subjects participated in the study, which had prior approval by the Sheffield Local Research Ethics Committee (Sheffield, U.K.). fMRI scans in both datasets are pre-treatment scans.

Baselines. We compare a) FMM_{TC} as an end-to-end predictive model against four TS baselines (LSTM [18], ResNet [13], BrainLM [4] w/wo pre-training) and two FC baselines (ResNet [13], BNT [15]), and b) FMM_{TC}-extracted feature \mathbf{R}_{TC} for linear machine learning classifiers against raw TS/FC data ($\mathbf{X}_T, \mathbf{X}_C$), PCA-extracted features (PCA(\mathbf{X}_T), PCA(\mathbf{X}_C)), and BrainLM-derived features (BrainLM(\mathbf{X}_T)) [4]. We also provide a random guessing experiment simulating conventional random participant recruitment in clinical trials. Random guessing is repeated 1,000 times. Other reported results are under 5-fold cross-validation.

Table 3. FMM_{TC} vs. unimodal baselines on drug-agnostic response prediction. The results of the ablation study are also shown by comparing Exp. 5, 7, 8, 9.

Dataset	Exp. ID	Model	Input modality	Pre-training	Metrics (%)			
					F1↑	BACC↑	AUROC↑	MCC↑
OpenNeuro [24]	1	Random	N/A	✗	49.57±11.39	50.64±11.49	50.64±11.49	1.29±22.96
	2	LSTM [18]	TS	✗	63.95±11.94	67.33±8.38	51.67±12.80	43.68±16.89
	3	BNT [15]	FC	✗	34.90±0.88	50.00±0.00	51.00±18.88	0.00±0.00
	4	ResNet [13]	TS	✗	63.36±9.21	67.00±7.58	50.00±14.91	42.79±12.90
	5	ResNet [13]	FC	✗	66.79±12.64	69.33±10.97	55.11±29.58	44.23±20.91
	6	BrainLM [4]	TS	✗	34.90±0.88	50.00±0.00	44.56±16.50	0.00±0.00
	7	BrainLM [4]	TS	✓	<u>79.38±8.77</u>	<u>80.00±8.16</u>	<u>71.44±15.28</u>	<u>63.36±14.47</u>
	8	FMM _{TC}	TS+FC	✗	70.18±14.04	73.67±11.87	60.67±25.43	54.93±18.50
	9	FMM _{TC}	TS+FC	✓	80.51±16.11	82.33±14.12	77.56±26.98	69.64±23.32
In-house	1	Random	N/A	✗	47.96±11.61	49.12±11.66	49.12±11.66	-1.77±23.35
	2	LSTM [18]	TS	✗	37.72±1.38	50.00±0.00	64.84±37.41	0.00±0.00
	3	BNT [15]	FC	✗	34.90±0.88	50.00±0.00	51.00±18.88	0.00±0.00
	4	ResNet [13]	TS	✗	66.98±12.74	68.67±11.63	59.67±15.20	41.86±22.42
	5	ResNet [13]	FC	✗	75.68±9.33	75.64±9.27	68.96±11.74	<u>60.92±12.58</u>
	6	BrainLM [4]	TS	✗	34.90±0.88	50.00±0.00	44.56±16.50	0.00±0.00
	7	BrainLM [4]	TS	✓	74.83±16.81	75.71±15.32	73.88±24.06	59.01±26.53
	8	FMM _{TC}	TS+FC	✗	66.92±21.59	72.14±16.95	71.43±21.29	48.32±32.64
	9	FMM _{TC}	TS+FC	✓	80.32±13.89	81.07±13.87	75.20±26.14	63.72±26.77

Implementation of FMM_{TC}. We use fMRIPrep [7] in fMRI pre-processing. We use A424 [23] atlas in parcellation following BrainLM [4]. The number of ROIs N is 424. \mathbf{X}_T is truncated or padded to a fixed length of 200 and patchified with a patch size of 20. The feature dimension M is 256.

Metrics. We report four metrics: F1, balanced accuracy (BACC), AUROC, and Matthew correlation coefficient (MCC). Because MCC fairly measures TP, TN, FP, and FN, we use MCC as the primary metric to select the best models. The bests are **bolded** and the second bests are underlined.

Experiment design of drug response prediction. We perform drug-agnostic and drug-specific response prediction in experiments. *Drug-agnostic response prediction* is to predict whether a patient is a responder regardless of the types of drugs and can be mathematically described as $\hat{y} = f_{\theta}(\mathbf{X}_{raw})$. It is the first step in clinical trials, reducing trial costs and time by focusing on likely responders [26]. *Drug-specific response prediction* is to predict the responses for specific drugs and can be described as $\hat{y} = f_{\theta}(\mathbf{X}_{raw}|d)$, where d is the type of given drugs. It can help clinicians select participants that are more targeted and fine-grained.

3.2 Experiment Results

Superiority of FMM_{TC} in end-to-end prediction. We evaluate FMM_{TC} on *drug-agnostic response prediction* against: a) four baselines for TS data (LSTM [18], ResNet [13], BrainLM [4] w/wo pre-training) and b) two baselines for FC data (ResNet [13], BNT [15]) on OpenNeuro and in-house datasets (see Table 3). BrainLM without pre-training (Exp. 6) and BNT (Exp. 3) struggle with limited training data, highlighting challenges in training advanced ML with small fMRI datasets. Other unimodal baselines (Exp. 2, 4, 5, 7) show inconsistent results

Table 4. We compare the representation ability (within-domain) and generalizability (out-of-domain) of FMM_{TC} and BrainLM on drug-specific response prediction using the OpenNeuro dataset. Given BrainLM’s expected generalizability as a foundation model and its rank as the second-best model on the OpenNeuro dataset in Table 3, we focus our comparison on FMM_{TC} and BrainLM.

Setup	Training drug	Testing drug	Model	Metrics (%)			
				F1↑	BACC↑	AUROC↑	MCC↑
Within-domain	Duloxetine	Duloxetine	BrainLM [4]	81.33±29.21	85.00±22.36	90.00±22.36	71.55±43.98
		Duloxetine	FMM _{TC}	84.00±14.60	86.67±12.63	85.00±23.36	74.64±23.14
	Placebo	Placebo	BrainLM [4]	77.84±12.87	79.17±12.84	78.33±21.17	60.37±25.68
		Placebo	FMM _{TC}	88.95±6.29	89.17±6.32	89.17±9.93	80.59±11.00
Out-of-domain	Duloxetine	Placebo	BrainLM [4]	55.40±16.43	58.19±13.43	71.36±9.21	16.67±27.20
		Placebo	FMM _{TC}	92.30±4.97	92.27±5.41	91.59±7.73	86.02±8.91
	Placebo	Duloxetine	BrainLM [4]	52.58±6.00	56.61±4.91	66.02±2.32	17.26±13.51
		Duloxetine	FMM _{TC}	83.24±14.93	84.59±12.20	89.77±9.93	71.85±19.25

across datasets. FMM_{TC} consistently outperforms all baselines, demonstrating superior prediction performance and cross-dataset adaptability.

We evaluate FMM_{TC} on *drug-specific response prediction* against BrainLM [4] (see Table 4). We divide the OpenNeuro dataset into two subsets according to the given drugs. We fine-tune BrainLM [4] and train FMM_{TC} on one subset and test them on both subsets. Results of *within-domain* tests show that FMM_{TC} outperforms BrainLM when training and testing on the same drug, illustrating FMM_{TC}’s superior representation ability. FMM_{TC}’s strong generalizability is demonstrated by *out-of-domain* evaluation: when training with one drug and transferring to another, FMM_{TC} outperforms BrainLM by a significant margin.

Effectiveness of multimodal feature \mathbf{R}_{TC} for linear classifiers. We evaluate FMM_{TC}’s representation ability as a feature extractor for linear machine learning classifiers on *drug-agnostic response prediction* with the OpenNeuro dataset. We compare multimodal feature \mathbf{R}_{TC} from FMM_{TC} against unimodal features from raw data, PCA, and BrainLM [4] with four commonly used linear classifiers [20]: SVM, Ridge, k-NN, and XGBoost. The results in Fig. 2(a) show that statistical machine learning approaches struggle to capture the complex patterns in fMRI with their limited capacity. In contrast, FMM_{TC} outperforms all other feature extractors by a significant margin, indicating the effectiveness of multimodal feature \mathbf{R}_{TC} and the strong representation ability of FMM_{TC}.

Ablation study. We now demonstrate the effectiveness of multimodal learning and FM-powered external knowledge transfer (see Exp. 5, 7, 8, 9 in Table 3).

Multimodal learning improves MCC by at least **2.80%**. The effectiveness of multimodal learning is demonstrated by comparing Exp. 5, 7, and 9 in Table 3, where ResNet with FC and BrainLM with TS serve as two separate streams of FMM_{TC}. The results indicate that the performance of ResNet (Exp. 5) and BrainLM (Exp. 7) varies across datasets, whereas FMM_{TC} (Exp. 9) consistently outperforms them, achieving strong and stable performance across datasets. This highlights the superiority and cross-dataset adaptability of multimodal learning. *FM-powered external knowledge transfer* improves MCC by at least **14.71%**. We evaluate the effectiveness of external knowledge in FMM_{TC} by comparing

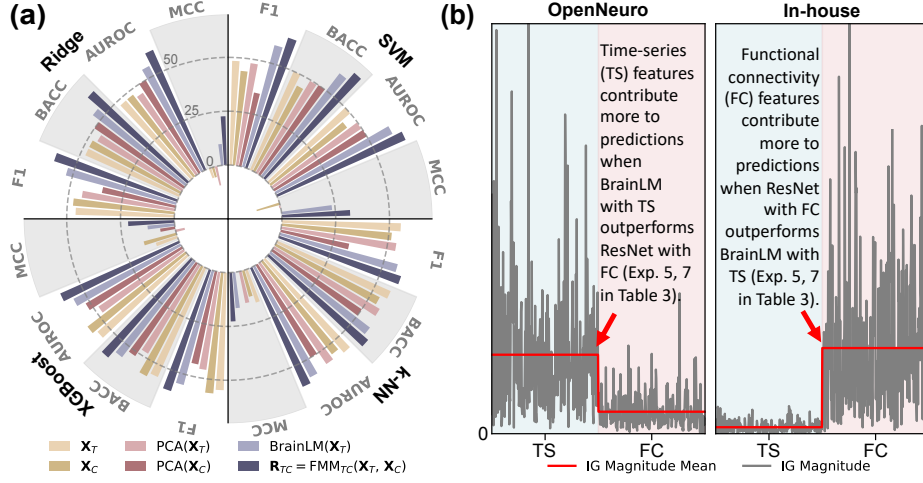


Fig. 2. (a) Multimodal feature R_{TC} from FMM_{TC} outperforms unimodal features from other feature extractors with linear classifiers on drug-agnostic response prediction. (b) Feature importance via integrated gradients (IG) illustrates how FMM_{TC} ’s cross-dataset dynamic behaviors enhance adaptability: FMM_{TC} flexibly prioritizes the most effective modality in predictions. The light blue background is IG values for time-series feature R_T and the light red region is for functional connectivity feature R_C .

FMM_{TC} w/o pre-training (Exp. 8, 9 in Table 3). The performance gap between the two models (69.64%/63.72% vs. 54.93%/48.32% on MCC) confirms the importance of external knowledge and pre-training to address data scarcity.

Interpretation to FMM_{TC} ’s cross-dataset adaptability. The high adaptability of FMM_{TC} is due to its flexible dynamic behaviors. We use integrated gradients (IG) [29] to explain how FMM_{TC} ’s flexible reliance on multimodal features enhances its adaptability across datasets. A larger IG magnitude signifies a higher impact of input features on the model’s predictions. We present IG magnitudes for the final predictor of FMM_{TC} , which takes the fused multimodal features R_{TC} as inputs and outputs predictions \hat{y} , shown in Fig. 2(b).

On the OpenNeuro dataset, TS features contribute more to the predictions than FC features, while FC features have a stronger impact on the in-house dataset. This observation aligns with the results of Exp. 5 and 7 in Table 3. When BrainLM with TS outperforms ResNet with FC on the OpenNeuro dataset, FMM_{TC} relies more on the TS features and capitalizes on the strengths of TS data. In the in-house dataset, ResNet with FC performs better than BrainLM with TS, causing FMM_{TC} to emphasize FC features in prediction.

These findings highlight FMM_{TC} ’s cross-dataset dynamic behaviors and explain FMM_{TC} ’s high adaptability. By combining TS and FC modalities, FMM_{TC} can dynamically adjust its reliance on various modalities based on the dataset and task. This flexibility enables FMM_{TC} to adapt effectively across diverse conditions and data distributions, improving performance in various scenarios.

4 Conclusion

To address the severe data scarcity in fMRI-based neuropathic pain drug response prediction, we propose FMM_{TC}, a foundation-model-boosted multimodal learning framework with multimodal learning and foundation-model-powered external knowledge transfer. FMM_{TC} shows strong representation ability, generalization, and cross-dataset adaptability, enabling FMM_{TC} with abilities to predict treatment effects and enhance stratification in the participant recruitment process, and with the potential to significantly reduce the time and money costs associated with random methods in traditional clinical trials.

Acknowledgments. This study is funded by the Engineering and Physical Sciences Research Council (EP/Y017544/1) for the "A Novel Artificial Intelligence Powered Neuroimaging Biomarker for Chronic Pain" project.

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

References

1. Baskozos, G., Hébert, H.L., Pascal, M.M., Themistocleous, A.C., Macfarlane, G.J., Wynick, D., Bennett, D.L., Smith, B.H.: Epidemiology of neuropathic pain: an analysis of prevalence and associated factors in UK Biobank. *Pain Reports* **8**(2), e1066 (2023)
2. Biswal, B., Yetkin, F.Z., Haughton, V.M., Hyde, J.S.: Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic Resonance in Medicine* **34** (1995)
3. Borsook, D., Becerra, L., Hargreaves, R.: A role for fMRI in optimizing CNS drug development. *Nature Reviews Drug Discovery* **5**(5), 411–425 (2006)
4. Caro, J.O., de Oliveira Fonseca, A.H., Rizvi, S.A., Rosati, M., Averill, C., Cross, J.L., Mittal, P., Zappala, E., Dhodapkar, R.M., Abdallah, C., van Dijk, D.: BrainLM: A foundation model for brain activity recordings. In: *The Twelfth International Conference on Learning Representations* (2024)
5. Cole, D.M., Smith, S.M., Beckmann, C.F.: Advances and pitfalls in the analysis and interpretation of resting-state fMRI data. *Frontiers in Systems Neuroscience* **4**, 1459 (2010)
6. Colloca, L., Ludman, T., Bouhassira, D., Baron, R., Dickenson, A.H., Yarnitsky, D., Freeman, R., Truini, A., Attal, N., Finnerup, N.B., et al.: Neuropathic pain. *Nature Reviews Disease Primers* **3**(1), 1–19 (2017)
7. Esteban, O., Markiewicz, C.J., Blair, R.W., Moodie, C.A., Isik, A.I., Erramuzpe, A., Kent, J.D., Goncalves, M., DuPre, E., Snyder, M., et al.: fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nature Methods* **16**(1), 111–116 (2019)
8. Finnerup, N.B., Attal, N., Haroutounian, S., McNicol, E., Baron, R., Dworkin, R.H., Gilron, I., Haanpää, M., Hansson, P., Jensen, T.S., et al.: Pharmacotherapy for neuropathic pain in adults: a systematic review and meta-analysis. *The Lancet Neurology* **14**(2), 162–173 (2015)
9. Fitridge, R., Thompson, M.: *Mechanisms of vascular disease: a reference book for vascular specialists*. University of Adelaide Press (2011)
10. Fonov, V.S., Evans, A.C., McKinstry, R.C., Almlí, C.R., Collins, D.: Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *NeuroImage* **47**, S102 (2009)

11. Fox, M.D., Raichle, M.E.: Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience* **8**(9), 700–711 (2007)
12. Han, X., Nguyen, H., Harris, C., Ho, N., Saria, S.: FuseMoE: Mixture-of-experts transformers for fleximodal fusion. arXiv preprint arXiv:2402.03226 (2024)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 770–778 (2016)
14. International Association for the Study of Pain (IASP): IASP pain taxonomy: Peripheral neuropathic pain (2015), retrieved January 13, 2015
15. Kan, X., Dai, W., Cui, H., Zhang, Z., Guo, Y., Yang, C.: Brain network transformer. *Advances in Neural Information Processing Systems* **35**, 25586–25599 (2022)
16. Kawahara, J., Brown, C.J., Miller, S.P., Booth, B.G., Chau, V., Grunau, R.E., Zwicker, J.G., Hamarneh, G.: BrainNetCNN: Convolutional neural networks for brain networks; towards predicting neurodevelopment. *NeuroImage* **146**, 1038–1049 (2017)
17. Kim, Y.g., Ravid, O., Zheng, X., Kim, Y., Neria, Y., Lee, S., He, X., Zhu, X.: Explaining deep learning-based representations of resting state functional connectivity data: focusing on interpreting nonlinear patterns in autism spectrum disorder. *Frontiers in Psychiatry* **15**, 1397093 (2024)
18. Li, H., Fan, Y.: Brain decoding from functional MRI using long short-term memory recurrent neural networks. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16–20, 2018, Proceedings, Part III* 11. pp. 320–328. Springer (2018)
19. Liu, Q., Wang, X.: Bidirectional feature fusion and enhanced alignment based multimodal semantic segmentation for remote sensing images. *Remote Sensing* **16**(13), 2289 (2024)
20. Lötsch, J., Ultsch, A.: Machine learning in pain research. *Pain* **159**(4), 623–630 (2018)
21. Lötsch, J., Ultsch, A., Mayer, B., Kringel, D.: Artificial intelligence and machine learning in pain research: a data scientometric analysis. *Pain Reports* **7**(6), e1044 (2022)
22. Miller, K.L., Alfaro-Almagro, F., Bangerter, N.K., Thomas, D.L., Yacoub, E., Xu, J., Bartsch, A.J., Jbabdi, S., Sotiropoulos, S.N., Andersson, J.L., et al.: Multimodal population brain imaging in the UK Biobank prospective epidemiological study. *Nature Neuroscience* **19**(11), 1523–1536 (2016)
23. Nemati, S., Akiki, T.J., Roscoe, J., Ju, Y., Averill, C.L., Fouda, S., Dutta, A., McKie, S., Krystal, J.H., Deakin, J.W., et al.: A unique brain connectome fingerprint predates and predicts response to antidepressants. *IScience* **23**(1) (2020)
24. P, T., A, M., E, V.P., TJ, S., AV, A., MN, B.: Brain connectivity predicts placebo response across chronic pain clinical trials (2022)
25. Scheinost, D., Noble, S., Horien, C., Greene, A.S., Lake, E.M., Salehi, M., Gao, S., Shen, X., O’Connor, D., Barron, D.S., et al.: Ten simple rules for predictive modeling of individual differences in neuroimaging. *NeuroImage* **193**, 35–45 (2019)
26. Simon, R.: Clinical trial designs for evaluating the medical utility of prognostic and predictive biomarkers in oncology. *Personalized Medicine* **7**(1), 33–47 (2010)
27. Smith, S.M., Miller, K.L., Salimi-Khorshidi, G., Webster, M., Beckmann, C.F., Nichols, T.E., Ramsey, J.D., Woolrich, M.W.: Network modeling methods for fMRI. *NeuroImage* **54**(2), 875–891 (2011)

28. Specht, K.: Current challenges in translational and clinical fMRI and future directions. *Frontiers in Psychiatry* **10**, 924 (2020)
29. Sundararajan, M., Taly, A., Yan, Q.: Axiomatic attribution for deep networks. In: *International Conference on Machine Learning*. pp. 3319–3328. PMLR (2017)
30. Teh, K., Armitage, P., Tesfaye, S., Selvarajah, D.: Deep learning classification of treatment response in diabetic painful neuropathy: a combined machine learning and magnetic resonance neuroimaging methodological study. *Neuroinformatics* **21**(1), 35–43 (2023)
31. Wei, Y., Abrol, A., Hassanzadeh, R., Calhoun, V.: Hierarchical spatio-temporal state-space modeling for fMRI analysis. *arXiv preprint arXiv:2408.13074* (2024)
32. Xiong, Y., Zeng, Z., Chakraborty, R., Tan, M., Fung, G., Li, Y., Singh, V.: Nyströmformer: A nyström-based algorithm for approximating self-attention. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 35, pp. 14138–14148 (2021)
33. Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong, H., He, Q.: A comprehensive survey on transfer learning. *Proceedings of the IEEE* **109**(1), 43–76 (2020)