

Prompt Your Brain: Scaffold Prompt Tuning for Efficient Adaptation of fMRI Pre-trained Model

Zijian Dong^{1,2,3}, Yilei Wu^{1,2}, Zijiao Chen^{1,2}, Yichi Zhang^{1,2}, Yueming Jin^{3,4},
and Juan Helen Zhou^{1,2,3}(✉)

¹ Centre for Sleep and Cognition & Centre for Translational Magnetic Resonance Research, Yong Loo Lin School of Medicine, National University of Singapore, Singapore

² Human Potential Translational Research Program and Department of Medicine, Yong Loo Lin School of Medicine, National University of Singapore, Singapore

³ Department of Electrical and Computer Engineering & Integrative Sciences and Engineering Programme (ISEP), NUS Graduate School, National University of Singapore, Singapore

⁴ Department of Biomedical Engineering, National University of Singapore, Singapore

helen.zhou@nus.edu.sg

Abstract. We introduce **Scaffold Prompt Tuning (ScaPT)**, a novel prompt-based framework for adapting large-scale functional magnetic resonance imaging (fMRI) pre-trained models to downstream tasks, with high parameter efficiency and improved performance compared to fine-tuning and baselines for prompt tuning. The full fine-tuning updates all pre-trained parameters, which may distort the learned feature space and lead to overfitting with limited training data which is common in fMRI fields. In contrast, we design a hierarchical prompt structure that transfers the knowledge learned from high-resource tasks to low-resource ones. This structure, equipped with a Deeply-conditioned Input-Prompt (DIP) mapping module, allows for efficient adaptation by updating only 2% of the trainable parameters. The framework enhances semantic interpretability through attention mechanisms between inputs and prompts, and it clusters prompts in the latent space in alignment with prior knowledge. Experiments on public resting state fMRI datasets reveal ScaPT outperforms fine-tuning and multitask-based prompt tuning in neurodegenerative diseases diagnosis/prognosis and personality trait prediction, even with fewer than 20 participants. It highlights ScaPT’s efficiency in adapting pre-trained fMRI models to low-resource tasks.

Keywords: FMRI pre-trained models · Prompt tuning · Neurodegenerative disease · Neuroticism.

1 Introduction

In the realm of neuroimaging, the emergence of large-scale, self-supervised pre-trained models/foundation models for fMRI, has demonstrated a significant

promise in improving performance across a variety of downstream tasks through fully fine-tuning [11,16]. However, fine-tuning in fMRI models requires updating all the pre-trained parameters given target task training data, which is computationally intensive and time-consuming. Moreover, when confronted with a scarcity of training data, large pre-trained models are at risk of overfitting. It has also been shown that fine-tuning could distort the learned feature space [14].

In natural language processing (NLP), the concept of soft prompt tuning presents an efficient alternative for adapting large language models (LLMs) [10]. This technique involves keeping the original model frozen while only training soft prompts prepended to the input. Such an approach has been demonstrated to enhance the decoding of information represented in the human brain for language understanding using fMRI [14]. Soft prompts have also proven capable of distinguishing between different communities within brain networks [2]. Despite its efficiency, prompt tuning often results in decreased task performance compared to fine-tuning [7] and fails to leverage the extensive knowledge embedded in a diverse array of high-resource tasks [1]. Building predictive models using neuroimaging data of disease cohorts (e.g., neurodegenerative disease such as Alzheimer’s disease (AD)) is typically a low-resource task due to the small sample size. Given the emerging large-scale neuroimaging data from healthy cohorts with comprehensive behavioral phenotyping, brain-behavior mapping in healthy populations is usually categorized as high-resource tasks. The knowledge gained from these high-resource tasks is invaluable and could significantly bolster the model performance in low-resource applications (e.g., disease prognosis).

Several works in NLP have designed prompt tuning strategies under multi-task transfer learning framework [1,15,17], which transfer knowledge from high-resource tasks to low-resource ones. Their source prompts are trained through multitask learning and subsequently utilized to initialize the model for a specific target task. However, they either overlook the intricate relationship between prompts and input [15,17], or map an input to a task-agnostic prompt space without capturing the input’s different aspects for various tasks [1]. Furthermore, in NLP, soft prompts function as a “black box”, lacking a semantic understanding of the information encapsulated within the embeddings. Lacking of interpretability makes it suboptimal for clinical application.

In this study, we introduce **Scaffold Prompt Tuning (ScaPT)**, the first prompt-based adaptive framework for fMRI pre-trained model with remarkable parameter efficiency and superior performance using limited downstream training data. Our contributions are four-fold: 1) We have designed a hierarchical prompt structure that evolves soft prompts into three levels: *modular prompts* at the super-domain level (treated as fundamental skills of the model) for basic fMRI knowledge embedded in the model; *phenotype prompts* combining modular prompts to represent different phenotypes in the human brain; and *vertex prompts* merging phenotype prompts and a newly-initialized target prompt for target tasks. Integration is guided by attention between inputs and prompts. 2) We propose a shared Deeply-conditioned Input-Prompt (DIP) mapping module for both source and target training, designed to uniquely map inputs to different




Fig. 1. Schematic overview of Scaffold Prompt Tuning (ScaPT) framework. ScaPT operates in two stages: **Source Training (ST)**, where it creates phenotype prompts by interpolation of modular prompts, and **Target Training (TT)**, where it blends phenotype prompts with a new target prompt for downstream tasks with fewer resources. Interpolation weights are determined by the attention between the input and prompts.

prompt spaces. 3) We demonstrated that our proposed phenotype prompts form clusters, which align well with prior knowledge. Moreover, the proposed attention mechanism between inputs and prompts offers significant semantic interpretability. 4) Through experiments on two public resting state fMRI datasets, ScaPT demonstrates remarkable parameter efficiency, outperforming both fine-tuning and baseline prompt tuning methods by updating only 2% of the trainable parameters. It achieves superior results in neurodegenerative disease diagnosis/prognosis, as well as personality trait prediction, with very limited training data.

2 Method

Problem Setup. Given an fMRI pre-trained model $f(\cdot)$ with parameters θ , and a high-resource fMRI dataset $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}$, where each data $\mathbf{x}_i$ is paired with $\mathbf{y}_i = [y_{(1,i)}, y_{(2,i)}, \dots, y_{(n,i)}, \dots, y_{(N,i)}]$ corresponding to N high-resource tasks $\mathcal{T} = \{T_1, \dots, T_n, \dots, T_N\}$, our goal is to learn a new low-resource task T_{target} by efficiently updating parameters ϕ given the target task dataset $\mathcal{D}' = \{(\mathbf{x}', y')\}$ ($|\mathcal{D}| > |\mathcal{D}'|$). The number of updating parameters ϕ is much smaller than that of trainable parameters θ in the pre-trained model $f(\cdot)$.

Overview. Prompts in ScaPT are trainable embeddings that direct the model’s responses without changing its architecture. Shown in Figure 1, ScaPT operates

in two stages: Source Training (ST), in which trainable modular prompts (encoding abstract information and fundamental skills required for fMRI understanding) are interpolated to one participant-wise phenotype prompt. Participant-wise prompts corresponding to the same phenotype are averaged to formulate group-wise phenotype prompt, serving as knowledge base for the subsequent stage: Target Training (TT). It combines group-level phenotype prompts with a new target prompt to formulate a participant-wise vertex prompt for the target task. The weights for interpolation of prompts are from attention between projected input and prompts, after mapping the input to prompt space through Deeply-conditioned Input-Prompt (DIP) mapping.

Deeply-conditioned Input-Prompt (DIP) Mapping. Previous research indicates that soft prompts might not match input embedding spaces [9], making direct attention between input and prompts unreliable. A proposed solution involves a network for projecting inputs into prompt spaces [1], but it falls short in multitasking scenarios by projecting inputs uniformly, thus failing to capture task-specific information. It also increases training parameters due to the addition of a separate network and relies on linear projections, which may not effectively represent the complex relationships between inputs and prompts.

To address the above issues, we propose Deeply-conditioned Input-Prompt (DIP) mapping $\mathcal{M}$ through reusing $f(\cdot)$ (Figure 1). We introduce learnable deep conditioning (DC) tokens $\mathcal{C} = \{\mathbf{c}_1, \dots, \mathbf{c}_n, \mathbf{c}_N | \mathbf{c}_n \in \mathbb{R}^{e \times m}\}$ that are prepended to the input, to guide $f(\cdot)$ to map the input to an appropriate prompt space conditioned on the given task. The prediction output from $f(\cdot)$ is then fed into a linear projection layer and a Layer Norm $\text{LN}(\cdot)$ to avoid gradient explosion [1]. To deeply guide the conditional mapping from input to prompt spaces, we inject DC to every layer of $f(\cdot)$. Formally, the input-prompt mapping is defined as:

$$\mathbf{H}_{(n,i)} = f([\mathbf{c}_n; \mathbf{X}_i]); \quad \hat{\mathbf{H}}_{(n,i)} = \text{LN}(\text{NonLinear}(\mathbf{W}^T \mathbf{H}_{(n,i)})) \quad (1)$$

where $\mathbf{X}_i \in \mathbb{R}^{l \times m}$ is the “text-like” representation generated from fMRI following [16], describing the signal of each of m brain networks for each time point l . $\mathbf{X}_i$ prepended by DC $[\mathbf{c}_n; \mathbf{X}_i] \in \mathbb{R}^{(l+e) \times m}$ is input to the frozen $f(\cdot)$. $\mathbf{H}_{(n,i)} \in \mathbb{R}^h$ is the conditional output from $f(\cdot)$, with h representing the hidden dimension. $\mathbf{W} \in \mathbb{R}^{h \times h}$ is the projection parameter to be updated during training, and $\hat{\mathbf{H}}_{(n,i)} \in \mathbb{R}^h$ is the projected input.

Source Training (ST). A healthy cohort encompasses a large number of participants, with various phenotypes associated (high-resource tasks). In the first stage - Source Training (ST), we aim to train a set of phenotype prompts (PheP) that encapsulate information on different behavior-relevant brain phenotypes. It will serve as a source of knowledge for downstream tasks.

To better capture the relationship between input fMRI and prompts, and boost the capacity of PheP to match the complexity of input, we model each participant-wise PheP as an interpolation of a set of modular prompts (MoP),

with the computed attention between input and MoP as the weights for interpolation, as shown in Figure 1. Each MoP can be seen as a basic skill; solving a task involves combining these fundamental skills [15]. Formally, for a high-resource task $T_n \in \mathcal{T}$, our training objective is to maximize the likelihood of predicting the label $y_{(n,i)}$ as follows:

$$\max_{\mathbf{P}_{(n,i)}, \mathcal{M}} p_{\theta}(y_{(n,i)} | [\mathbf{P}_{(n,i)}; \mathbf{X}_i]) \quad (2)$$

where $\mathbf{P}_{(n,i)} \in \mathbb{R}^{d \times m}$ is the participant-wise PheP with length d for task T_n . $\mathbf{X}_i$ prepended by the prompt $[\mathbf{P}_{(n,i)}; \mathbf{X}_i] \in \mathbb{R}^{(l+d) \times m}$ is input to the frozen $f(\cdot)$ to predict $\hat{y}_{(n,i)}$. $\mathbf{P}_{(n,i)}$ is generated from the interpolation of K MoP:

$$\mathbf{P}_{(n,i)} = \sum_{k=1}^K \alpha_k^{(n,i)} \cdot \dot{\mathbf{P}}_k, \quad \alpha_k^{(n,i)} = \frac{e^{\langle \dot{\mathbf{P}}_k, \hat{\mathbf{H}}_{(n,i)} \rangle / \tau}}{\sum_{i=1}^K e^{\langle \dot{\mathbf{P}}_i, \hat{\mathbf{H}}_{(n,i)} \rangle / \tau}}; \quad \bar{\mathbf{P}}_n = \frac{1}{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{D}|} \mathbf{P}_{(n,i)} \quad (3)$$

where we compute the attention score $\alpha_k^{(n,i)}$ between MoP $\hat{\mathbf{P}}_k \in \mathbb{R}^h$ (max-pool of $\dot{\mathbf{P}}_k$ from $\mathbb{R}^{d \times m}$ to $\mathbb{R}^m$, followed by a linear transformation to $\mathbb{R}^h$) and $\hat{\mathbf{H}}_{(n,i)}$. τ is the temperature. $\mathbf{P}_{(n,i)}$ is a participant-wise prompt for $\mathbf{X}_i$. To transfer the learned prompts to the next stage, we average $\mathbf{P}_{(n,i)}$ to formulate one group-wise PheP $\bar{\mathbf{P}}_n$, which is utilized subsequently for next stage.

Target Training (TT). Prompts for a new task could be blended with pre-trained prompts to incorporate gained knowledge [1]. In the second stage - Target Training (TT), we first initialize a target prompt $\mathbf{P}_{\text{target}}$ tailored for a target task. To capitalize on the insights embedded in $\bar{\mathbf{P}}_n$, we learn a vertex prompt $\mathbf{P}^*$, by interpolating $\bar{\mathbf{P}}_n$ and $\mathbf{P}_{\text{target}}$ given attention computed by $\mathcal{M}$ (Figure 1). Similar to ST, the goal of TT is to maximize the likelihood of predicting the correct target task label y' , given the concatenation of $\mathbf{P}^*$ and input $\mathbf{X}'$:

$$\max_{\mathbf{P}_{\text{target}}, \mathcal{M}} p_{\theta}(y' | [\mathbf{P}^*; \mathbf{X}']); \quad \mathbf{P}^* = \mathbf{P}_{\text{target}} + \sum_{n=1}^N \beta_n \cdot \bar{\mathbf{P}}_n \quad (4)$$

where $\mathbf{P}^*$ is the interpolation of $\mathbf{P}_{\text{target}}$ and $\bar{\mathbf{P}}_n$. β_n is the attention score between $\mathbf{X}'$ and the $\bar{\mathbf{P}}_n$ computed by $\mathcal{M}$.

3 Experiments

Datasets. In ST, resting state fMRI data from 656 participants of the Lifespan Human Connectome Project Aging (HCP-A) [3,6] were analyzed to predict 38 phenotypes (see details in the supplementary material), alongside sex and age. The brain-behavior phenotypes established at this stage include three domains: cognition, personality, and social emotion. We hypothesized that ScaPT would perform well in tasks relevant to these domains. In TT, ScaPT was assessed on

two classification tasks for neurodegenerative disease diagnosis/prognosis (related to cognition) using Alzheimer’s Disease Neuroimaging Initiative (ADNI) [8], and a regression task for personality trait prediction (related to personality and social emotion) using UK Biobank (UKB) [13]. fMRI in MNI space were preprocessed using fMRIPrep [5] (20.2.3) with the default settings and parcellated to 1024 networks using DiFuMo with Nilearn library (0.10.4).

The experiments of TT were performed similar to [16], using limited data to showcase the adaptation performance and varying the size to demonstrate how performance scales. (1) Control Normal (CN) v.s. Mild Cognition Impairment (MCI) classification: 340 ADNI participants, with 170 for each class. We first allocated an independent test set consisting of 100 CNs and 100 MCIs. From the remaining 140, we randomly sampled 3/5/10 data per class for training, using the rest for validation. Sampling was repeated 10 times, and accuracy and F1 score were reported with mean and standard deviation. (2) Amyloid Positive v.s. Negative classification: 100 ADNI participants with normal cognition (50 for each class). It was performed under the similar experimental setting with an independent test set consisting of 25 for each class. (3) Neuroticism Score Prediction: Neuroticism describes a tendency to experience unsettling feelings. The score ranges from 0 to 12 (normalized to 0-1 for training), with higher score corresponding to a greater tendency. We have 1000 UKB participants with neuroticism scores (800 for training and 200 for testing). Training samples were increased to 30/50/100 (randomly sampled from 800), following the same settings as the prior tasks.

Training Details. We adopt the state-of-the-art fMRI language model [16], with causal sequence modeling structure, for our downstream adaptation. It was pre-trained using 11,980 runs of 1,726 individuals across 34 datasets. The pre-trained model contains 4 GPT-2 layers [12], with 12 attention heads in each self-attention module. To utilize the pre-trained model, the input must be parcellated by Dictionaries of Functional Modes (DiFuMo) [4]. The preprocessed input to the model is $\mathbf{X} \in \mathbb{R}^{l \times m}$ obtained by DiFuMo, where l is the input sequence length and $m = 1024$ networks. The hidden dimension in the model is $h = 768$. We utilized $K = 5$ modular prompts. Prompt/DC length is $d = e = 5$. Refer to supplementary for more hyper-parameters and experimental settings.

Main Results. In a scenario with limited training data, we evaluated ScaPT against fine-tuning and three multitask-based prompt-tuning approaches: SPoT [17], MP² [15], and ATTEMPT [1] (Table 1 and 2). For fine-tuning, $f(\cdot)$ underwent direct fine-tuning using TT datasets (ADNI/UKB). Meanwhile, for the multitask-based methods, prompts were initially trained on ST datasets (HCP-A), which were then served as the prompt initialization in the TT stage.

ScaPT demonstrated superior performance over both fine-tuning and other prompt tuning methods across various sizes of training datasets, scaling well with the number of training data. This underscores ScaPT’s effectiveness in transferring knowledge from high-resource tasks to those with scarce resources.

Table 1. Accuracy (Acc) and F1 score on CN v.s. MCI and Amyloid $a\beta + ve$ v.s. $a\beta - ve$ classification (% , mean(standard deviation) for 10 independent runs), trained on varying training dataset size ($I = 3, 5, 10$ per class). The best results are in **bold**, with * denoting significant improvement ($p < 0.05$).

Methods	$I = 3$		$I = 5$		$I = 10$	
	Acc $\uparrow$	F1 $\uparrow$	Acc $\uparrow$	F1 $\uparrow$	Acc $\uparrow$	F1 $\uparrow$
<i>CN v.s. MCI</i>						
Fine-tuning [16]	53.0(3.1)	47.3(3.7)	61.1(3.9)	57.1(5.4)	65.3(2.7)	62.0(3.2)
SPoT [17]	56.3(.70)	51.4(1.6)	65.6(.35)	62.9(.66)	64.8(2.0)	61.2(2.5)
MP ² [15]	65.0(2.9)	59.2(3.8)	71.4(3.1)	68.3(3.4)	74.4(2.0)	71.3(2.5)
ATTEMPT [1]	64.3(1.5)	64.4(2.1)	75.0(1.9)	74.9(2.0)	77.3(2.0)	76.5(2.2)
ScaPT	69.7*(2.8)	67.4*(2.9)	75.1(2.1)	73.3(2.8)	80.3*(2.2)	79.4*(2.2)
<i>Amyloid $a\beta + ve$ v.s. $a\beta - ve$</i>						
Fine-tuning [16]	51.1(.15)	52.3(1.1)	64.2(0.2)	66.1(3.2)	71.1(.15)	72.2(.01)
SPoT [17]	55.9(.49)	50.7(.88)	58.9(.20)	51.1(.71)	64.5(.35)	58.1(.26)
MP ² [15]	51.3(2.2)	53.2(1.2)	70.6(.13)	71.2(2.1)	78.4(.15)	79.1(.21)
ATTEMPT [1]	57.6(2.6)	61.7(1.3)	68.1(.37)	70.2(.69)	80.6(1.6)	81.5(1.4)
ScaPT	61.2*(1.2)	63.1*(1.1)	71.2(.53)	72.1*(.25)	86.0*(.10)	87.7*(.10)

Table 2. Mean Absolute Error (MAE, $\times 10^{-1}$) and Pearson Correlation (ρ) on Neuroticism score prediction.

Methods	$I = 30$		$I = 50$		$I = 100$	
	MAE $\downarrow$	ρ $\uparrow$	MAE $\downarrow$	ρ $\uparrow$	MAE $\downarrow$	ρ $\uparrow$
Fine-tuning [16]	3.03(.43)	0.37(.008)	2.57(.21)	0.40(.002)	2.35(.00)	0.45(.002)
SPoT [17]	5.29(.74)	0.35(.005)	4.43(.69)	0.38(.004)	4.37(.58)	0.39(.003)
MP ² [15]	3.92(.15)	0.40(.004)	3.59(.16)	0.41(.003)	3.41(.12)	0.42(.002)
ATTEMPT [1]	2.85(.18)	0.38(.004)	2.37(.19)	0.43(.006)	2.11(.21)	0.45(.004)
ScaPT	2.53*(.27)	0.40(.003)	2.40(.25)	0.45*(.001)	1.90*(.06)	0.49*(.003)

Prompt Interpretation. In the ST stage, we created 40 phenotype prompts for 38 phenotype predictions, alongside age and sex determinations. Remarkably, these prompts naturally formed into three clusters - Personality, Social Emotion, and Cognition - without prior supervision, indicating they effectively capture different pillars of brain-behavior associations (Figure 2-1).

During the TT stage, the attention scores between input and phenotype prompts (Figure 2-2) aid in interpreting the target task. Attention score vectors were averaged across inputs, and then attributes within each group were averaged (without the one for $\mathbf{P}_{\text{target}}$). These five values were normalized for analysis. Aligning well with the literature, ScaPT shows a focus on "cognition" or "age" in neurodegenerative disease diagnosis/prognosis task, while it focuses on "personality" for neuroticism score prediction.




Fig. 2. Further analysis: (1) visualization of 2D phenotype prompts space, (2) interpretation of target tasks through attention distribution, (3) ablation study, and (4) comparison of model performance versus numbers of parameters.

Ablation Study and Parameter-efficiency. We evaluated ScaPT against its ablations (Figure 2-3), including ScaPT w/o PheP (using MoP directly for $\mathbf{P}^*$ formulation without high-resource task training), ScaPT w/o MoP (learning $\bar{\mathbf{P}}_n$ without prompt width expansion), and ScaPT w/o DIP (utilizing ATTEMPT’s subnetwork for input-prompt mapping). The absence of PheP led to a significant performance drop, underscoring the importance of high-resource task knowledge in boosting low-resource task performance. ScaPT outperformed its counterparts lacking MoP, demonstrating MoP’s role in enhancing expressive capacity by widening prompts. Additionally, ScaPT’s DIP module surpassed ATTEMPT’s subnetwork in mapping inputs to prompts, effectively capturing complex input-prompt relationships using $f(\cdot)$. In Figure 2-4, we compared ScaPT’s performance with other models relative to their trainable parameters. ScaPT significantly outperformed fine-tuning, MP^2 , and ATTEMPT, despite updating only 2% of total parameters. Although SPoT had the fewest trainable parameters, its performance lagged, likely due to its limited feature capture capability.

4 Conclusion

We introduce Scaffold Prompt Tuning (ScaPT), the first prompt-based adaptation framework for large-scale fMRI pre-trained models, compatible with very limited training data. ScaPT features a hierarchical prompt structure that facilitates knowledge transfer from high-resource tasks to those with fewer resources. Moreover, we develop a Deeply-conditioned Input-Prompt (DIP) mapping to capture the complex relationship between the input and prompt spaces. Our experiments demonstrate ScaPT’s exceptional parameter efficiency and its superior performance in neurodegenerative disease diagnosis or prognosis, as well as personality trait prediction from resting-state fMRI data. In addition, our attention mechanism offers semantic interpretation for target tasks. Future studies could expand ScaPT’s reach to longitudinal data, and explore its possibility for prompting vision models.

Acknowledgments. This study was supported by the Singapore National Medical Research Council (NMRC\OFLCG19May-0035, NMRC\CIRG\1485\2018, NMRC\CSA-SI\0007\2016, NMRC\MOH-00707-01, NMRC\CG\435 M009\2017-NUH\NUHS, CIRG21nov-0007 and HLCA23Feb-0004), RIE2020 AME Programmatic Fund from A*STAR (No. A20G8b0102), Ministry of Education (MOE-T2EP40120-0007 & T2EP2-0223-0025, MOE-T2EP20220-0001), and Yong Loo Lin School of Medicine Research Core Funding, National University of Singapore, Singapore. Yueming Jin was supported by MoE Tier 1 Start up grant (WBS: A-8001267-00-00).

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

References

1. Asai, A., Salehi, M., Peters, M.E., Hajishirzi, H.: Attempt: Parameter-efficient multi-task tuning via attentional mixtures of soft prompts. In: Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing. pp. 6655–6672 (2022)
2. Bannadabhavi, A., Lee, S., Deng, W., Ying, R., Li, X.: Community-aware transformer for autism prediction in fmri connectome. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 287–297. Springer (2023)
3. Bookheimer, S.Y., Salat, D.H., Terpstra, M., Ances, B.M., Barch, D.M., Buckner, R.L., Burgess, G.C., Curtiss, S.W., Diaz-Santos, M., Elam, J.S., et al.: The lifespan human connectome project in aging: an overview. *Neuroimage* **185**, 335–348 (2019)
4. Dadi, K., Varoquaux, G., Machlouzarides-Shalit, A., Gorgolewski, K.J., Wassermann, D., Thirion, B., Mensch, A.: Fine-grain atlases of functional modes for fmri analysis. *NeuroImage* **221**, 117126 (2020)
5. 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.: fmripipeline: a robust preprocessing pipeline for functional mri. *Nature methods* **16**(1), 111–116 (2019)

6. Harms, M.P., Somerville, L.H., Ances, B.M., Andersson, J., Barch, D.M., Bastiani, M., Bookheimer, S.Y., Brown, T.B., Buckner, R.L., Burgess, G.C., et al.: Extending the human connectome project across ages: Imaging protocols for the lifespan development and aging projects. *Neuroimage* **183**, 972–984 (2018)
7. He, J., Zhou, C., Ma, X., Berg-Kirkpatrick, T., Neubig, G.: Towards a unified view of parameter-efficient transfer learning. In: *International Conference on Learning Representations* (2021)
8. Jack Jr, C.R., Bernstein, M.A., Fox, N.C., Thompson, P., Alexander, G., Harvey, D., Borowski, B., Britson, P.J., L. Whitwell, J., Ward, C., et al.: The alzheimer’s disease neuroimaging initiative (adni): Mri methods. *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine* **27**(4), 685–691 (2008)
9. Khashabi, D., Lyu, X., Min, S., Qin, L., Richardson, K., Welleck, S., Hajishirzi, H., Khot, T., Sabharwal, A., Singh, S., et al.: Prompt waywardness: The curious case of discretized interpretation of continuous prompts. In: *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. pp. 3631–3643 (2022)
10. Lester, B., Al-Rfou, R., Constant, N.: The power of scale for parameter-efficient prompt tuning. In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. pp. 3045–3059 (2021)
11. Ortega Caro, J., Oliveira Fonseca, A.H., Averill, C., Rizvi, S.A., Rosati, M., Cross, J.L., Mittal, P., Zappala, E., Levine, D., Dhodapkar, R.M., et al.: Brainlm: A foundation model for brain activity recordings. *bioRxiv* pp. 2023–09 (2023)
12. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners
13. Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., Downey, P., Elliott, P., Green, J., Landray, M., et al.: Uk biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS medicine* **12**(3), e1001779 (2015)
14. Sun, J., Moens, M.F.: Fine-tuned vs. prompt-tuned supervised representations: which better account for brain language representations? In: *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*. pp. 5197–5205 (2023)
15. Sun, T., He, Z., Zhu, Q., Qiu, X., Huang, X.J.: Multitask pre-training of modular prompt for chinese few-shot learning. In: *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. pp. 11156–11172 (2023)
16. Thomas, A., Ré, C., Poldrack, R.: Self-supervised learning of brain dynamics from broad neuroimaging data. *Advances in Neural Information Processing Systems* **35**, 21255–21269 (2022)
17. Vu, T., Lester, B., Constant, N., Al-Rfou, R., Cer, D.: Spot: Better frozen model adaptation through soft prompt transfer. In: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. pp. 5039–5059 (2022)