Investigation of Information Processing Mechanisms in the Human Brain during Reading Tanka Poetry

Anna Sato[1] [0009−0003−0297−938X], Junichi Chikazoe[2] [0000−0003−4155−1370], Shotaro Funai[2] [0000−0002−2451−1935], Daichi Mochihashi[3] [0000−0003−0344−5382], Yutaka Shikano[4,5] [0000−0003−2107−7536], Masayuki Asahara[6] [0000−0002−5178−7275], Satoshi Iso[7] [0000−0002−2492−5396], and Ichiro Kobayashi[1] [0000−0001−7789−475X]

[1] Ochanomizu University, Tokyo, Japan
{g1920519,koba}@is.ocha.ac.jp
[2] Araya Inc., Tokyo, Japan
{chikazoe_junichi, funai_shotaro}@araya.org
[3] Institute of Statistical Mathematics, Tokyo, Japan
daichi@ism.ac.jp
[4] Gunma University, Gunma, Japan
yshikano@gunma-u.ac.jp
[5] Institute for Quantum Studies, Chapman University, Tokyo, Japan
shikano@chapman.edu
[6] National Institute for Japanese Language and Linguistics, Tokyo, Japan
masayu-a@ninjal.ac.jp
[7] High Energy Accelerator Research Organization, Ibaraki, Japan
iso@post.kek.jp

Abstract. Recent advances in non-invasive brain function measurement technologies, such as functional magnetic resonance imaging (fMRI) and magnetoencephalography (MEG), and the development of machine learning techniques, including deep learning, have led to increased research on the elucidation and quantitative understanding of information processing processes in the human brain. Since the emergence of word2vec, which represents the meaning of natural language words as vectors, features of language stimuli given to the human brain have been represented using large language models in natural language processing and used to estimate brain states. In this study, we used GPT-2, which is known to perform well as a feature for predicting brain states, and investigated the information processing processes in the human brain when reading Japanese short poems, i.e., tanka poetry. In particular, we investigated the hubness of the regions of interest in the brain by applying the PageRank algorithm. As a result, we have found that the cingulate cortex and the insula, which are said to be related to emotion, have hubness in brain regions, while occipital lobe, which are not said to be related to emotion, have also hubness.

Keywords: Language Processing · Deep learning models · Neuroscience · fMRI · Emotion.

1 Introduction

In recent years, there has been a growing interest in exploring the relationship between deep learning models (DLMs) and the human brain. DLMs have been developed and inspired by the structure and behavior of the human brain, and have been trained on massive amounts of data, which has allowed them to achieve impressive result in various fields such as computer vision [1–3], natural language processing [4–8], speech recognition [9, 10], and many others.

They also helps in understanding human cognitive processes and brain functions. The objective of this study is to elucidate the processing prosess of emotion in the human brain. In particular, we will investigate the information processing induced by the stimulus of reading tanka poetry.

2 Related Work

Recent studies suggest that the hidden representations of various DLMs have shown to linearly predict human brain activity [11–13]. Since the introduction of word2vec [7], large language models have been utilized in neuroimaging research to investigate how the human brain processes language [14–20]. By presenting human participants with language stimuli, such as sentences or individual words, and measuring their brain activity, researchers can build models that predict the neural activity patterns in response to these stimuli and decode words and sentences from brain states [20].

Those DLMs have a hierarchical structure, and researchers have been investigating whether the hierarchical structure of DLMs corresponds to the hierarchical organization of the human brain. Convolutional neural networks (CNNs) have been suggested to correspond to the concept of hierarchical processing in human visual processing [21]. Kawasaki et al [22] used a deep learning model to predict brain states under visual and verbal stimuli, and investigated the processing transitions through analysis of these states using representation similarity analysis (RSA) [23]. However, there has been limited research on the relationship between the hierarchy of DLMs and the hierarchy in the human brain under language stimuli.

In this study, we chose GPT-2 as a model for estimating brain activity because it has been shown that the prediction of brain states by using GPT-2 as language features significantly correlates with semantic comprehension [16, 18, 24]. We will construct an encoding model that predicts brain activity from the feature of each layer of GPT-2 and examine hierarchical processing in the brain when reading a text. By using a deep learning model that processes natural language as a working model, we investigate emotional activities in the human brain that are similar to the sensitivity and sensation induced by tanka, a verbal art form, instead of dealing with emotions expressed as direct responses in the human brain under various stimuli.

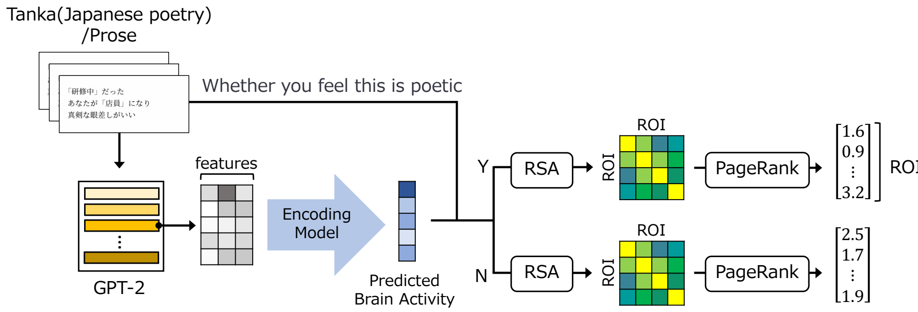

Fig. 1: Overall process: the features of tanka are represented by GPT-2. representation similarity analysis (RSA) is applied to the brain activity data predicted by the encoding model based on the features. The representation dissimilarity matrix (RDM) obtained by RSA is regarded as an adjacency matrix, and the PageRank algorithm is applied to determine the hubness of the regions of interest (ROI) of brain activity stimulated by poetic and unpoetic sentences.

3 Analysis of Brain Activity under Language Stimuli

Figure 1 shows an overview of the method in this study.

To compare brain activities evoked by tanka poetry, and unpoetic sentences, we 1) extract features using tanka poetry that presented in the experiment as input; 2) construct an encoding model to predict brain activity from the features; 3) separate the predicted brain activity under stimuli labeled as poetic or unpoetic; 4) apply RSA to the brain activity and create a RDM that shows the connectivity of ROIs in the brain; and 5) apply PageRank algorithm [25] to the matrix and find the hubness in ROIs.

3.1 Brain Activity Data and Experimental Tasks

Brain activity data were collected from 15 male subjects and 17 female subjects. All subjects are right-handed and native Japanese speakers, aged from 18 to 34 years. Functional scans were collected using a 3.0 T scanner at the National Institute for Physiological Sciences (NIPS) with repetition time (TR) = 750 ms, and voxel size = 2.0mm × 2.0mm × 2.0mm.

As experimental tasks, subjects read and evaluate 300 sentences from the Balanced Corpus of Contemporary Written Japanese (BCCWJ) [26]. The sentence dataset consists of 150 tanka, 31-syllable poems originated in Japan, and 150 prose that have approximately 31 characters. Each tanka/prose was divided into three lines, where only the first line was displayed for the first three seconds, up to the second line for the next three seconds, and all three lines for the following three seconds. In the next three seconds, a slide with the question “Do you feel that this is poetic?” written on it is displayed, and the subjects answers “Yes” or “No” by pressing a button. One trial is conducted in 12 seconds. These data were collected during six scanning sessions with breaks in between, and each session had 50 trials (25 tanka and 25 prose). The sentences were presented in a different random order for each subject.

The study was approved by the Ethical Committee of the National Institute for Japanese Language and Linguistics, the National Institute for Physiological Sciences of Japan and the Institute of Statistical Mathematics.

3.2 Encoding Model

The method employed in this study to construct an encoding model is the one by Naselaris et al [27]. As a method for constructing the encoding model, linear regression is performed between the features extracted from the data that stimulate the human brain and the brain activity state under stimulation, and weights are learned so that the measured brain activity pattern and the predicted brain activity pattern are close to each other. In general, ridge regression is applied to linear regression, and by observing the regression coefficients, it is possible to observe the behavior of the voxels.

3.3 Representational Similarity Analysis

Representational Similarity Analysis (RSA) uses correlation matrices to measure the degree of similarity or dissimilarity between representational patterns, proposed by Kriegeskorte et al [23, 28]. RSA measures the distance between the representational geometries of different brain regions or conditions and calculates the dissimilarity between different neural activation patterns to create Representational Dissimilarity Matrix (RDM) which refers to the matrix of pairwise distance between representational patterns.

In this study, correlation distance (1 − Pearson's correlation coefficient) is used as a measure of similarity. We measure the functional dissimilarity at each time point for each ROI, and then calculate the dissimilarity between the created time × time RDMs to create a ROI × ROI RDM, which represents the connectivity among ROIs (see, Figure 2). The brain region division atlas used in the experiments was Destrieux Atlas [29] provided by FreeSurfer[8].

3.4 PageRank Algorithm

We used PageRank algorithm to identify regions in the brain that serve as hubs in the linkage among regions of interest (ROIs). The PageRank algorithm [25] is a widely used algorithm for ranking web pages, and outputs a PageRank score which indicates a ranking of importance. The application of the PageRank algorithm is not limited to Web networks but can also be extended to social and brain networks [30–32]. The PageRank score of $r_{k+1}(P_i)$ is obtained by a power method as follows.

$$r_{k+1}(P_i) = \sum_{P_j \in B_{P_i}} \frac{r_k(P_j)}{|P_j|}$$

[8] https://surfer.nmr.mgh.harvard.edu/

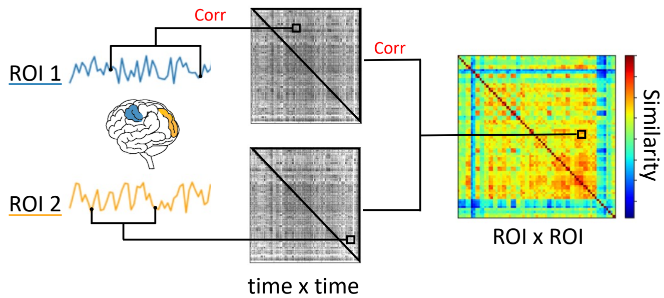

Fig. 2: Creation of a ROI × ROI RDM matrix by RSA

The PageRank score for page P_i is dependent on the PageRank scores for each page P_j contained in B_{P_i}, the set of all pages linking to page P_j, dividing by $|P_i|$, the number of outbound links page P_j.

4 Experiments

4.1 Encoding models for the experiments

Encoding models are constructed to estimate the brain state from the features of linguistic stimuli given to the human brain. Specifically, we performed ridge regression from the features extracted from each layer of GPT-2 to brain activity. When a tanka/prose slide was displayed, the language stimulus of t_i is the presented sentence, and when a question slide was shown, the language stimulus is the sentence which is added the question sentence to the preceding tanka/prose.

The features of language stimuli extracted from each layer of GPT-2 are 1024 dimensions, but we reduced them to 300 dimensions using principal component analysis (PCA) to mitigate the risk of overfitting the model. In order to account for the hemodynamic delay when observing blood oxygenation level dependent (BOLD) signal by fMRI, we concatenated seven features from t_{i-9} to t_{i-3} to predict brain activity at time t_i. Of the six sessions, one session was test data and the remaining five sessions were training data for building an encoding model, and then predicted brain activities for all sessions were collected. This regularization coefficient is chosen amongst ten values log-spaced between 10^1 and 10^6 by 5-fold Cross-Validation. For each layer and for each subject, the regularization coefficient is found as the value that led to the best performance with a Pearson correlation between the predicted and measured brain activities.

4.2 Experimental settings

The features of tanka and prose are represented by GPT-2 in this study. We used a pretrained Japanese GPT-2 model with 24 layers[9], provided by Hugging

[9] https://huggingface.co/rinna/japanese-gpt2-medium

Face. We fine-tuned GPT-2 on 3571 tanka in BCCWJ that were not used in the experiment. We further fine-tuned GPT-2 for binary text classification for each subject, whether a subject felt it was poetic or not in the experiment. In the second fine-tuning, the training data consisting of 250 sentences, which were displayed during five scanning sessions used for encoding model training. To investigate how information can be changed across each layer of GPT-2, we extracted text features from each layer.

4.3 Results

We analyzed brain activity data from 19 subjects of 32 subjects who participated in the experiment, selected with a statistical threshold of $p < 0.01$. The p-values were calculated as the fraction of 200 random sequences in which the correlation with the measured data was greater than or equal to the correlation between the predicted and measured data. Furthermore, the voxels that were corrected with FDR at $q < 0.05$ and had positive correlation with the measured brain activity were extracted and used for analysis.

In Figure 3, we show the accuracy predicted by an encoding model. Pearson correlation coefficient was used for evaluating the accuracy. Figure 3a visualizes the prediction accuracy on the cortical surface for one subject. The correlation coefficient between the predicted and mesured brain activity patterns was computed for each voxel and for each scanning session. The average correlation coefficient across all sessions is then calculated. The map shows only voxels that were multiple-corrected with a false discovery rate (FDR) correction at $q < 0.05$. Figure 3b shows that the prediction accuracy of each layer of GPT-2. The correlation coefficients for each layer were averaged across subjects, and across

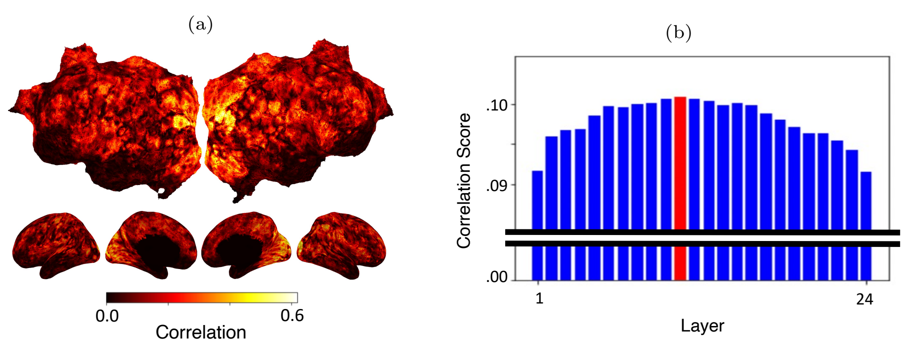

Fig. 3: Performance of the encoding model: (a)Brain map of prediction accuracy for each voxel at the 24th layer. Each correlation coefficient was averaged across the six sessions. (b)Prediction accuracy for each layer. Each correlation coefficient was averaged across the voxels that were FDR-corrected and have positive correlation between the predicted and measured brain activity.

those voxels that are FDR-corrected and have positive correlation between the predicted and measured brain activity. The layer with the highest correlation coefficients, which was the 10th layer, is highlighted in red. Prior studies have also reported better prediction accuracy in the middle layer of DLMs [33–35].

We created RDMs for each layer representing the behavior of the ROI in each of the predicted brain activities labeled "poetic" and "unpoetic" based on our experiments. The dissimilarity between 48 RDMs of 2(poetic/unpopetic) $\times$ 24(the number of hierarchies), each averaged across subjects, was calculated (see, Figure 4a). Figure 4b shows the results of dimensional compression of the above RDM into three dimensions using UMAP [36].

Next, we applied the PageRank algorithm to a similarity matrix of a ROI $\times$ ROI RDM to obtain PageRank scores for each ROI and detect hub regions for each poetic and unpoetic state. PageRank scores were calculated for each state and averaged across subjects. The result of subtracting the PageRank score of the unpoetic state from the PageRank score of the poetic state are visualized in Figure 5a; the 1st layer as the lowest layer, the 13th layer as middle layer and the 24th layer as the highest layer. The regions colored red have higher PageRank score when the brain feels poetic, indicating that the regions have higher "hubness", while the regions colored blue have higher hubness when the brain did not feel poetic. Figure 5b shows the subtracted PageRank scores after averaged in each regions. We selected four regions where there were significant differences: the cingulate cortex (green), the occipital lobe (orange), the parietal lobe (red) and the frontal lobe (blue). Layers in which significant differences were found were marked with a star ($p < 0.01$, a two-sample t-test).

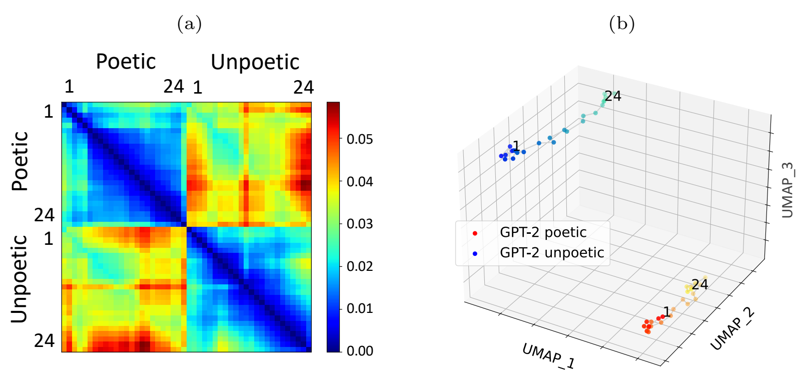

Fig. 4: (a)RDM of hidden layer for each state. (b)The results of dimensional compression of a layer $\times$ layer RDM

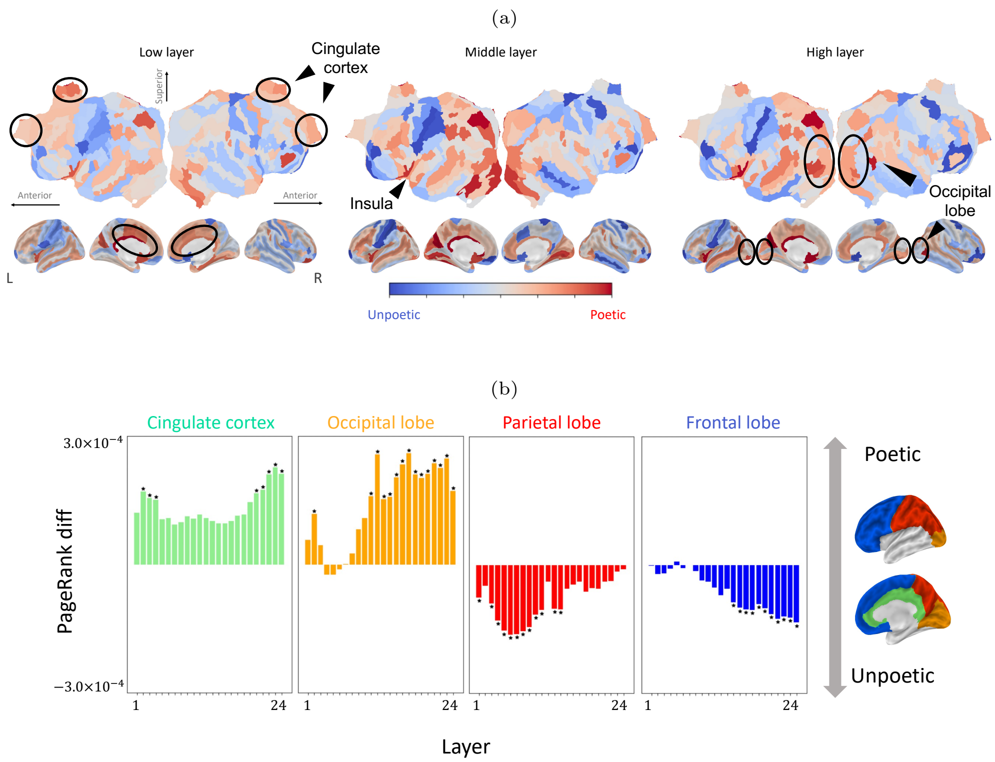

Fig. 5: PageRank scores of poetic state vs. unpoetic state: (a)Brain map of the subtracted PageRank scores at 1st, 13th and 24th layer as low, middle and high layer respectively. (b)The subtracted PageRank scores after averaged in each region(*$p <$ 0.01).

4.4 Discussion

The information is gradually transitioning in Figure 4b, indicating that the brain activity that represented by each layer of GPT-2 changes as the hierarchy shifts.

The cingulate cortex that had a high PageRank score in Figure 5a is involved in a wide range of cognitive and emotional processes. It is said to be the one which plays a key role in attention, decision making, memory, and emotion regulation [37, 38]. In a previous study of brain activity during reading of poetry and prose [39], a part of the cingulate gyrus was cited as a brain region that was activated as the emotionality of the text increases. The insula, which also had the same high PageRank score is also said to be an area involved in emotions and experiences [40, 41]. In the region around the cingulate cortex and the left insula, the subtracted values were consistently positive through almost all layers.

The occipital lobe's PageRank score was also elevated when participants felt poetic. However, this region is primarily responsible for processing visual information and is typically activated while reading a text, and those well-known functions do not explain why the PageRank score was high when feeling poetic.

In the parietal and frontal lobes, PageRank scores were high in all layers when not feeling poetic (see, Figure 5b). Further investigation is needed to fully understand these results.

5 Conclusion

In this study, using brain activity data measured by fMRI while reading Japanese tanka poetry and prose, we constructed an encoding model to predict brain state from the features extracted from the hidden layers of a DLM. We especially used GPT-2 to represent the feature of those tanka poetry and prose.

We investigated the ROIs that work as hubs of information connectivity in the brain when feeling poetic by applying the PageRank algorithm to a matrix of connectivity between ROIs. By this, we found a significant importance of the cingulate cortex and the left insula, which has been previously reported to be activated when a person feeling more poetic. On the other hand, we also found results that could not be explained by previous studies, such as the finding of hubness in the occipital lobe at higher layers of GPT-2 and hubness in the parietal and frontal lobes when not feeling poetic. As future work, we will further investigate the results obtained in this study.

Acknowledgement

This study was supported by Inter-University Research Institute Corporation (I-URIC) and the Cooperative Study Program (20-640, 21-543 and 22NIPS630) of National Institute for Physiological Sciences (NIPS). It was also supported by JSPS KAKENHI (21H05060 and 21H05061), JST MOONSHOT (JPMJMS2012) and JST PRESTO (JPMJPR20M4).

References

1. Kolesnikov, A., Beyer, L., Zhai, X., Puigcerver, J., Yung, J., Gelly, S., Houlsby, N.: Big transfer (bit): General visual representation learning (2020)
2. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2021), https://openreview.net/forum?id=YicbFdNTTy
3. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N., Weinberger, K. (eds.) Advances in Neural Information Processing Systems. vol. 27. Curran Associates, Inc. (2014)

4. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I.: Language models are unsupervised multitask learners. OpenAI Blog (2018), https://d4mucfpksywv.cloudfront.net/better-language-models/language-models.pdf
5. Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D.M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., Amodei, D.: Language models are few-shot learners (2020)
6. Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. In: Empirical Methods in Natural Language Processing (EMNLP). pp. 1532–1543 (2014), http://www.aclweb.org/anthology/D14-1162
7. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. Advances in neural information processing systems **26** (2013), https://proceedings.neurips.cc/paper/2013/file/9aa42b31882ec039965f3c4923ce901b-Paper.pdf
8. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., Liu, P.J.: Exploring the limits of transfer learning with a unified text-to-text transformer (2020)
9. Radford, A., Kim, J.W., Xu, T., Brockman, G., McLeavey, C., Sutskever, I.: Robust speech recognition via large-scale weak supervision (2022)
10. Amodei, D., Anubhai, R., Battenberg, E., Case, C., Casper, J., Catanzaro, B., Chen, J., Chrzanowski, M., Coates, A., Diamos, G., Elsen, E., Engel, J., Fan, L., Fougner, C., Han, T., Hannun, A., Jun, B., LeGresley, P., Lin, L., Narang, S., Ng, A., Ozair, S., Prenger, R., Raiman, J., Satheesh, S., Seetapun, D., Sengupta, S., Wang, Y., Wang, Z., Wang, C., Xiao, B., Yogatama, D., Zhan, J., Zhu, Z.: Deep speech 2: End-to-end speech recognition in english and mandarin (2015)
11. Holdgraf, C., Rieger, J., Micheli, C., Martin, S., Knight, R., Theunissen, F.: Encoding and decoding models in cognitive electrophysiology. Frontiers in Systems Neuroscience **11**, 61 (09 2017). https://doi.org/10.3389/fnsys.2017.00061
12. Caucheteux, C., King, J.R.: Language processing in brains and deep neural networks: computational convergence and its limits. bioRxiv (2021). https://doi.org/10.1101/2020.07.03.186288, https://www.biorxiv.org/content/early/2021/01/14/2020.07.03.186288
13. Yu, T., Shinji, N.: High-resolution image reconstruction with latent diffusion models from human brain activity. bioRxiv (2022)
14. Toneva, M., Stretcu, O., Poczos, B., Wehbe, L., Mitchell, T.M.: Modeling task effects on meaning representation in the brain via zero-shot meg prediction. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) Advances in Neural Information Processing Systems. vol. 33, pp. 5284–5295. Curran Associates, Inc. (2020)
15. Goldstein, A., Zada, Z., Buchnik, E., Schain, M., Price, A., Aubrey, B., Nastase, S.A., Feder, A., Emanuel, D., Cohen, A., Jansen, A., Gazula, H., Choe, G., Rao, A., Kim, S.C., Casto, C., Fanda, L., Doyle, W., Friedman, D., Dugan, P., Melloni, L., Reichart, R., Devore, S., Flinker, A., Hasenfratz, L., Levy, O., Hassidim, A., Brenner, M., Matias, Y., Norman, K.A., Devinsky, O., Hasson, U.: Thinking ahead: spontaneous prediction in context as a keystone of language in humans and machines. bioRxiv (2021). https://doi.org/10.1101/2020.12.02.403477, https://www.biorxiv.org/content/early/2021/09/30/2020.12.02.403477

16. Schrimpf, M., Blank, I., Tuckute, G., Kauf, C., Hosseini, E.A., Kanwisher, N., Tenenbaum, J., Fedorenko, E.: Artificial neural networks accurately predict language processing in the brain. bioRxiv (2020). https://doi.org/10.1101/2020.06.26.174482, https://www.biorxiv.org/content/early/2020/06/27/2020.06.26.174482
17. Schrimpf, M., Blank, I.A., Tuckute, G., Kauf, C., Hosseini, E.A., Kanwisher, N.G., Tenenbaum, J.B., Fedorenko, E.: The neural architecture of language: Integrative modeling converges on predictive processing. Proceedings of the National Academy of Sciences (PNAS) (2021). https://doi.org/10.1073/pnas.2105646118, https://www.pnas.org/content/118/45/e2105646118
18. Caucheteux, C., Gramfort, A., King, J.R.: Gpt-2's activations predict the degree of semantic comprehension in the human brain. bioRxiv (2021). https://doi.org/10.1101/2021.04.20.440622, https://www.biorxiv.org/content/early/2021/04/21/2021.04.20.440622
19. Schwartz, D., Toneva, M., Wehbe, L.: Inducing brain-relevant bias in natural language processing models. Advances in neural information processing systems **32** (10 2019)
20. Tang, J., LeBel, A., Jain, S., Huth, A.G.: Semantic reconstruction of continuous language from non-invasive brain recordings. bioRxiv (2022). https://doi.org/10.1101/2022.09.29.509744, https://www.biorxiv.org/content/early/2022/09/29/2022.09.29.509744
21. Eickenberg, M., Gramfort, A., Varoquaux, G., Thirion, B.: Seeing it all: Convolutional network layers map the function of the human visual system. Neuroimage **152**, 184–194 (May 2017). https://doi.org/10.1016/j.neuroimage.2016.10.001
22. Kawasaki, H., Nishida, S., Kobayashi, I.: Hierarchical processing of visual and language information in the brain. In: Findings of the Association for Computational Linguistics: AACL-IJCNLP 2022. pp. 405–410. Association for Computational Linguistics, Online only (Nov 2022), https://aclanthology.org/2022.findings-aacl.38
23. Kriegeskorte, N., Mur, M., Bandettini, P.: Representational similarity analysis - connecting the branches of systems neuroscience. Frontiers in Systems Neuroscience **2**, 4 (2008). https://doi.org/10.3389/neuro.06.004.2008, https://www.frontiersin.org/article/10.3389/neuro.06.004.2008
24. Caucheteux, C., Gramfort, A., King, J.R.: Evidence of a predictive coding hierarchy in the human brain listening to speech. Nature Human Behaviour **7**, 430–441 (2023). https://doi.org/10.1038/s41562-022-01516-2
25. Page, L., Brin, S., Motwani, R., Winograd, T.: The PageRank Citation Ranking: Bringing Order to the Web. Tech. rep., Stanford Digital Library Technologies Project (1998), http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.31.1768
26. The balanced corpus of contemporary written japanese, https://clrd.ninjal.ac.jp/bccwj/en/
27. Naselaris, T., Kay, K., Nishimoto, S., Gallant, J.: Encoding and decoding in fmri. NeuroImage **56**(2), 400–410 (May 2011). https://doi.org/10.1016/j.neuroimage.2010.07.073
28. Kriegeskorte, N., Kievit, R.A.: Representational geometry: integrating cognition, computation, and the brain. Trends Cogn Sci **17**(8), 401–412 (Aug 2013). https://doi.org/10.1016/j.tics.2013.06.007
29. Destrieux, C., Fischl, B., Dale, A.M., Halgren, E.: Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature. NeuroImage **53**(1), 1–15 (2010)

30. Zhu, Q., Yang, J., Xu, B., Hou, Z., Sun, L., Zhang, D.: Multimodal brain network jointly construction and fusion for diagnosis of epilepsy. Frontiers in Neuroscience (2021). https://doi.org/10.3389/fnins.2021.734711
31. Biaobin, J., Kyle, K., David, G., Gribskov, M.: Aptrank: an adaptive pagerank model for protein function prediction on bi-relational graphs. Bioinformatics (2017). https://doi.org/10.1093/bioinformatics/btx029
32. Bin, L., Shuangyan, J., Quan, Z.: Hits-pr-hhblits: protein remote homology detection by combining pagerank and hyperlink-induced topic search. Brief Bioinform (2020). https://doi.org/10.1093/bib/bby104
33. Caucheteux, C., King, J.R.: Brains and algorithms partially converge in natural language processing. Communications Biology **5** (2022). https://doi.org/10.1038/s42003-022-03036-1
34. Caucheteux, C., Gramfort, A., King, J.R.: Deep language algorithms predict semantic comprehension from brain activity. Scientific Reports **12**(1) (2022). https://doi.org/10.1038/s41598-022-20460-9
35. Huth, A.G., de Heer, W.A., Griffiths, T.L., Theunissen, Frédéric E and G̃allant, J.L.: Natural speech reveals the semantic maps that tile human cerebral cortex. Nature **532**(7600), 453–458 (Apr 2016). https://doi.org/10.1038/nature17637
36. McInnes, L., Healy, J., Melville, J.: Umap: Uniform manifold approximation and projection for dimension reduction (2018), http://arxiv.org/abs/1802.03426
37. Bush, G., Luu, P., Posner, M.I.: Cognitive and emotional influences in anterior cingulate cortex. Trends in Cognitive Sciences **4**, 215–222 (2000)
38. Shackman, A., Salomons, T., Slagter, H., Fox, A., Winter, J., Davidson, R.: The integration of negative affect, pain and cognitive control in the cingulate cortex. NATURE REVIEWS NEUROSCIENCE **12**(3), 154–167 (2011)
39. Zeman, A., Milton, F., Smith, A., Rylance, R.: By heart an fmri study of brain activation by poetry and prose. Journal of Consciousness Studies **20**(9-10), 9–10 (2013)
40. Kurth, F., Zilles, K., Fox, P., Laird, A., Eickhoff, S.: Kurth f, zilles k, fox pt, laird ar, eickhoff sb. a link between the systems: functional differentiation and integration within the human insula revealed by meta-analysis. brain struct funct 214: 519-534. Brain structure & function **214**, 519–34 (06 2010). https://doi.org/10.1007/s00429-010-0255-z
41. Uddin, L., Nomi, J., Hébert-Seropian, B., Ghaziri, J., Boucher, O.: Structure and function of the human insula. Journal of Clinical Neurophysiology **34**, 300–306 (07 2017). https://doi.org/10.1097/WNP.0000000000000377