

# Do Large Language Models Mirror Cognitive Language Processing?

Yuqi Ren<sup>1</sup>, Renren Jin<sup>1</sup>, Tongxuan Zhang<sup>2</sup>, Deyi Xiong<sup>1\*</sup>

<sup>1</sup> College of Intelligence and Computing, Tianjin University

<sup>2</sup> College of Computer and Information Engineering, Tianjin Normal University

{ryq20, rrjin, dyxiong}@tju.edu.cn

txzhang@tjnu.edu.cn

## Abstract

Large Language Models (LLMs) have demonstrated remarkable abilities in text comprehension and logical reasoning, indicating that the text representations learned by LLMs can facilitate their language processing capabilities. In neuroscience, brain cognitive processing signals are typically utilized to study human language processing. Therefore, it is natural to ask how well the text embeddings from LLMs align with the brain cognitive processing signals, and how training strategies affect the LLM-brain alignment? In this paper, we employ Representational Similarity Analysis (RSA) to measure the alignment between 23 mainstream LLMs and fMRI signals of the brain to evaluate how effectively LLMs simulate cognitive language processing. We empirically investigate the impact of various factors (e.g., pre-training data size, model scaling, alignment training, and prompts) on such LLM-brain alignment. Experimental results indicate that pre-training data size and model scaling are positively correlated with LLM-brain similarity,<sup>1</sup> and alignment training can significantly improve LLM-brain similarity. Explicit prompts contribute to the consistency of LLMs with brain cognitive language processing, while nonsensical noisy prompts may attenuate such alignment. Additionally, the performance of a wide range of LLM evaluations (e.g., MMLU, Chatbot Arena) is highly correlated with the LLM-brain similarity.

## 1 Introduction

Large language models, e.g., chatGPT (OpenAI, 2023), have demonstrated linguistic capabilities that progressively approach human-level language comprehension and generation. The text representations learned by LLMs transform text into a high-dimensional semantic space (Mikolov et al., 2013).

\* Corresponding author.

<sup>1</sup>For notational simplicity, we refer to the similarity between LLM representations and brain cognitive language processing signals as LLM-brain similarity.

In neuroscience, brain cognitive processing signals, e.g., EEG, fMRI signals, have been widely acknowledged to be able to represent a snapshot of cognitive representations of language in human brains (Xu et al., 2016). Such signals record human brain activities during the cognitive processing of language and play a key role in revealing the internal workings of human language cognition (Friederici, 2011). Thus, it is natural to ask: how “similar” are these two different modalities of language representations (LLMs representations and cognitive language processing signals in the brain)? Or do LLMs mirror human cognitive language processing? Answering these questions might facilitate us to interpret LLMs from the perspective of neuroscience.

Previous studies have explored the assessment of consistency between representations learned by pre-trained language models and cognitive processing signals (Xu et al., 2016; Hollenstein et al., 2019; Abnar et al., 2019). Along this research line, recent studies compare human neuroimaging data with representations learned by 43 pre-trained language models that are smaller than 1B-sized LLMs, revealing that language models with stronger predictive abilities for the next word exhibit greater alignment with the human brain (Schrimpf et al., 2021). However, current LLMs have scaled up their model size by hundreds of times (e.g. LLaMa3-400B+ (Meta, 2024)), and been strengthened by alignment training methods such as supervised fine-tuning (SFT) and reinforcement learning with human feedback (RLHF) (Ouyang et al., 2022). These factors contribute significantly to the enhanced capabilities of LLMs (Wei et al., 2022; Zheng et al., 2023b). Nevertheless, it is yet to be tested whether the latest LLMs with scales of 7B and above maintain consistency with human cognitive data. Additionally, the impact of pre-training data size, alignment training, prompts, and other training strategies on LLM-brain similarity has not been investigated.

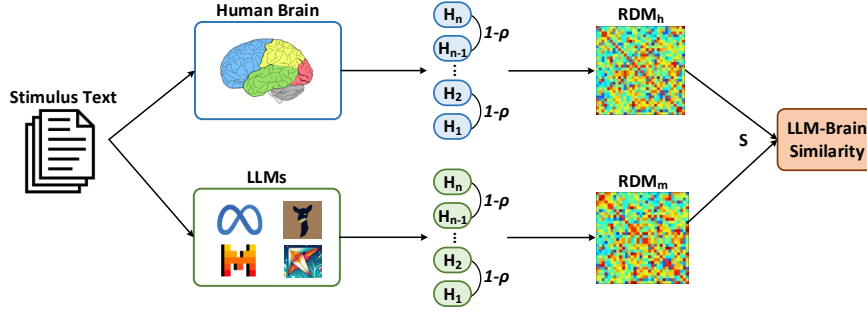


Figure 1: Diagram of the proposed LLM-brain similarity estimation framework. ‘H’ denotes sentence representations from different modalities. ‘ $\rho$ ’ denotes the pearson correlation coefficient. ‘S’ denotes similarity measurement method.

This paper aims to assess the extent to which LLMs mirror cognitive language processing activity and subsequently analyze the impacts of different factors (e.g., pre-training data size, model scaling, alignment training) on such LLM-brain alignment. Specifically, we adopt functional magnetic resonance imaging (fMRI) signals as human cognitive language processing signals, and RSA (Kriegeskorte et al., 2008), widely used in neuroscience (Connolly et al., 2012; Diedrichsen and Kriegeskorte, 2017), to compute similarity scores between LLM representations and human cognitive language processing signals. The framework for LLM-brain similarity computation is illustrated in Figure 1.

Our experiments comprise 23 open-source LLMs, trained with various pre-training data size, scaling and training strategies (i.e., pre-training, SFT, RLHF), to investigate the impact of these factors on LLM-brain similarity. We further analyze the impact of prompt strategy on the human intention understanding abilities of LLMs and the consistency between LLM sentiment and human sentiment. Additionally, we investigate the relationship between the capability of LLM evaluations and LLM-brain similarity.

Our main contributions can be summarized as follows.

- We employ human cognitive language processing signals to assess the resemblance between LLMs and human cognition. By analyzing the similarity of 23 LLMs to human cognitive language processing, we investigate the impact of various factors, such as pre-training data size, model scaling, alignment training and prompt strategy, on the LLM-brain similarity.

- We find that the explicit-prompts result in a higher LLM-brain similarity compared to no-prompts, indicating that explicit prompts can enhance the alignment of LLMs with human intentions.
- Our findings indicate that LLMs exhibit high similarity to human cognitive language processing in positive sentiment, suggesting that LLMs may encode more positive sentimental text during training.
- We observe a high degree of consistency between the performance of a wide range of LLM evaluations (e.g., MMLU, Chatbot Arena) and LLM-brain similarity, indicating that the LLM-brain similarity holds substantial potential to evaluate the capabilities of LLMs.

## 2 Related Work

**Neuroscience-Inspired NLP** Human brain contains intricate neural networks and hierarchical structures, that encompass a broader cognitive complexity and generalization, providing a more comprehensive and profound perspective to understand the performance of LLMs in language cognition (Denève et al., 2017; Thibaut de Schotten and Forkel, 2022; Ren and Xiong, 2021). Thus, aligning brain cognitive processing signals with text embeddings to evaluate LLMs has remained a prominent area of research. Numerous studies explore how embedding representations of words and sentences extracted from NLP models correspond to fMRI or MEG recordings (Mitchell et al., 2008; Pereira et al., 2018). Hollenstein et al. (2019) propose a framework for cognitive word embedding evaluation called CogniVal. This framework as-

asses the extent to which language model representations reflect semantic information in the human brain across three modalities: eye-tracking, EEG, and fMRI. Xu et al. (2016) introduce a lightweight tool named BrainBench, designed to evaluate the word semantics in distributional models, using fMRI images corresponding to 60 concrete nouns. Toneva and Wehbe (2019) interpret the differences among 4 language models concerning layer depth, context length, and attention types using fMRI data. By leveraging insights from attention experiments, the performance of syntax-related tasks is enhanced.

**NLP-Inspired Neuroscience** Leveraging the superior semantic expressiveness of text embeddings from pre-trained models, a few studies utilize the computational mapping between brain activity and text embeddings to decode human brains (Mitchell et al., 2008; Pereira et al., 2018; Ren and Xiong, 2022). Sun et al. (2019) have explored the effectiveness of different types of distributed representation models in sentence-level brain decoding based on fMRI, and found that supervised structured models exhibited remarkable performance in brain decoding. Additionally, Wehbe et al. (2014) drew an analogy between the recurrent neural network language model and the working mechanisms of the brain during discourse comprehension. They utilize the predictive performance of magnetoencephalography on hidden states at different moments to reveal what information is encoded by brain regions at specific times.

### 3 Methodology

To measure the extent to which LLMs align with human cognitive language processing, we employ the RSA (Kriegeskorte et al., 2008) method to assess the similarity in representations between human brains and LLMs when presented with the same stimulus text. Furthermore, we investigate how LLM-brain similarity varies with different prompts and analyze resemblances between LLMs and human cognitive language processing regarding sentimental polarity using crafted prompts.

#### 3.1 LLM-Brain Similarity Estimation

RSA is a widely-used technique for comparing cross-modal representation spaces and has found extensive applications in neuroscience (Connolly et al., 2012; Diedrichsen and Kriegeskorte, 2017). In NLP, this method is utilized to calculate the

consistency between two different representation spaces encoding information (Abdou et al., 2019; Chrupała and Alishahi, 2019). Moreover, it can be employed to quantify the relationship between human brains and neural network models (King et al., 2019).

We hence use RSA to measure the LLM-brain similarity. Firstly, we compute the representation dissimilarity matrix (RDM) separately for brain cognitive processing signals and LLMs embeddings. For each pair of stimulus sentences, we calculate the 1 - Pearson correlation coefficient ( $\rho$ ) of their representations (either LLM representations or brain cognitive processing signals) as dissimilarity, following the previous works (Dwivedi and Roig, 2019; Luo et al., 2022). Consequently, both the human brain ( $RDM_h$ ) and the LLM ( $RDM_m$ ) yield a diagonal symmetric dissimilarity matrix of size  $n * n$ , where  $n$  is the number of stimulus sentences. We formulate it as follows:

$$RDM = \begin{bmatrix} 0 & \cdots & 1 - \rho(H_1, H_n) \\ \vdots & \ddots & \vdots \\ 1 - \rho(H_n, H_1) & \cdots & 0 \end{bmatrix} \quad (1)$$

We then estimate the similarity between  $RDM_m$  and  $RDM_h$  using the Pearson correlation coefficient, Euclidean distance, Cosine distance and Spearman correlation coefficient. In addition to Euclidean distance, a high correlation score implies that representations learned by the LLM are highly aligned with brain cognitive processing signals.

$$Sim = S(RDM_h, RDM_m) \quad (2)$$

#### 3.2 Prompt Strategy

Instruction tuning significantly enhances the in-context learning ability of LLMs by understanding task requirements well through task definitions and corresponding examples (Ouyang et al., 2022; Peng et al., 2023). We aim to verify the effect of prompts on human intentions by examining whether appended prompts increase LLM-brain similarity.

To assess the sensitivity of the LLM-brain similarity to appended prompts, we compare the effects of three conditions: no prompts (vanilla input), explicit prompts, and noisy prompts. No-prompt-appending involves directly feeding text into LLMs, allowing the LLM to automatically continue the text without any prompts. To evaluate the effect of prompt on the same task, we prepend the phrase "Please complete the following text:" to the input

Positive	Negative
great	war
charming	sin
successful	stupid
pleasure	prison
laugh	pain
elegance	liar
kindness	angry
smiling	damage
accomplished	sad
impress	poor

Table 1: 10 positive concepts vs. 10 negative concepts, serving as proxies to positive and negative sentiment.

text as an explicit prompt. Furthermore, we introduce a noisy prompt consisting of 5 randomly sampled English words, e.g., “Harmony illuminate umbrella freedom like.” to analyze the impact of nonsensical prompts under the same length.

### 3.3 Sentimental Polarities

LLMs have demonstrated significant achievements in value alignment (Han, 2023; Ganguli et al., 2023), with sentimental expressions emerging as a prominent feature in their outputs. Therefore, it is important to analyze the consistency of LLMs and human sentimental tendencies to assess the alignment ability of LLMs. Current studies typically utilize psychological scales to analyze the similarity of sentimental expression between LLMs and humans (Huang et al., 2023; Wang et al., 2023). In this paper, we focus on the intrinsic representations of sentimental expressions in LLMs. We delve into the similarity between the embeddings of sentimental concepts learned by LLMs and brain signals under the same concepts to investigate the consistency between LLMs and humans across different sentimental polarities. It’s worth noting that brain signals, which record physiological responses accompanied by emotions, are commonly used in psychology to analyze human sentiment (Davidson, 2003).

Specifically, we determine the sentimental polarities (i.e., positive vs. negative) of 180 concepts in fMRI stimulus texts by the sentiment analysis tool VADER (Valence Aware Dictionary and sentiment Reasoner) included in NLTK (Bird et al., 2009). After inputting a concept into the VADER, it will return a sentimental polarity score ranging from -1 (negative) to +1 (positive), with 0 indicating neutrality. Subsequently, we select the top 10

concepts with the highest positive sentiment and the top 10 concepts with the highest negative sentiment to serve as proxies for positive and negative sentiment, as shown in Table 1. To obtain LLM representations for the selected concepts, we generate prompts using the template “The sentiment of the word <concept> is”, focusing on extracting embeddings solely for the “<concept>” token. The similarity score for positive sentiment and negative sentiment are computed by averaging the LLM-brain similarity scores across all concepts within the each polarity.

## 4 Experiments

We conducted experiments on 23 mainstream LLMs to investigate the alignment between LLM representations and brain signals, and the impact of various factors on this alignment. Additionally, we studied the consistency between LLMs and human cognitive language processing regarding sentimental polarity.

### 4.1 Cognitive Language Processing Data

fMRI signals map brain activity by detecting changes associated with blood flow. The fMRI data utilized in our experiments was from Pereira (Pereira et al., 2018).<sup>2</sup> These fMRI signals are recorded on a whole-body 3-Tesla Siemens Trio scanner with a 32-channel head coil. The fMRI data consist of neural activity representations in cubic millimeter-sized voxels, with each fMRI recording containing a substantial number of voxels. This dataset comprises three experiments: in Experiment 1, 16 subjects read 180 concepts; in Experiment 2, 9 subjects read 384 sentences; and in Experiment 3, 6 subjects read 243 sentences. However, only 5 subjects participated in all three experiments. To ensure the fairness of our experimental results, we selected the fMRI signals from these 5 subjects.

We flattened the 3-dimensional fMRI images into 1-dimensional vectors and randomly selected 1,000 voxels as the brain cognitive language processing signal for each subject. The stimulus texts presented to subjects consist of 180 isolated stimuli (concepts) and 627 continuous stimuli (natural language sentences). Due to individual variations among human subjects, we calculated LLM-brain similarity separately for each individual. The overall similarity scores were derived by averaging the individual scores.

<sup>2</sup>The dataset is publicly available at <https://osf.io/crwz7/>.

## 4.2 Large Language Models

We used 23 LLMs from 8 LLM families.

**Amber** is a 7B LLM pre-trained on 1.3 TB of tokens, with all training codes, training parameters, and system configurations open-sourced to enhance the reproducibility and scalability of LLMs (Liu et al., 2023). Amber evenly divides the pre-training data into 360 data chunks, saving a checkpoint after training each chunk. Each checkpoint involves training the same amount of data as the previous one, with no differences in the number of training epochs. Additionally, several fine-tuned versions, including Amberchat and Ambersafe, have been released. Amberchat is fine-tuned using instruction data from WizardLM (Xu et al., 2023), while Ambersafe conducts direct parameter optimization (DPO) based on Amberchat to align with human preferences. In this paper, we evenly select 10 checkpoints from the publicly available 360 checkpoints to investigate the impact of pre-training data size on LLM-brain similarity.

**LLaMA** is a series of open-source and powerful base language models, trained on trillions of tokens from publicly available datasets. We selected the 7B and 13B variants of LLaMA to investigate the impact of model scaling on LLM-brain similarity.

**LLaMA2** is a collection of pre-trained and fine-tuned LLMs, with model size ranging from 7B to 70B (Touvron et al., 2023). Compared to LLaMA, the pre-training data size has increased by 40%, and the context length has expanded from 2k to 4k. The fine-tuned version, LLaMA2-Chat, aligns with human preferences through SFT and RLHF. In this paper, we selected LLMs with sizes of 7B, 13B, and 70B from the LLaMA2 family.

**LLaMA3** is an improved version of LLaMA2, available in both pre-trained and instruction-tuned versions (Meta, 2024). Compared to LLaMA2, LLaMA3 has made significant advancements, including the number of trained tokens has increased from 2TB to 15TB, the context window size has expanded from 4096 to 8192, and the vocabulary has enlarged from 32,000 to 128,000. During fine-tuning stage, LLaMA3 integrates SFT, rejection sampling, proximal policy optimization (PPO), and DPO. These enhancements have led to substantial improvements in inference capabilities, code generation, and instruction following. For this paper, we employed both 8B and 70B scales for the LLaMA3

family.

**Vicuna** To investigate the effect of SFT on LLMs, we employed the Vicuna family (Zheng et al., 2023a). Vicuna v1.3 is fine-tuned from LLaMA using a training dataset comprising 125k user-shared conversations collected from ShareGPT.com. Additionally, Vicuna v1.5 is fine-tuned from LLaMA2, with the same training dataset as Vicuna v1.3. Both versions of Vicuna are available in 7B and 13B scales.

**Mistral-7B** is a LLM with 7B parameters that outperforms the best open-source 13B model (LLaMA2) across multiple evaluation benchmarks. This model employs grouped query attention to achieve faster inference and utilizes sliding window attention to extend the processed sequence length (Jiang et al., 2023).

**Mistral-7B-sft** comprises SFT versions of Mistral-7B. Mistral-7B-sft-alpha is fine-tuned on the UltraChat dataset (Ding et al., 2023), which consists of 1.47M multi-turn dialogues generated by GPT-3.5-Turbo. Similarly, Mistral-7B-sft-beta is fine-tuned on a cleaned version of the UltraChat dataset, containing 200k samples.

**Zephyr-7B** is built upon Mistral-7B-sft-beta and trained using distilled supervised fine-tuning (dSFT) and distilled direct preference optimization (dDPO) to better alignment with human intentions in interactions. This model demonstrates comparable performance to a 70B chat model aligned using human feedback (Tunstall et al., 2023).

The average embeddings of all tokens from the final layer of LLMs were utilized as sentence representations. All the LLMs used in this paper are available from <https://huggingface.co/models>.

## 4.3 Results

The LLM-brain similarity results for 23 LLMs are shown in Table 2. We observe that under different similarity measurement methods, the trend of LLM-brain similarity remains roughly consistent. This indicates that our method demonstrates robustness across different similarity evaluation methods.

**Pre-training data size** The scaling law of LLMs plays a crucial role in the performance of LLMs (Kaplan et al., 2020; Hoffmann et al., 2022). Hoffmann et al. (2022) suggest that with a fixed scaling of LLMs, the loss gradually decreases as the

Scaling	LLM	Training Stage	Pearson	Euclidean	Cosine	Spearman
7B	Amber	Pre-training	0.2038 ± 0.013	128.26 ± 7.45	0.9785 ± 0.002	0.1440 ± 0.031
	Amberchat	SFT	0.2263 ± 0.015	105.41 ± 7.62	0.9831 ± 0.003	0.1477 ± 0.038
	Ambersafe	SFT+DPO	0.2302 ± 0.014	100.51 ± 6.93	0.9856 ± 0.003	0.1597 ± 0.032
	LLaMA-7B	Pre-training	0.2103 ± 0.016	108.17 ± 6.88	0.9773 ± 0.001	0.1565 ± 0.033
	Vicuna-7B-v1.3	SFT	0.2217 ± 0.012	99.62 ± 7.34	0.9796 ± 0.005	0.1824 ± 0.030
	LLaMA2-7B	Pre-training	0.2207 ± 0.016	105.12 ± 7.81	0.9800 ± 0.002	0.1653 ± 0.030
	Vicuna-7B-v1.5	SFT	0.2346 ± 0.014	98.40 ± 6.95	0.9850 ± 0.004	0.1890 ± 0.031
	LLaMA2-7B-chat	SFT+RLHF	0.2410 ± 0.017	95.12 ± 7.44	0.9876 ± 0.003	0.1948 ± 0.034
	Mistral-7B	Pre-training	0.2533 ± 0.013	100.53 ± 7.29	0.9767 ± 0.001	0.1730 ± 0.037
	Mistral-7B-sft-alpha	SFT	0.2481 ± 0.016	93.40 ± 7.53	0.9808 ± 0.003	0.1769 ± 0.030
Mistral-7B-sft-beta	SFT	0.2573 ± 0.015	91.07 ± 7.66	0.9834 ± 0.003	0.1896 ± 0.032	
Zephyr-7B	dSFT+dDPO	0.2670 ± 0.014	89.24 ± 7.35	0.9876 ± 0.006	0.1945 ± 0.034	
8B	LLaMA3-8B	Pre-training	0.2540 ± 0.015	99.95 ± 7.62	0.9795 ± 0.005	0.1799 ± 0.035
	LLaMA3-8B-instruct	SFT+PPO+DPO	0.2697 ± 0.014	87.39 ± 7.11	0.9859 ± 0.002	0.1932 ± 0.034
13B	LLaMA-13B	Pre-training	0.2221 ± 0.017	105.78 ± 6.86	0.9832 ± 0.004	0.1783 ± 0.033
	Vicuna-13B-v1.3	SFT	0.2427 ± 0.016	94.77 ± 7.49	0.9853 ± 0.005	0.1850 ± 0.037
	LLaMA2-13B	Pre-training	0.2340 ± 0.013	102.50 ± 7.59	0.9840 ± 0.002	0.1799 ± 0.032
70B	Vicuna-13B-v1.5	SFT	0.2563 ± 0.012	93.68 ± 7.05	0.9851 ± 0.005	0.1776 ± 0.031
	LLaMA2-13B-chat	SFT+RLHF	0.2694 ± 0.013	90.89 ± 7.64	0.9869 ± 0.006	0.1831 ± 0.031
	LLaMA2-70B	Pre-training	0.2568 ± 0.015	94.83 ± 7.08	0.9838 ± 0.002	0.1857 ± 0.034
70B	LLaMA2-70B-chat	SFT+RLHF	0.2762 ± 0.016	83.97 ± 6.98	0.9886 ± 0.003	0.1988 ± 0.033
	LLaMA3-70B	Pre-training	0.2778 ± 0.015	92.67 ± 7.46	0.9877 ± 0.003	0.1930 ± 0.032
	LLaMA3-70B-instruct	SFT+PPO+DPO	0.2938 ± 0.012	80.94 ± 7.69	0.9923 ± 0.005	0.2098 ± 0.035

Table 2: The LLM-brain similarity of 23 LLMs with standard deviation. The RSA similarity calculation method includes the Pearson correlation coefficient, Euclidean distance, Cosine similarity, and Spearman correlation coefficient.




Figure 2: The LLM-brain similarity of 10 different checkpoints on Amber. “ckpt” is the abbreviation for “checkpoint”.

amount of pre-training data increases, thereby enhancing the performance of LLMs. In this paper, we delve into how the pre-training data size affects LLM-brain similarity scores. The LLM-brain similarity scores for Amber across 10 checkpoints representing different pre-training data sizes are shown in Figure 2, with the scores calculated by Pearson correlation coefficient. It can be observed that as the size of pre-training data increases, the LLM-brain similarity scores demonstrate gradually improvement. This indirectly suggests a negative correlation between the loss of LLMs and LLM-brain similarity scores.

**Scaling of LLMs** The LLM-brain similarity scores across different sizes of LLMs are illustrated

in Table 2. We observe that the LLM-brain similarity increases with the size of LLMs. For example, when LLaMA2 is scaled from 7B to 70B, the corresponding LLM-brain similarity calculated by Pearson correlation coefficient increases from 0.2207 to 0.2568. This suggests that larger LLMs better mirror cognitive language processing of the brain than smaller LLMs.

**Alignment of LLMs** Alignment training can significantly enhance the alignment capability of LLMs with human values (Ouyang et al., 2022). Therefore, we want to explore whether LLM alignment brings changes in the representation space of LLMs to narrow the gap with representations of the brain cognitive language processing. The LLM-brain similarity results for different training stages are presented in Table 2. From these results, we observe that:

- LLMs with alignment training exhibit higher similarity to brain cognitive processing signals compared to the pre-trained versions, with LLaMA3-70B-instruct achieving the highest performance. This indicates that alignment training guides LLMs to better align with human brain.
- Some LLMs, after alignment training, approach or surpass larger pre-trained LLMs,




Figure 3: The LLM-brain similarity calculated by Pearson correlation coefficient of different prompt appending strategies on LLMs.

such as the LLaMA2-13B-chat exceeds the LLaMA2-70B by 0.0126 in Pearson correlation coefficient. This suggests that alignment training has a more significant impact on the LLM-brain similarity than model scaling.

- Mistral-7B-sft-alpha (M-alpha) and Mistral-7B-sft-beta (M-beta) are both derived from Mistral-7B through SFT. Notably, M-alpha demonstrates lower LLM-brain similarity than Mistral-7B (0.2481 vs. 0.2573). While both models utilize SFT data from the same dataset, M-beta employs a reduced dataset by 7 fold, containing only high-quality, filtered samples. This indicates that for aligning with humans, the quality of SFT data is more important than its quantity, resonating with conclusions drawn in prior research (Touvron et al., 2023).

**Sensitivity of Prompts** The LLM-brain similarity results under the three different prompt appending strategies are illustrated in Figure 3. It is evident that explicit-prompt-appending generally outperforms the other two types of prompt appending strategies. While the difference between no-prompt-appending and noisy-prompt-appending is relatively small, no-prompt-appending tends to outperform noisy-prompt-appending for most LLMs. This suggests that explicit prompts contribute to the consistency of LLMs with brain cognitive language processing, while nonsensical noisy prompt may attenuate such alignment.

Simultaneously, across these three types of appended prompts, LLMs consistently exhibit a positive correlation between scaling and LLM-brain similarity, indicating that the relationship between scaling and LLM-brain similarity remains insensitive to appended prompts. Furthermore, we explored the sensitivity of different training stages to appended prompt. We calculated the difference between the explicit-prompt-appending and




Figure 4: The LLM-brain similarity calculated by Pearson correlation coefficient of LLMs across different sentimental polarities.

no-prompt-appending, as shown in Appendix A.1. It is observed that, the alignment version of LLMs demonstrates a more significant improvement in LLM-brain similarity after introducing explicit prompts compared to the pre-trained version. This might be because prompt helps LLMs align more closely with human intentions, thereby enhancing LLMs’s sensitivity to prompts.

#### 4.4 Consistency with Human Sentiment

The pre-training data for LLMs are predominantly from the Internet, a realm abundant with a plethora of subjective texts infused with sentimental nuances. Thus, LLMs may inherently contain sentiment-yielding capabilities, and it is worth exploring whether the sentiment of LLMs is consistent with human moral sentiment. To probe the disparities in sentimental expressions between LLMs and humans, we conducted an analysis of the LLM-brain similarity across various sentimental polarities.

The experimental results are illustrated in Figure 4. We observe that the LLM-brain similarity in positive sentiment is significantly higher than that in negative sentiment. This may be due to the data governance of training data in LLMs (Touvron et al., 2023; Liu et al., 2023), such as toxicity filtering (Friedl, 2023; Gargee et al., 2022), which retains high-quality and positive sentiment text, resulting in a deeper encoding of positive sentiment compared to negative sentiment by the LLMs. Furthermore, alignment training enhances LLM-brain similarity in both positive and negative sentiments, indicating that LLMs encode human sentimental tendency when aligning with human sense of worth.

#### 4.5 Correlation with LLM Evaluations

We conducted a comparative analysis between LLM-brain similarity and performance on various

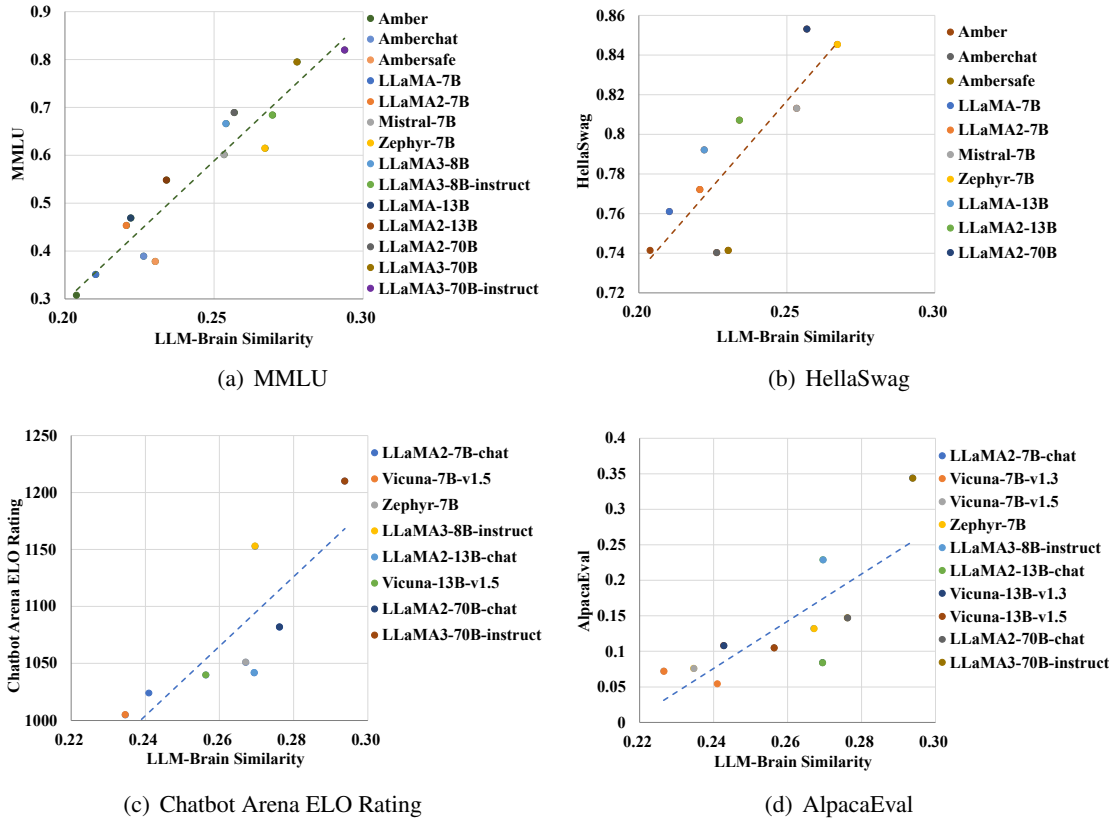


Figure 5: Correlation between the performance of LLMs on evaluations and the LLM-brain similarity calculated by Pearson correlation coefficient.

LLM evaluations. To assess the knowledge and capability, we selected two benchmarks: Massive Multilingual Language Understanding (MMLU) (Hendrycks et al., 2021) and HellaSwag (Zellers et al., 2019). MMLU spans 57 subjects, including STEM, humanities, and social sciences, designed to evaluate the knowledge reservoir of LLMs. HellaSwag requires LLMs to complete sentences, serving as an evaluation of commonsense reasoning abilities. For alignment evaluation, we employed Chatbot Arena Elo Rating (Zheng et al., 2023a), which uses crowdsourced methods for adversarial assessments of LLM outputs, and AlpacaEval 2.0 (Li et al., 2023), which adopts powerful LLMs such as GPT-4 to evaluate LLMs. All performance results of these LLM evaluations are from publicly available benchmarks.

Figures 5(a) and 5(b) respectively illustrate the correlation between LLM-brain similarity and the two knowledge and capability evaluations. The results suggest a positive correlation between LLM-brain similarity and performance on these evaluations, aligning with similar findings from earlier studies (Hollenstein et al., 2019). An intriguing

finding is that LLM-brain similarity not only assesses the consistency between LLMs and human brain but also evaluates the capabilities of LLMs.

The correlation between LLM-brain similarity and alignment evaluations is illustrated in Figures 5(c) and 5(d). The results reveal that, with the exception of LLaMA3-70B-instruct, the LLM-brain similarity and alignment capability of LLMs generally exhibit a positive correlation. LLaMA3-70B-instruct consistently achieves the highest performance in both LLM-brain similarity and alignment evaluations, significantly surpassing the linear trend line in the figures. This suggests that, although LLaMA3-70B-instruct excels in addressing artificially constructed alignment evaluations, its intrinsic representation still maintains a perceptible distance from the human brain.

#### 4.6 The Impact of LLM Layers

The representations in different layers of LLMs capture different semantic information. Intermediate layers tend to focus on abstract concepts (Jawahar et al., 2019), while the final layer is primarily responsible for predicting the next token. We further



Scaling	LLM	Training Stage	Pearson	Euclidean	Cosine	Spearman
7B	Amber	Pre-training	0.2121 ± 0.014	125.59 ± 6.96	0.9901 ± 0.001	0.1489 ± 0.029
	Amberchat	SFT	0.2426 ± 0.012	101.37 ± 7.54	0.9910 ± 0.002	0.1525 ± 0.031
	Ambersafe	SFT+DPO	0.2492 ± 0.013	98.54 ± 6.69	0.9925 ± 0.001	0.1576 ± 0.033
	LLaMA-7B	Pre-training	0.2201 ± 0.014	102.70 ± 7.15	0.9902 ± 0.002	0.1443 ± 0.032
	Vicuna-7B-v1.3	SFT	0.2359 ± 0.015	96.23 ± 6.57	0.9910 ± 0.002	0.1587 ± 0.034
	LLaMA2-7B	Pre-training	0.2278 ± 0.013	102.93 ± 6.89	0.9902 ± 0.003	0.1486 ± 0.028
	Vicuna-7B-v1.5	SFT	0.2395 ± 0.011	96.25 ± 7.15	0.9907 ± 0.001	0.1598 ± 0.030
	LLaMA2-7B-chat	SFT+RLHF	0.2663 ± 0.013	96.95 ± 7.22	0.9926 ± 0.001	0.1586 ± 0.032
	Mistral-7B	Pre-training	0.2326 ± 0.015	99.45 ± 7.64	0.9909 ± 0.002	0.1530 ± 0.031
	Mistral-7B-sft-alpha	SFT	0.2564 ± 0.012	90.74 ± 6.38	0.9908 ± 0.003	0.1589 ± 0.030
	Mistral-7B-sft-beta	SFT	0.2662 ± 0.014	89.65 ± 6.98	0.9924 ± 0.002	0.1620 ± 0.033
	Zephyr-7B	dSFT+dDPO	0.2802 ± 0.011	86.09 ± 7.14	0.9925 ± 0.003	0.1705 ± 0.031
8B	LLaMA3-8B	Pre-training	0.2634 ± 0.015	94.52 ± 7.70	0.9908 ± 0.002	0.1559 ± 0.034
	LLaMA3-8B-instruct	SFT+PPO+DPO	0.2709 ± 0.012	89.44 ± 6.71	0.9924 ± 0.001	0.1684 ± 0.031
13B	LLaMA-13B	Pre-training	0.2631 ± 0.013	104.98 ± 7.32	0.9914 ± 0.002	0.1653 ± 0.032
	Vicuna-13B-v1.3	SFT	0.2703 ± 0.012	90.60 ± 7.27	0.9914 ± 0.001	0.1659 ± 0.034
	LLaMA2-13B	Pre-training	0.2642 ± 0.014	99.72 ± 7.30	0.9916 ± 0.001	0.1659 ± 0.030
	Vicuna-13B-v1.5	SFT	0.2823 ± 0.011	91.53 ± 6.45	0.9914 ± 0.001	0.1779 ± 0.033
	LLaMA2-13B-chat	SFT+RLHF	0.2774 ± 0.012	88.67 ± 6.68	0.9926 ± 0.002	0.1867 ± 0.030
70B	LLaMA2-70B	Pre-training	0.2776 ± 0.014	90.34 ± 7.30	0.9923 ± 0.001	0.1793 ± 0.031
	LLaMA2-70B-chat	SFT+RLHF	0.3188 ± 0.012	82.70 ± 6.77	0.9929 ± 0.001	0.1982 ± 0.032
	LLaMA3-70B	Pre-training	0.2906 ± 0.015	88.23 ± 7.22	0.9921 ± 0.001	0.1837 ± 0.034
	LLaMA3-70B-instruct	SFT+PPO+DPO	0.3377 ± 0.013	75.85 ± 7.38	0.9937 ± 0.001	0.2175 ± 0.030

Table 3: The LLM-brain similarity of 23 LLMs with standard deviation, when using the intermediate layer as text representations. The RSA similarity calculation method includes the Pearson correlation coefficient, Euclidean distance, Cosine similarity, and Spearman correlation coefficient.

analyze the LLM-brain similarity for 23 LLMs using text embeddings from the intermediate layers. The intermediate layer is selected by halving the total number of layers in the corresponding LLM. For example, for a 32-layer LLM, we select the 16th layer as the intermediate layer. The results are shown in Table 3, demonstrating that as the model scaling increases and alignment training is applied, the LLM-brain similarity in the intermediate layers improves, consistent with the trend observed in the final layer. This suggests that our findings are robust across different layers of LLMs.

## 5 Conclusions

This paper has presented a framework to estimate how well LLMs mirror human cognitive language processing. We have investigated the impact of pre-training data size, model scaling, alignment training, and prompts on the LLM-brain similarity, and explored the consistency of LLMs with humans in sentimental polarities. Experimental results reveal that expanding the size of pre-training data, scaling up models, and employing alignment training contribute to enhancing LLM-brain similarity. Moreover, explicit prompts aids LLMs in understanding human intentions, and alignment training enhances the sensitivity to prompts. Notably, LLMs exhibit a stronger resemblance to humans in positive senti-

ment. The strong correlation between LLM-brain similarity and various LLM evaluations suggests that the proposed LLM-brain similarity could serve as a new way to evaluate LLMs.

## Limitations

This study, relying on text representations from open-source LLMs, is unable to assess the LLM-brain similarity of closed-source LLMs such as ChatGPT. We aim for future investigations to extend these findings to a broader spectrum of LLMs. Additionally, the fMRI stimulus texts used in this paper are exclusively in English, limiting the generalization of the observed relationships between LLMs and human cognition to other linguistic environments. Future endeavors should explore LLM-brain similarity in LLMs across diverse languages.

## Acknowledgments

The present research was supported by the National Key Research and Development Program of China (Grant No. 2023YFE0116400) and the Natural Science Foundation of China (Grant No. 62306213). We would like to thank the anonymous reviewers for their insightful comments.

## References

- Mostafa Abdou, Artur Kulmizev, Felix Hill, Daniel M Low, and Anders Søgaard. 2019. Higher-order comparisons of sentence encoder representations. *arXiv preprint arXiv:1909.00303*.
- Samira Abnar, Lisa Beinborn, Rochelle Choenni, and Willem Zuidema. 2019. Blackbox meets blackbox: Representational similarity and stability analysis of neural language models and brains. *arXiv preprint arXiv:1906.01539*.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural Language Processing with Python*. O'Reilly.
- Grzegorz Chrupała and Afra Alishahi. 2019. Correlating neural and symbolic representations of language. *arXiv preprint arXiv:1905.06401*.
- Andrew C Connolly, J Swaroop Guntupalli, Jason Gors, Michael Hanke, Yaroslav O Halchenko, Yu-Chien Wu, Hervé Abdi, and James V Haxby. 2012. The representation of biological classes in the human brain. *Journal of Neuroscience*, 32(8):2608–2618.
- Richard J Davidson. 2003. Affective neuroscience and psychophysiology: Toward a synthesis. *Psychophysiology*, 40(5):655–665.
- Sophie Denève, Alireza Alemi, and Ralph Bourdoukan. 2017. The brain as an efficient and robust adaptive learner. *Neuron*, 94(5):969–977.
- Jörn Diedrichsen and Nikolaus Kriegeskorte. 2017. Representational models: A common framework for understanding encoding, pattern-component, and representational-similarity analysis. *PLoS computational biology*, 13(4):e1005508.
- Ning Ding, Yulin Chen, Bokai Xu, Yujia Qin, Shengding Hu, Zhiyuan Liu, Maosong Sun, and Bowen Zhou. 2023. Enhancing chat language models by scaling high-quality instructional conversations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023*, pages 3029–3051.
- Kshitij Dwivedi and Gemma Roig. 2019. Representation similarity analysis for efficient task taxonomy & transfer learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12387–12396.
- Angela D Friederici. 2011. The brain basis of language processing: from structure to function. *Physiological reviews*.
- Paul Friedl. 2023. Dis/similarities in the design and development of legal and algorithmic normative systems: the case of perspective api. *Law, Innovation and Technology*, 15(1):25–59.
- Deep Ganguli, Amanda Askell, Nicholas Schiefer, Thomas I Liao, Kamilė Lukošiuūtė, Anna Chen, Anna Goldie, Azalia Mirhoseini, Catherine Olsson, Danny Hernandez, et al. 2023. The capacity for moral self-correction in large language models. *arXiv preprint arXiv:2302.07459*.
- SK Gargee, Pranav Bhargav Gopinath, Shridhar Reddy SR Kancharla, CR Anand, and Anoop S Babu. 2022. Analyzing and addressing the difference in toxicity prediction between different comments with same semantic meaning in google’s perspective api. In *ICT Systems and Sustainability: Proceedings of ICT4SD 2022*, pages 455–464. Springer.
- Xiaochuang Han. 2023. In-context alignment: Chat with vanilla language models before fine-tuning. *arXiv preprint arXiv:2308.04275*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. In *9th International Conference on Learning Representations, ICLR 2021*.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. 2022. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.
- Nora Hollenstein, Antonio de la Torre, Nicolas Langer, and Ce Zhang. 2019. Cognival: A framework for cognitive word embedding evaluation. In *Proceedings of the 23rd Conference on Computational Natural Language Learning, CoNLL 2019*, pages 538–549.
- Jen-tse Huang, Man Ho Lam, Eric John Li, Shujie Ren, Wenxuan Wang, Wenxiang Jiao, Zhaopeng Tu, and Michael R Lyu. 2023. Emotionally numb or empathetic? evaluating how LLMs feel using emotion-bench. *arXiv preprint arXiv:2308.03656*.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. What does bert learn about the structure of language? In *ACL 2019-57th Annual Meeting of the Association for Computational Linguistics*.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Marcie L. King, Iris I. A. Groen, Adam Steel, Dwight J. Kravitz, and Chris I. Baker. 2019. Similarity judgments and cortical visual responses reflect different properties of object and scene categories in naturalistic images. *NeuroImage*, 197:368–382.

- Nikolaus Kriegeskorte, Marieke Mur, and Peter A Bandettini. 2008. Representational similarity analysis—connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, page 4.
- Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. AlpacaEval: An automatic evaluator of instruction-following models.
- Zhengzhong Liu, Aurick Qiao, Willie Neiswanger, Hongyi Wang, Bowen Tan, Tianhua Tao, Junbo Li, Yuqi Wang, Suqi Sun, Omkar Pangarkar, Richard Fan, Yi Gu, Victor Miller, Yonghao Zhuang, Guowei He, Haonan Li, Fajri Koto, Liping Tang, Nikhil Ranjan, Zhiqiang Shen, Xuguang Ren, Roberto Iriando, Cun Mu, Zhiting Hu, Mark Schulze, Preslav Nakov, Tim Baldwin, and Eric P. Xing. 2023. LLM360: Towards fully transparent open-source llms. *arXiv preprint arXiv:2312.06550*.
- Yifei Luo, Minghui Xu, and Deyi Xiong. 2022. Cog-taskonomy: Cognitively inspired task taxonomy is beneficial to transfer learning in NLP. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 904–920.
- Meta. 2024. Introducing meta LLaMA 3: The most capable openly available llm to date. web link.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, 26.
- Tom M Mitchell, Svetlana V Shinkareva, Andrew Carlson, Kai-Min Chang, Vicente L Malave, Robert A Mason, and Marcel Adam Just. 2008. Predicting human brain activity associated with the meanings of nouns. *science*, 320(5880):1191–1195.
- OpenAI. 2023. Introducing chatgpt. web link.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. Instruction tuning with gpt-4. *arXiv preprint arXiv:2304.03277*.
- Francisco Pereira, Bin Lou, Brianna Pritchett, Samuel Ritter, Samuel J Gershman, Nancy Kanwisher, Matthew Botvinick, and Evelina Fedorenko. 2018. Toward a universal decoder of linguistic meaning from brain activation. *Nature communications*, 9(1):963.
- Yuqi Ren and Deyi Xiong. 2021. Cogalign: Learning to align textual neural representations to cognitive language processing signals. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3758–3769.
- Yuqi Ren and Deyi Xiong. 2022. Bridging between cognitive processing signals and linguistic features via a unified attentional network. In *Proceedings of the AACL Conference on Artificial Intelligence*, volume 36, pages 49–58.
- Martin Schrimpf, Idan Asher Blank, Greta Tuckute, Carina Kauf, Eghbal A Hosseini, Nancy Kanwisher, Joshua B Tenenbaum, and Evelina Fedorenko. 2021. The neural architecture of language: Integrative modeling converges on predictive processing. *Proceedings of the National Academy of Sciences*, 118(45):e2105646118.
- Jingyuan Sun, Shaonan Wang, Jiajun Zhang, and Chengqing Zong. 2019. Towards sentence-level brain decoding with distributed representations. In *Proceedings of the AACL Conference on Artificial Intelligence*, volume 33, pages 7047–7054.
- Michel Thiebaut de Schotten and Stephanie J Forkel. 2022. The emergent properties of the connected brain. *Science*, 378(6619):505–510.
- Mariya Toneva and Leila Wehbe. 2019. Interpreting and improving natural-language processing (in machines) with natural language-processing (in the brain). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019*, pages 14928–14938.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023. LLaMA 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Lewis Tunstall, Edward Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro von Werra, Cl  mentine Fourrier, Nathan Habib, Nathan Sarrazin, Omar Sanseviero, Alexander M. Rush, and Thomas Wolf. 2023. *Zephyr: Direct distillation of LM alignment*. *CoRR*, abs/2310.16944.
- Xuena Wang, Xueting Li, Zi Yin, Yue Wu, and Jia Liu. 2023. Emotional intelligence of large language models. *Journal of Pacific Rim Psychology*, 17:18344909231213958.
- Leila Wehbe, Ashish Vaswani, Kevin Knight, and Tom Mitchell. 2014. Aligning context-based statistical models of language with brain activity during reading. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 233–243.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama,

- Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. 2023. Wizardlm: Empowering large language models to follow complex instructions. *arXiv preprint arXiv:2304.12244*.
- Haoyan Xu, Brian Murphy, and Alona Fyshe. 2016. Brainbench: A brain-image test suite for distributional semantic models. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2017–2021. The Association for Computational Linguistics.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023a. Judging LLM-as-a-judge with MT-Bench and Chatbot Arena. *arXiv preprint arXiv:2306.05685*.
- Rui Zheng, Shihan Dou, Songyang Gao, Yuan Hua, Wei Shen, Binghai Wang, Yan Liu, Senjie Jin, Qin Liu, Yuhao Zhou, Limao Xiong, Lu Chen, Zhiheng Xi, Nuo Xu, Wenbin Lai, Minghao Zhu, Cheng Chang, Zhangyue Yin, Rongxiang Weng, Wensen Cheng, Haoran Huang, Tianxiang Sun, Hang Yan, Tao Gui, Qi Zhang, Xipeng Qiu, and Xuanjing Huang. 2023b. Secrets of RLHF in large language models part I: PPO. *CoRR*, abs/2307.04964.

## **A Appendix**

### **A.1 The Sensitivity of Training Stages to Prompt Strategy**

To investigate the sensitivity of different training stages to prompt strategies, we compared the disparity of LLM-brain similarity between the explicit-prompt-appending and no-prompt-appending under the pre-trained LLMs and the alignment LLMs. The experimental results are shown in Table 4.

Scaling	LLM	Training Stage	$\Delta$
7B	Amber	Pre-training	0.0121
	Amberchat	SFT	0.0077
	Ambersafe	SFT+DPO	<b>0.0174</b>
	LLaMA-7B	Pre-training	0.0088
	Vicuna-7B-v1.3	SFT	<b>0.0153</b>
	LLaMA2-7B	Pre-training	0.0061
	Vicuna-7B-v1.5	SFT	0.0051
	LLaMA2-7B-chat	SFT+RLHF	<b>0.0086</b>
	Mistral-7B	Pre-training	0.0089
	Mistral-7B-sft-alpha	SFT	<b>0.0107</b>
	Mistral-7B-sft-beta	SFT	0.0096
	Zephyr-7B	dSFT+dDPO	0.0106
8B	LLaMA3-8B	Pre-training	0.0014
	LLaMA3-8B-chat	SFT+PPO+DPO	<b>0.0168</b>
13B	LLaMA-13B	Pre-training	<b>0.0072</b>
	Vicuna-13B-v1.3	SFT	0.0071
	LLaMA2-13B	Pre-training	0.0086
	Vicuna-13B-v1.5	SFT	0.0075
70B	LLaMA2-13B-chat	SFT+RLHF	<b>0.0091</b>
	LLaMA2-70B	Pre-training	0.0108
	LLaMA2-70B-chat	SFT+RLHF	<b>0.0204</b>
	LLaMA3-70B	Pre-training	0.0241
	LLaMA3-70B-chat	SFT+PPO+DPO	<b>0.0351</b>

Table 4: The disparity in the LLM-brain similarity between the explicit-prompt-appending and no-prompt-appending strategy on different training stages. ‘ $\Delta$ ’ denotes the increment of the explicit-prompt-appending vs. the no-prompt-appending strategy.