

Fine-tuned vs. Prompt-tuned Supervised Representations: Which Better Account for Brain Language Representations?

Jingyuan Sun and Marie-Francine Moens

Department of Computer Science, KU Leuven, Belgium

{jingyuan.sun, sien.moens}@kuleuven.be

Abstract

To decipher the algorithm underlying the human brain’s language representation, previous work probed brain responses to language input with pre-trained artificial neural network (ANN) models fine-tuned on NLU tasks. However, full fine-tuning generally updates the entire parametric space and distorts pre-trained features, cognitively inconsistent with the brain’s robust multi-task learning ability. Prompt-tuning, in contrast, protects pre-trained weights and learns task-specific embeddings to fit a task. Could prompt-tuning generate representations that better account for the brain’s language representations than fine-tuning? If so, what kind of NLU task leads a pre-trained model to better decode the information represented in the human brain? We investigate these questions by comparing prompt-tuned and fine-tuned representations in neural decoding, that is predicting the linguistic stimulus from the brain activities evoked by the stimulus. We find that on none of the 10 NLU tasks, full fine-tuning significantly outperforms prompt-tuning in neural decoding, implicating that a more brain-consistent tuning method yields representations that better correlate with brain data. Moreover, we identify that tasks dealing with fine-grained concept meaning yield representations that better decode brain activation patterns than other tasks, especially the syntactic chunking task. This indicates that our brain encodes more fine-grained concept information than shallow syntactic information when representing languages.

1 Introduction

To decipher the mechanism underlying human language representation by analyzing brain activation patterns is a fundamental goal of language neurobiology [Schrimpf *et al.*, 2021; Huth *et al.*, 2016]. The recent upsurge of ANNs for language processing, exemplified by the Transformer [Vaswani *et al.*, 2017; Devlin *et al.*, 2018] and its mutants [Ouyang *et al.*, 2022; Kojima *et al.*, 2022], has yielded significant advances in solving a wide range of downstream NLP tasks. It is logical to investigate the potential of ANNs to assist in reverse-

engineering the algorithms that underlie language representation in the human brain.

Existing studies have employed representations produced by ANN models to explain the brain activities related to language processing [Anderson *et al.*, 2016; Sun *et al.*, 2019]. Some of the best-performing models such as GPT-2 [Radford *et al.*, 2019] are found to account for more than 99% of explainable variance in neural responses for a corresponding linguistic stimulus [Schrimpf *et al.*, 2021]. Transformer-based large-scale language models show a consistent advantage over other models for predicting and deciphering brain activities [Sun *et al.*, 2020b]. Among these works, unsupervised methods optimized for language modeling or similar context-predicting goals take a major part in the evaluated ANN models, while supervised models are far less explored. Could additional information learned from task supervision lead to better candidate models to account for the brain’s language representation?

To study this question, previous work mostly exclusively fine-tuned pre-trained models, producing supervised representations to match with brain activation patterns [Gauthier and Levy, 2019; Oota *et al.*, 2022]. To fit a new NLU task, the full fine-tuning generally updates the entire parameter space of the model which has been proven to distort the pre-trained features [Kumar *et al.*, 2022]. However, this is potentially inconsistent with the human brain’s mechanism which does not require a reformation of the entire brain’s language network to learn a single new task.

In this study, we propose to link the human brain and task-supervised artificial representations with prompt-tuning. We take 10 NLU tasks involving diverse linguistic-cognitive skills from widely used NLU benchmarks such as GLUE [Wang *et al.*, 2018] and SuperGLUE [Sarlin *et al.*, 2020]. A pre-trained BERT model is tuned on the 10 tasks with full fine-tuning and prompt-tuning, respectively. The supervised representations generated by the tuned models are then employed to decipher the brain activation patterns through neural decoding. We comprehensively compare the prompt-tuned against the fine-tuned representations. We also demonstrate which task leads to representations that are better at decoding the brain responses. We find that:

- (i) On the regions of interest (ROIs) comprising the brain

language and semantic network¹, prompt-tuned representations consistently and significantly outperform the fully fine-tuned peers in decoding accuracy.

- (ii) Tuning on tasks dealing with fine-grained concept meaning including Word Sense Disambiguation and Co-reference Resolution yields representations that are better at neural decoding than tuning on other tasks with both fine-tuning and prompt-tuning.

The proposed prompt-tuning-based framework helps a better understanding of the relationship between supervised artificial and brain language representations: the more similar to the brain working mechanism, the better correlated with the brain data. Furthermore, our experimental findings contribute to deciphering the mechanism supporting the human brain’s language processing.

2 Related Work

2.1 Fine-tuning and Prompt-tuning

In recent years, there have been significant developments in pre-trained language models [Brown *et al.*, 2020; Sun *et al.*, 2020a], resulting in remarkable enhancements in performance across various NLP tasks. One widely used technique, known as full fine-tuning, involves updating all model parameters for a specific task, as depicted in Figure 1[a]. While full fine-tuning generally yields good task performance, it can be computationally intensive due to the requirement of storing gradients and optimizer states for all parameters. Moreover, the cognitive plausibility of fine-tuning can be hindered by its potential distortion of general domain knowledge acquired during pre-training. To address these concerns, an approach called discrete prompting has been developed as an alternative of fine-tuning [Liu *et al.*, 2022a]. With discrete prompting, the parameters of a pre-trained model are fixed, and an NLP task is solved by querying the model with a discrete natural language prompt. For example, in machine translation, one might prepend a sample like "A lovely day" with a prompt such as "Translate it to Chinese" and ask the pre-trained language model to predict the subsequent tokens.

While discrete prompting does not require additional training and only necessitates storing a single set of model parameters, its task performance can be highly dependent on the design of the prompt template. As a result, it may be sub-optimal compared to fine-tuning in certain cases [Qiu *et al.*, 2020; Lester *et al.*, 2021]. Due to this limitation, we do not adopt discrete prompting as a baseline in this paper. Prompt-tuning, an extension of discrete prompting, has been introduced as a technique that fixes the parameters of a pre-trained model while learning continuous prompt embeddings [Qin and Eisner, 2021]. An example of prompt-tuning is prefix-tuning [Li and Liang, 2021], which involves adding trainable embeddings called *prefixes* to the original input word embeddings and optimizing them for downstream tasks. Prompt-tuning has shown superior performance compared to discrete prompting on multiple tasks [Chen *et al.*, 2022] and achieves

¹A brain functional network is a collection of brain regions that consistently show coordinated activities for certain cognitive functions or during behavioral tasks.

performance levels close to fine-tuning. In this paper, we leverage the characteristics of prompt-tuning to preserve the knowledge learned by a pre-trained model. This enables us to investigate whether the additional supervised information obtained from downstream tasks improves neural decoding performance or not. Specifically, we employ the P-tuning V2 method, as illustrated in Figure 1[b] [Liu *et al.*, 2022b], which is an optimized version of prefix-tuning claiming comparable performance to fine-tuning across a range of natural language understanding tasks.

2.2 Matching Artificial and Neural Representations of Language

The use of artificial neural network (ANN) representations has proven successful in encoding and decoding neural responses to linguistic stimuli [Pereira *et al.*, 2018; Caucheteux *et al.*, 2021]. Early studies have demonstrated their ability to recover simple verbal stimuli, such as words [Pereira *et al.*, 2011; Handjaras *et al.*, 2016], as well as phrases [Fyshe, 2015; Huth *et al.*, 2016], from brain activation patterns recorded by functional magnetic resonance imaging (fMRI). More recent work have expanded such capability to decoding sentences using similar methods [Matsuo *et al.*, 2016; Oota *et al.*, 2018; Pereira *et al.*, 2018; Wang *et al.*, 2022].

The majority of ANN models evaluated in previous studies have focused on unsupervised representations optimized for language modeling or context-predicting tasks, while supervised models have received comparatively less attention. One relevant research [Oota *et al.*, 2022] fine-tuned a pre-trained BERT model on a series of downstream tasks and compared their neural encoding performances. However, since fine-tuning differs from how the human brain adapts to new tasks, we question whether fine-tuning is the most suitable framework for obtaining supervised representations to probe the human brain. To investigate this, we conduct a comprehensive comparison of experiments using fine-tuned and prompt-tuned representations for neural decoding. It’s important to note that our focus in this paper is more on deciphering the information represented in the human brain for language understanding, so we do not conduct neural encoding experiments as done in [Oota *et al.*, 2022].

3 Methods

In the following subsections, we will first introduce how prompt-tuning with the P-tuning method works. We then demonstrate how the representations produced by prompt-tuned models are used to build a neural decoder.

3.1 Prompt-Tuning

We build upon the P-tuning method [Liu *et al.*, 2022b] to prompt-tune a pre-trained model on downstream tasks. In Figure 1, we illustrate the mechanism of P-tuning with a Transformer-based language model. P-tuning prepends a continuous trainable prefix to different layers of an ANN to guide its fitting on downstream tasks.

We now explain the mechanism of P-tuning, with tuning a Transformer-based language model on a conditional generation task as an example, following [Li and Liang, 2021]. For

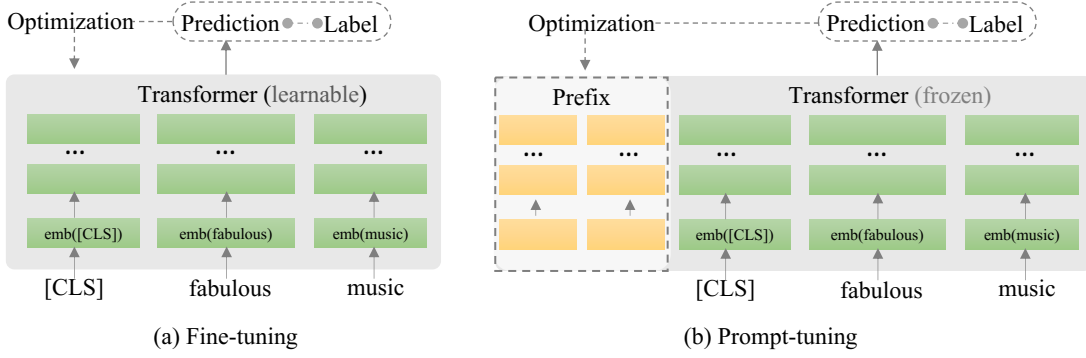


Figure 1: Tuning Transformer on downstream tasks [a] with fine-tuning [b] with prompt-tuning.

a given conditional generation task, where the input context is denoted as x and the output token sequence is denoted as y , we consider a language model $LM = p_\phi(y|x)$, which is parameterized by ϕ . To facilitate the P-tuning process, we define the variable z as the concatenation of x and y , and we represent the corresponding sequence of indices for x and y as id_x and id_y , respectively. At each time step i , the activation is denoted as h_i , which is obtained by concatenating the activation layers at that specific time step, represented as $h_i = [h_i^1, h_i^2, \dots, h_i^n]$. Here, h_i^j refers to the activation of the j -th Transformer layer at time step i . The computation of h_i relies on z_i and the past activations within its context, and it can be formulated as follows: $h_i = \mathbf{LM}_\phi(z_i, h)$. The distribution for the next token is computed as follows: $p_\phi(z_{i+1} | h_{\leq i}) = \text{softmax}(W_\phi h_i^{(n)})$. In this equation, $h_i^{(n)}$ represents the hidden states of the last layer within h_i . The matrix W_ϕ is pre-trained to map these hidden states $h_i^{(n)}$ to logits across the vocabulary.

During P-tuning, a prefix, denoted as Pr , is prepended to the input sequence, resulting in the concatenation $z = [Pr; x; y]$. Pr is a continuous trainable embedding, its corresponding sequence of indices is denoted as Pr_{id_x} with length $|Pr_{id_x}|$. A trainable matrix M_θ^{Pr} , with dimensions $|Pr_{id_x}| \times \dim(h_i)$, is initialized to store the prefix parameters. And

$$h_i = \begin{cases} M_\theta^{Pr}[i, :], & \text{if } i \in Pr_{id_x} \\ \mathbf{LM}_\phi(z_i, h_{<i}), & \text{otherwise.} \end{cases} \quad (1)$$

Training P-tuning is then to optimize the following log-likelihood objective:

$$\max_\phi \log p_\phi(y | x) = \sum_{i \in id_y} \log p_\phi(z_i | h_{<i}) \quad (2)$$

with ϕ being fixed and θ being trainable.

3.2 Neural Decoder

The neural decoder aims to categorize or reconstruct the perceived stimulus by mining the corresponding neural activities. As shown in Figure 2, we tune the pre-trained Transformer model respectively with fine-tuning and prompt-tuning. The tuned models generate sentence embeddings to

represent the sentences in the brain imaging dataset. We then build neural decoders with sentence embeddings to decipher the brain images. Here we use regression-based decoding where a semantic vector is directly estimated from the voxels by regression.

Specifically, given the voxel matrix $V_d \in \mathbb{R}^{N_E \times N_V}$ and sentence embedding matrix $Z_d \in \mathbb{R}^{N_E \times N_D}$ in the training set, where N_E represents the number of examples, N_V represents the number of voxels, and N_D represents the number of dimensions of the sentence embedding, the regression coefficients of the decoder W_d are estimated by minimizing

$$\|W_d V - z_i\|_2^2 + \lambda \|W_d\|_2^2 \quad (3)$$

for each column z_i in Z , i.e., each dimension of the sentence vectors. λ is the regularization parameter.

The decoder is trained using a 5-fold cross-validation approach. The regression parameters for each dimension are learned by mapping the brain images of the training sentences to distributed representations. Then, the performance is tested by predicting semantic vectors from the brain images of the left-out testing sentences in each fold. This cross-validation procedure is applied to data from all subjects with each supervised representation. Upon completion of the 5 folds, this process yields decoded semantic vectors for each subject and supervised representation combination.

4 Experimental Setup

In this section, we will first briefly introduce the brain imaging dataset we use. We then demonstrate the principle of task selection and details of the selected tasks, as well as tuning on these tasks to generate supervised representations. Codes, checkpoints and other supplementary materials are available online².

4.1 Brain Imaging Data

We utilize the fMRI data published in the study conducted by [Pereira *et al.*, 2018]. The fMRI neural activation data were acquired using a 3-Tesla Siemens Trio scanner equipped with a 32-channel head coil. Two experiments were conducted in

²https://github.com/soinx0629/sup_fmri_dec_eng

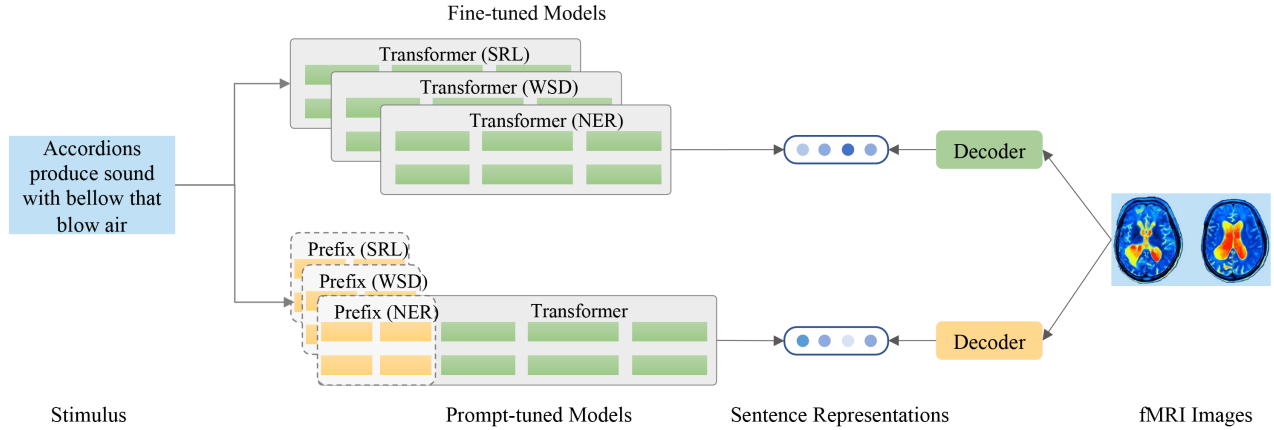


Figure 2: Neural decoding with fully fine-tuned and prompt-tuned supervised representations.

this research, and 5 subjects participated in both of the 2 experiments. The stimuli used in both experiments followed a hierarchical structure of topic-passage-sentence. In the first experiment, a total of 96 passages, resembling Wikipedia-style content, were presented. Each passage consisted of 4 sentences that focused on a specific concept. These concepts were selected from a pool of 24 broad topics, providing a diverse range of content. The second experiment involved 72 passages, including 48 Wikipedia-style passages and 24 passages written in first/third-person narratives. Each passage contains 3 or 4 sentences about a particular concept unrelated to those in the first experiment. The two experiments were designed to maintain similar semantic relatedness within and between passages/topics. In this study, we combine the data from the two experiments to make full use of the brain images. There will be imaging data from 5 subjects scanned with 672 sentences presented as stimuli. Such a setting is consistent with previous work [Sun *et al.*, 2019; Oota *et al.*, 2022]. The full details about the presented materials and experimental setup can be found online³.

4.2 Neural Decoding Evaluation

The evaluation of the decoding results primarily relies on the pairwise matching task, a widely used metric in neural decoding research. The correlation between the decoded vectors and the ground-truth sentence embeddings is first computed. A successful matching is scored when the decoded semantic vectors have a higher degree of similarity with their corresponding brain activation patterns compared to alternative pairings. Formally, for each possible pair of sentence stimuli S_i and S_j , let X_{S_i} and X_{S_j} denote the brain images of neural responses to S_i and S_j . Let Z_{S_i} and Z_{S_j} denote the sentence embeddings produced by a tuned model, while D_{S_i} and D_{S_j} denote the decoded semantic vectors from brain images X_{S_i} and X_{S_j} . We score 1 for a matching if

$$\begin{aligned} \text{corr}(D_{S_i}, Z_{S_i}) + \text{corr}(D_{S_j}, Z_{S_j}) > \\ \text{corr}(D_{S_i}, Z_{S_j}) + \text{corr}(D_{S_j}, Z_{S_i}), \end{aligned} \quad (4)$$

³<https://osf.io/crwz7/wiki/home/>

else 0. corr denotes the correlation function employed which is Pearson’s correlation in this study. An examination of all possible pairs of sentences will be conducted. The final matching accuracy for each participant is determined by calculating the proportion of correctly matched pairs. The expected chance-level accuracy, assuming a random selection by the model, is 0.50.

In the experimental results, we will mainly display the pairwise matching performance which is most widely adopted in related work [Pereira *et al.*, 2018; Sun *et al.*, 2019; Oota *et al.*, 2022]. Except for pair-wise matching, we also use the MSE loss between the decoding value and ground truth as a supplementary metric. Please refer to the appendix for MSE results.

4.3 Tasks for Tuning

The tasks chosen for this study are based on the two major principles taking into account the setting of previous work [Oota *et al.*, 2022]:

- (i) tasks involving a diverse range of cognitive-linguistic skills are desired;
- (ii) tasks with corpora that are commonly used in popular natural language processing (NLP) benchmarks.

Based on the above principles, we select tasks analyzing emotion and sentiment polarity (*Emotion Categorization and Sentiment Classification*), tasks dealing with concept meaning (*Word Sense Disambiguation, Co-reference Resolution and Named Entity Recognition*), tasks that require reasoning semantic relation (*Paraphrase Detection and Semantic Role Labeling*), tasks requiring knowledge and logical representations (*Question Answering and Natural Language Inference*) as well as tasks that parse the syntactic structure (*Syntactic Chunking*). We present these tasks in Table 1.

4.4 Generating Sentence Embeddings

The pre-trained Transformer model employed in this study is the open-source BERT-base-based. As shown in Figure 2, task-specific tuning of the BERT model with fine-tuning and prompt-tuning are respectively conducted on each of the 10 tasks. These tuned models generate the supervised sentence

Task	Abbr.	Corpus
Emotion Categorization	Emo	emotion
Sentiment Classification	SC	SST2
Question Answering	QA	SQuAD
Semantic Role Labeling	SRL	Conll05
Word Sense Disambiguation	WSD	WiC
Co-reference Resolution	Coref.	WSC
Named Entity Recognition	NER	Conll03-NER
Natural Language Inferencing	NLI	MNLI
Syntactic Chunking	Chunk	Conll03-chunk
Paraphrase Detection	Para	MRPC

Table 1: 10 tasks used for tuning supervised representations. Abbr. denotes the abbreviation for the task name.

embeddings needed for this study. Sentences presented for collecting the brain images are input to the tuned models one by one. The hidden states of the last layer are averaged as sentence embeddings. The generated sentence embeddings have a dimension of 768. We train on 4 Nvidia 3080TI GPUs with 12GB VRAM. Hyper-parameter searching is conducted to achieve optimal task performance. The influences of hyper-parameters settings including the size of prompts on the task performance are recorded in the appendix.

5 Experimental Results

In this section, we compare the decoding performance of fine-tuned and prompt-tuned representations. We demonstrate how supervised representations trained on the 10 tasks differently decode the voxels of 4 major brain functional networks and the regions of interests (ROIs). We analyse the representational similarity structure of the supervised representations to explain the experimental results. We validate the findings from analyzing the decoding performance by paired t-test with the Bonferroni correction and report the results throughout with a significance level of $\alpha = 0.01$.

5.1 Decoding Within Functional Networks

Following previous work [Pereira *et al.*, 2018; Sun *et al.*, 2020b], we select (i) the fronto-temporal language selective network [Fedorenko *et al.*, 2011], (ii) the ensemble semantic system [Binder *et al.*, 2009], (iii) the visual network [Buckner *et al.*, 2008] (including scene, face, body and object network) and (iii) the default mode network (DMN⁴). We conduct neural decoding on the voxels from each of the ROI comprising the functional networks and average the decoding performance across the entire network.

As shown in Figure 3, we first observe a consistent and obvious advantage of prompt-tuning representations over fine-tuned peers. Specifically, we find a stark contrast in average

⁴DMN is a network known for being active when the brain is at wakeful rest or during detailed thoughts related to external task performance.

decoding accuracy between fine-tuned and prompt-tuned representations in the tasks requiring reasoning with the syntactic structure (Chunk), and tasks analyzing the sentiment and emotion (SC and Emo). The accuracy differences are significant across the scanned subjects ($p < 10^{-3}$). On decoding the language and semantic network, although representations fine-tuned on the WSD task slightly outperform the prompt-tuned one in average decoding accuracy, their difference does not stand up to the significance test among the scanned subjects ($p > 10^{-2}$).

There are also tasks whose representations learned by fine-tuning and prompting yield similar decoding performance. For both prompt-tuning and fine-tuning, representations learned from tasks dealing with concept meaning (CoRef, WSD and NER) deliver impressive decoding accuracy across the 4 functional networks. On decoding the Language and Semantic network, learning the CoRef and WSD task yields average decoding accuracies higher than other tasks except NER across subjects for both types of tuned representations. The difference between WSD and CoRef is not consistently significant ($p > 10^{-2}$). On decoding the visual network and DMN, learning the CoRef task yields the best average decoding accuracy.

5.2 Decoding Within Regions of Interests

In this subsection, we present the decoding performance of fine-tuned and prompt-tuned representations on the ROIs comprising the language network.⁵ We conduct neural decoding on each ROI for all the subjects and show the averaged decoding performance among the subjects in Figure 4.

Though fine-tuning on CoRef and WSD yields representations that better decode some ROIs than prompt-tuning, after pooling across all the tasks, we find that on all the ROIs prompt-tuning exceeds fine-tuning in average decoding accuracy. For both fine-tuned and prompt-tuned representations, left posterior temporal cortex (LPostTemp) is the ROI where the highest average decoding accuracy is observed while left middle frontal gyrus (LMFG) is the lowest. Actually, if we rank different ROIs according to the average decoding accuracy, fine-tuning and prompt-tuning produce the same rank which is also consistent with the decoding performance of the naive BERT representations. We also find that on different ROIs, the supervised representations show different decoding performance patterns. For example for the fine-tuned representations, the CoRef task representations yield the best decoding accuracy on most language network ROIs. But for the prompt-tuned embeddings, the WSD task representations outperform.

5.3 Representational Similarity of Supervised Embeddings

In this subsection, we dive deep into the similarity relationship of supervised representations. We calculate the Pearson’s correlation between each two of the task-supervised representations and depict the correlation as heatmaps in Fig-

⁵There are 8 ROIs comprising the language network and their full names are detailed in the appendix.

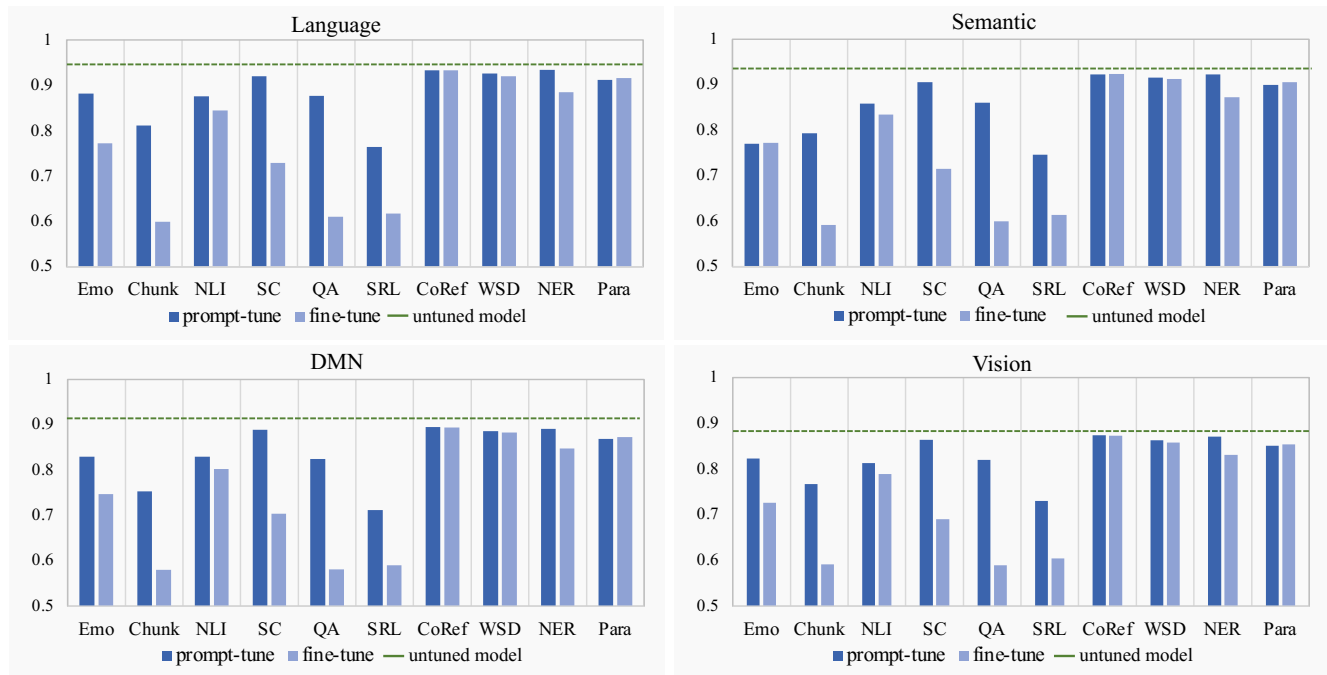


Figure 3: Decoding accuracy evaluated by pairwise classification with task-supervised representations on the brain language networks, the ensemble semantic system, DMN and visual networks. The green dotted line denotes the decoding performance of the original untuned BERT representations.

ure 5[a] (prompt-tuned representations) and Figure 5[b] (fine-tuned representations).

As shown in Figure 5, the fine-tuned representations largely diverge in similarity with the naive untuned BERT representations, while prompt-tuned representations correlate closer to the untuned representations. Representations fine-tuned on the QA, Chunk and SRL task even show zero or negative correlation with the untuned representations, in line with previous work [Kumar *et al.*, 2022] which proved fine-tuning can distort pre-trained features. Neural decoding accuracies with fine-tuning on the QA, Chunk and SRL tasks are also the lowest three observed in previous experiments, as depicted in Figure 3 and 4. Both fine-tuning and prompt-tuning on the CoRef task yield supervised representations that accurately decode brain activities. The CoRef representations with fine-tuning and prompt-tuning also show top correlations with the untuned model.

To gain a deeper insight into the relationship among the supervised representations, we conduct a hierarchical clustering and illustrate the results as dendrograms in Figure 5[c] and 5[d]. For both fine-tuning and prompt-tuning, CoRef representations share one cluster with the naive untuned representations. Chunk and SRL representations are closely clustered, distant from the naive and CoRef representations in the hierarchy. We further calculate the variance of neural decoding accuracies explained by the correlation between each of the supervised representations and the naive BERT representation. We find that such correlation explains 35.35% variance of prompt-tuned representations’ neural decoding accuracy within the brain language network, and explains 44.22%

variance of neural decoding accuracy within the ensemble semantic system.

6 Discussion

In this section we summarize the experimental findings and discuss possible explanations for the findings. First, we find that decoding performance with fully fine-tuned representations do not significantly exceed prompt-tuned representations on decoding any of the 4 functional network voxels. Full fine-tuning updates the entire set of model parameters to fit on a downstream NLU task, which has been proven to distort the pre-trained features. However, the human brain does not need to reformulate an entire functional network to learn a single new task [Cole *et al.*, 2013; Fedorenko *et al.*, 2011]. Learning from a new task also does not necessarily lead to distortion of previously acquired knowledge [Parisi *et al.*, 2019]. In this sense, full fine-tuning is cognitively inconsistent with the human brain’s mechanism of language representation. Prompt-tuning, in contrast with fine-tuning, freezes the pre-trained weights and learns only trainable embeddings to guide the base model to fit on a downstream task. As hypothesized, prompt-tuned embeddings better decode the information represented in the human brain for language understanding than fine-tuned peers.

Second, we find that for both fine-tuning and prompt-tuning, representations tuned on the CoRef and WSD task yield higher average decoding accuracy than other tasks, while the Chunk task yields the lowest. The CoRef and WSD tasks require complex reasoning involving concept meaning and semantic relation, while the Chunk task requires mod-




Figure 4: Decoding accuracy measured by pairwise matching with ROIs comprising the Language network. *FT* and *PT* respectively denote fine-tuning and prompt-tuning. *mean* denotes the average score of all the decoding accuracies with fine-tuned or prompt-tuned representations.

eling the shallow syntactic structure. This indicates that deep level semantic information could be involved in human brain’s language representation, possibly taking a larger proportion than the shallow syntactic information. We also find that the neural decoding performance of the prompt-tuned representations can be at least partially explained by their correlation with the untuned naive model. This means that general domain knowledge learned by the pre-trained model plays an important role in decoding the brain representations of language.

Third, we observe that representations tuned on half of the evaluated tasks do not significantly improve over the untuned model in neural decoding accuracy. Learning task-specific features leads to updating or even reformation of the general domain knowledge acquired in pre-training, which could be important for decoding human language representations. If the task-specific feature learned from tuning is not as important as the general domain knowledge in decoding brain representations, we can expect tuning yields decreased decoding accuracy. Since the tasks are from popular NLP benchmarks and the models are tuned with the hyper-parameter search to achieve their optimal task performance, this finding also suggests these benchmarks might not robustly and accurately reflect how well the tuned models resemble the brain’s mechanism for language understanding. It then urges further studies to build benchmarks that assess the cognitive and biological plausibility of the task-supervised models.

7 Conclusion

In this study, for the first time we propose to link the human brain and task-supervised artificial representations by

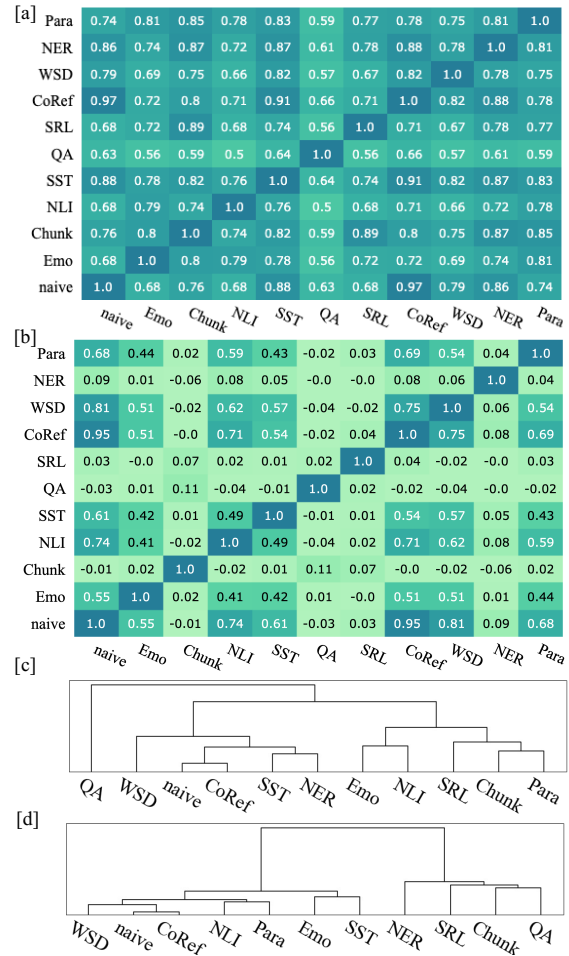


Figure 5: [a] and [b]: Pearson correlation between each two of the supervised representations, [a] for the prompt-tuned representations and [b] for the fine-tuned representations. [c] and [d]: Dendrograms constructed with hierarchical clustering on the Pearson correlations in [a] and [b], [c] for prompt-tuned representations and [d] for the fine-tuned representations.

prompt-tuning. We take 10 NLU tasks involving diverse linguistic-cognitive skills from widely used NLU benchmarks. A pre-trained BERT model is fitted on the 10 tasks with full fine-tuning and prompt-tuning, respectively. The supervised representations generated by the tuned models are then employed to decipher the brain activation patterns through neural decoding. We comprehensively compare the prompt-tuned supervised representations against the fine-tuned peers on 4 functional networks and their comprising ROIs. The comparison demonstrates that prompt-tuned representations could better decode the neural responses to a linguistic stimulus than fully fine-tuned representations in most cases. We also demonstrate how task representations differently decode brain activities. We believe that the proposed prompt-tuning-based framework and experimental findings will help to better understand the relationship between supervised artificial and brain language representations.

Acknowledgements

This work is funded by the CALCULUS project (European Research Council Advanced Grant H2020-ERC-2017-ADG 788506). We thank Dr. Shaonan Wang very much for proof-reading.

References

- [Anderson *et al.*, 2016] Andrew James Anderson, Jeffrey R Binder, Leonardo Fernandino, Colin J Humphries, Lisa L Conant, Mario Aguilar, Xixi Wang, Donias Doko, and Rajeev DS Raizada. Predicting neural activity patterns associated with sentences using a neurobiologically motivated model of semantic representation. *Cerebral Cortex*, 27(9):4379–4395, 2016.
- [Binder *et al.*, 2009] Jeffrey R Binder, Rutvik H Desai, William W Graves, and Lisa L Conant. Where is the semantic system? a critical review and meta-analysis of 120 functional neuroimaging studies. *Cerebral Cortex*, 19(12):2767–2796, 2009.
- [Brown *et al.*, 2020] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020.
- [Buckner *et al.*, 2008] Randy L Buckner, Jessica R Andrews-Hanna, and Daniel L Schacter. The brain’s default network. *Annals of the New York Academy of Sciences*, 1124(1):1–38, 2008.
- [Caucheteux *et al.*, 2021] Charlotte Caucheteux, Alexandre Gramfort, and Jean-Remi King. Model-based analysis of brain activity reveals the hierarchy of language in 305 subjects. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3635–3644, Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics.
- [Chen *et al.*, 2022] Xiang Chen, Ningyu Zhang, Xin Xie, Shumin Deng, Yunzhi Yao, Chuanqi Tan, Fei Huang, Luo Si, and Huajun Chen. Knowprompt: Knowledge-aware prompt-tuning with synergistic optimization for relation extraction. In *Proceedings of the ACM Web Conference 2022*, pages 2778–2788, 2022.
- [Cole *et al.*, 2013] Michael W Cole, Jeremy R Reynolds, Jonathan D Power, Grega Repovs, Alan Anticevic, and Todd S Braver. Multi-task connectivity reveals flexible hubs for adaptive task control. *Nature neuroscience*, 16(9):1348–1355, 2013.
- [Devlin *et al.*, 2018] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, 2018.
- [Fedorenko *et al.*, 2011] Evelina Fedorenko, Michael K Behr, and Nancy Kanwisher. Functional specificity for high-level linguistic processing in the human brain. *Proceedings of the National Academy of Sciences*, 108(39):16428–16433, 2011.
- [Fyshe, 2015] Alona Fyshe. *Corpora and Cognition: The Semantic Composition of Adjectives and Nouns in the Human Brain*. PhD thesis, Doctoral dissertation, Air Force Research Laboratory, 2015.
- [Gauthier and Levy, 2019] Jon Gauthier and Roger Levy. Linking artificial and human neural representations of language. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP) and the 9th International Joint Conference on Natural Language Processing*, 2019.
- [Handjaras *et al.*, 2016] Giacomo Handjaras, Emiliano Ricciardi, Andrea Leo, Alessandro Lenci, Luca Cecchetti, Mirco Cosottini, Giovanna Marotta, and Pietro Pietrini. How concepts are encoded in the human brain: a modality independent, category-based cortical organization of semantic knowledge. *Neuroimage*, 135:232–242, 2016.
- [Huth *et al.*, 2016] Alexander G Huth, Wendy A de Heer, Thomas L Griffiths, Frédéric E Theunissen, and Jack L Gallant. Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature*, 532(7600):453, 2016.
- [Kojima *et al.*, 2022] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- [Kumar *et al.*, 2022] Ananya Kumar, Aditi Raghunathan, Robbie Matthew Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pretrained features and underperform out-of-distribution. In *International Conference on Learning Representations*, 2022.
- [Lester *et al.*, 2021] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics.
- [Li and Liang, 2021] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Vol-*

- ume 1: Long Papers), pages 4582–4597, Online, August 2021. Association for Computational Linguistics.
- [Liu *et al.*, 2022a] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Comput. Surv.*, sep 2022. Just Accepted.
- [Liu *et al.*, 2022b] Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. P-tuning: Prompt tuning can be comparable to fine-tuning across scales and tasks. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 61–68, Dublin, Ireland, May 2022. Association for Computational Linguistics.
- [Matsuo *et al.*, 2016] Eri Matsuo, Ichiro Kobayashi, Shinji Nishimoto, Satoshi Nishida, and Hideki Asoh. Generating natural language descriptions for semantic representations of human brain activity. *Proceedings of the ACL 2016 Student Research Workshop*, pages 22–29, 2016.
- [Oota *et al.*, 2018] Subba Reddy Oota, Naresh Manwani, and Raju S Bapi. fmri semantic category decoding using linguistic encoding of word embeddings. In *International Conference on Neural Information Processing*, pages 3–15. Springer, 2018.
- [Oota *et al.*, 2022] Subba Reddy Oota, Jashn Arora, Veeral Agarwal, Mounika Marreddy, Manish Gupta, and Bapi Surampudi. Neural language taskonomy: Which NLP tasks are the most predictive of fMRI brain activity? In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3220–3237, Seattle, United States, July 2022. Association for Computational Linguistics.
- [Ouyang *et al.*, 2022] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Proceedings of Advances in neural information processing systems (NeurIPS)*, 2022.
- [Parisi *et al.*, 2019] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019.
- [Pereira *et al.*, 2011] Francisco Pereira, Greg Detre, and Matthew Botvinick. Generating text from functional brain images. *Frontiers in human neuroscience*, 5:72, 2011.
- [Pereira *et al.*, 2018] Francisco Pereira, Bin Lou, Brianna Pritchett, Samuel Ritter, Samuel J Gershman, Nancy Kanwisher, Matthew Botvinick, and Evelina Fedorenko. Toward a universal decoder of linguistic meaning from brain activation. *Nature communications*, 9(1):963, 2018.
- [Qin and Eisner, 2021] Guanghui Qin and Jason Eisner. Learning how to ask: Querying LMs with mixtures of soft prompts. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5203–5212, Online, June 2021. Association for Computational Linguistics.
- [Qiu *et al.*, 2020] Xipeng Qiu, Tianxiang Sun, Yige Xu, Yunfan Shao, Ning Dai, and Xuanjing Huang. Pre-trained models for natural language processing: A survey. *Science China Technological Sciences*, 63(10):1872–1897, 2020.
- [Radford *et al.*, 2019] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9, 2019.
- [Sarlin *et al.*, 2020] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4938–4947, 2020.
- [Schrimpf *et al.*, 2021] Martin Schrimpf, Idan Asher Blank, Greta Tuckute, Carina Kauf, Eghbal A Hosseini, Nancy Kanwisher, Joshua B Tenenbaum, and Evelina Fedorenko. The neural architecture of language: Integrative modeling converges on predictive processing. *Proceedings of the National Academy of Sciences*, 118(45):e2105646118, 2021.
- [Sun *et al.*, 2019] Jingyuan Sun, Shaonan Wang, Jiajun Zhang, and Chengqing Zong. Towards sentence-level brain decoding with distributed representations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7047–7054, 2019.
- [Sun *et al.*, 2020a] Jingyuan Sun, Shaonan Wang, Jiajun Zhang, and Chengqing Zong. Distill and replay for continual language learning. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3569–3579, Barcelona, Spain (Online), December 2020. International Committee on Computational Linguistics.
- [Sun *et al.*, 2020b] Jingyuan Sun, Shaonan Wang, Jiajun Zhang, and Chengqing Zong. Neural encoding and decoding with distributed sentence representations. *IEEE Transactions on Neural Networks and Learning Systems*, 32(2):589–603, 2020.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Proceedings of Advances in neural information processing systems (NIPS)*, 2017.
- [Wang *et al.*, 2018] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium, November 2018. Association for Computational Linguistics.
- [Wang *et al.*, 2022] Shaonan Wang, Yunhao Zhang, Xiaohan Zhang, Jingyuan Sun, Nan Lin, Jiajun Zhang, and Chengqing Zong. An fmri dataset for concept representation with semantic feature annotations. *Scientific Data*, 9(1):721, 2022.