

Learning Exemplar Representations in Single-Trial EEG Category Decoding

Jack A. Kilgallen
Hamilton Institute
Maynooth University
Maynooth, Ireland
jkilgallen@protonmail.com

Barak A. Pearlmutter
Department of Computer Science
Maynooth University
Maynooth, Ireland
barak@pearlmutter.net

Jeffrey Mark Siskind
Elmore Family School of Electrical
and Computer Engineering
Purdue University
West Lafayette, IN, USA
qobi@qobi.org

Abstract—Within neuroimaging studies it is a common practice to perform repetitions of trials in an experiment when working with a noisy class of data acquisition system, such as electroencephalography (EEG) or magnetoencephalography (MEG). While this approach can be useful in some experimental designs, it presents significant limitations for certain types of analyses, such as identifying the category of an object observed by a subject. In this study we demonstrate that when trials relating to a single object are allowed to appear in both the training and testing sets, almost any classification algorithm is capable of learning the representation of an object given only category labels. This ability to learn object representations is of particular significance as it suggests that the results of several published studies which predict the category of observed objects from EEG signals may be affected by a subtle form of leakage which has inflated their reported accuracies. We demonstrate the ability of both simple classification algorithms, and sophisticated deep learning models, to learn object representations given only category labels. We do this using two datasets; the Kaneshiro et al. (2015) dataset and the Gifford et al. (2022) dataset. Our results raise doubts about the true generalizability of several published models and suggests that the reported performance of these models may be significantly inflated.

Index Terms—EEG, machine learning, leakage, category decoding

I. INTRODUCTION

Within neuroimaging studies when working with a noisy modality such as electroencephalography (EEG), or magnetoencephalography (MEG), it is common practice to perform multiple trials using a single stimulus. With some experimental designs this repetition of trials can be useful for reducing the impact of noise on the analysis being performed, particularly when extracting features such as Event Related Potentials (ERPs) [1]. However, this practice presents significant limitations for certain types of analyses such as *single-trial category decoding*.

This publication emanated from research supported in part by Science Foundation Ireland research grants 16/RI/3399, 18/CRT/6049, and 20/FFP-P/8853, the United States National Science Foundation Grant 1734938-IIS, and the Defense Advance Research Projects Agency. This publication is approved for public release; distribution is unlimited. Any opinions, findings, views, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views, official policies, or endorsements, either expressed or implied, of the sponsors. Finally, we thank Hari M. Bharadwaj for suggesting that we look into the issue discussed in this paper as a potential source of leakage. Copyright notice: 979-8-3503-5298-6/24/\$31.00 ©2024 IEEE

In single-trial category decoding experiments, subjects are shown a series of objects from different categories while their EEG signals are recorded. The aim of the experiment is to identify the category of an object observed by a subject from a single-trial of their EEG signals. This is a challenging task as the EEG signals are noisy, and the features which are relevant to the category of an object may be subtle.

When performing single-trial category decoding, the use of multiple trials per object can allow a model to use information which we would not expect it to have access to at prediction time (*leakage*). When multiple trials per object are used, a classification algorithm may overfit to features of an EEG signal which are specific to an observed object from a category (*an exemplar representation*) rather than to features of neural responses which are shared by many or all objects from a category (*a category representation*). Moreover, if trials relating to specific stimuli are allowed to appear in both the training and testing sets, then algorithms which overfit to exemplar representations are likely to achieve higher accuracies. This occurs as their performance is evaluated using trials which are more similar to the trials they were trained on than would be expected in a real-world scenario. In this way a loss of generalizability is rewarded by the evaluation procedure.

An analogy to this form of leakage can be found in the case of a machine learning student who wants to classify images of cats and dogs. The student collects photos from friends and family of their cats and their dogs, but each person they ask sends multiple photos of each of their pets. If the student trains a computer vision algorithm to classify the images of cats and dogs, but images relating to specific cats or dogs are allowed to appear in both the training and testing sets, then the algorithm may be affected by a similar form of leakage. Instead of learning what a dog or a cat generally looks like, the algorithm may learn what a specific dog or cat looks like. It would then achieve higher accuracy by overfitting to the features of that specific dog or cat (i.e. learning the representation of an exemplar). The situation would be further exacerbated if the dataset didn't contain many breeds of dogs or cats, as the algorithm would have fewer opportunities to learn what features are shared by all dogs or cats. Similarly, in the case of single-trial category decoding, algorithms may achieve increased by overfitting to features of a signal a

subject’s brain generates when they observe a specific object.

While this form of leakage may seem contrived, it is increasingly present in recent literature. Since the publication of the Kaneshiro et al. (2015) [2] dataset (hereafter referred to as the Kaneshiro dataset) which features multiple trials per stimulus, 12 studies have made use of their dataset for single-trial EEG category decoding [3]–[14]. However, none of these studies make any mention of the potential for leakage in their experimental design, and so it is likely that the results of these studies are affected.

To investigate the ability of classification algorithms to exploit exemplar representations to inflate their performance at single-trial category decoding, we augmented two EEG category decoding datasets to remove any semantic meaning from their category labels. We remove semantic meaning by creating new category labels for which no reasonable distinction can be made between categories. Therefore, when performing category decoding on such a dataset, any model which achieves a decoding accuracy which is above chance, in a statistically significant manner, would suggest that the model has exploited knowledge of an exemplar representation.

Our analysis demonstrates not only that some bespoke EEG decoding algorithms are capable of learning exemplar representations given only category labels, but that even simple classification algorithms such as k-Nearest Neighbors (k-NN) are capable of exploiting exemplar representations. Moreover, we demonstrate that this form of leakage can occur using two separate datasets; the Kaneshiro dataset [2] and the Gifford et al. (2022) [15] dataset (hereafter referred to as the Gifford dataset).

II. RELATED WORK

While machine learning has a broad range of applications, training models on data without a strong understanding of the domain the data originates from can cause confounds and leakage to go unnoticed. This has occurred in a variety of contexts [16] and attention is being given to rooting out studies affected by such confounds [17]. Even within a domain which might be considered niche, such as neuroimaging, significant confounds, such as the Block Design Confound [18], have been discovered in recent literature. The work presented here is part of a sustained effort to identify erroneous or exaggerated claims resulting from a misapplication of machine learning techniques in neuroimaging studies.

III. MATERIALS AND METHODS

A. Datasets

The datasets used in this study are the Kaneshiro dataset [2] and the Gifford dataset [15]. They were selected as they were publically available, and possessed both multiple trials per stimulus, and a category hierarchy.

1) *The Kaneshiro Dataset [2]*: This dataset consists of preprocessed EEG recordings taken from 10 subjects while they viewed 72 images evenly distributed across 6 categories: Human Body (HB), Human Face (HF), Animal Body (AB), Animal Face (AF), Fruit/Vegetable (FV) and Inanimate Object

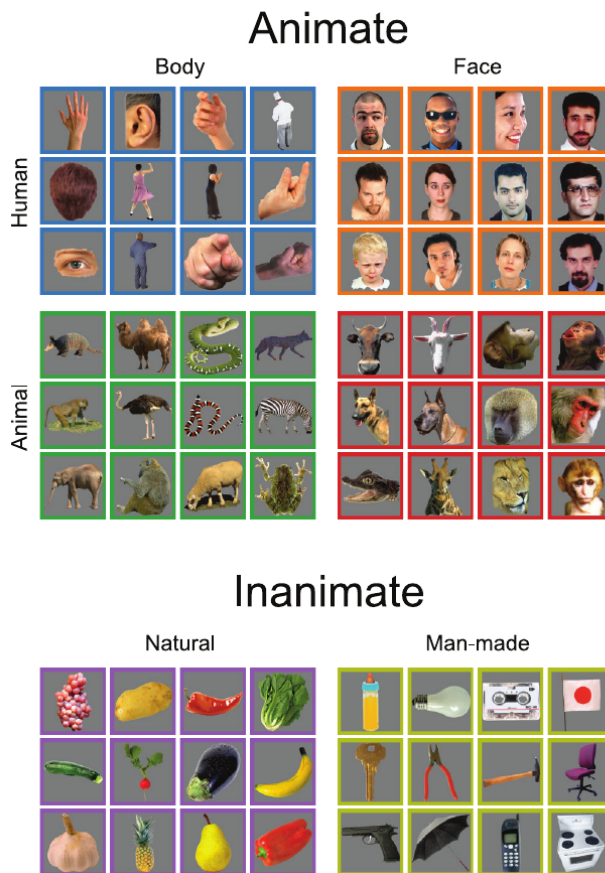


Fig. 1. The stimuli presented to subjects in the Kaneshiro dataset [2].

(IO). Fig. refkaneshiro-category-structure shows the stimuli that were presented to the subjects grouped by category. To reduce the impact that noise would have on their analysis the dataset contains 72 trials for each image per subject which were presented in random order. This gives a total of 5,184 trials per participant. The data was recorded using a 128 channel EEG system with a sampling rate of 1 kHz. The EEG signals were then preprocessed using a high-pass fourth-order Butterworth filter to attenuate frequencies below 1 Hz, and a low-pass Chebyshev Type I filter to attenuate frequencies above 25 Hz. Ocular artifacts were removed using the Bell-Sejnowski Infomax independent component analysis algorithm [19], and finally the data was subsampled to 62.5 Hz to reduce the computational cost of the analysis.

2) *The Gifford Dataset [15]*: This dataset is made up of preprocessed EEG recordings taken from 10 subjects while they viewed a selection of images from the THINGS image database [20]. The images were divided across 27 high-level categories, which span 1854 object concepts. The object categories were partitioned into training and test sets with 1654 object concepts used in the training set and 200 used in the test set. For the training set 10 exemplars per concept which were each presented to the subjects once for a total of 16,540 trials per subject. For the test set its 200 object concepts each

contained a single exemplar which was presented 80 times, for a total of 16,000 trials per subject. The experiment used a rapid serial visual presentation (RSVP) paradigm. During image presentation the subjects were asked to perform a target detection task in which they counted the number of times a predefined target image was presented. The data was recorded using a 64-channel EEG system sampled at 1000Hz, which was filtered using a 0.1Hz high-pass filter and a 100Hz low-pass filter, and baseline correction was applied using pre-stimulus recordings. The data finally subsampled to 100Hz, and 17 channels over the occipital and parietal cortex were selected. It should be noted that this dataset is not easily misused in the same way as the Kaneshiro dataset [2], as it explicitly supplies the data as separate training and testing sets with no exemplars appearing in both sets. However, its large size and repetition of exemplars in its test set make it a very useful dataset to investigate if category decoding algorithms can learn exemplar representations when the dataset features a higher number of exemplars per category.

B. Generating Pseudocategories

To demonstrate that category decoding algorithms can learn exemplar representations, we constructed *pseudocategories* from the stimulus sets used in the original datasets. By pseudocategory we mean a collection of exemplars with no semantic relationship to each other, but which are presented to the model as if they were from the same category. Importantly, we ensured that the pseudocategories were balanced in both the number of exemplars they contained, and the distribution of those exemplars among their true categories.

1) *The Kaneshiro Dataset [2] Pseudocategories:* To construct the pseudocategories for this dataset, we selected a single exemplar from each of the 6 categories and combined all trials relating to these exemplars into a single pseudocategory. This was repeated 12 times until all exemplars were used, resulting in 12 pseudocategories. As there are 12 pseudocategories, which are balanced in their number of trials, it follows chance accuracy for a model predicting the pseudocategory of a trial is $\frac{1}{12} = 8.25\%$.

2) *The Gifford Dataset [15] Pseudocategories:* To construct the pseudocategories for this dataset, we used only trials from the predefined test set for each subject as only this set contained multiple trials per exemplar. We then grouped them by their higher-level categories, and excluded any categories which contained fewer than five exemplars, or were deemed to be too similar to other high-level categories. Finally, we evenly divided exemplars from each category into 5 pseudocategories. The pseudocategories each contained: 3 animals, 1 bird, 2 items of clothing, 1 container, 1 electronic device, 6 items of food, 1 musical instrument, 2 items of sports equipment, 3 tools, and 3 vehicles for a total of 23 exemplars per pseudocategory. Fig. 2 shows an example of a pseudocategory constructed this way. Again, as the pseudocategories are balanced in their number of trials, and there are 5 categories, it follows that chance accuracy for a model predicting the pseudocategories of a trial is $\frac{1}{5} = 20\%$.



Fig. 2. Two of the pseudocategories generated for the Gifford dataset [15].

C. Category Decoding Algorithms

To test the ability of classification algorithms to learn exemplar representations, we selected a broad variety of models. Our selection was based on capturing the spectrum of complexity of classification algorithms commonly used in EEG decoding, from simple classical machine learning algorithms to state-of-the-art deep learning models. While many recent models have been developed specifically for single-trial EEG category decoding the reproducibility of these models is limited by the availability of their implementations. Therefore, we selected models which are well known and have implementations available in popular libraries. The models used in this study are:

- Support Vector Machine (SVM) [21].
- Linear Discriminant Analysis (LDA) [22].
- EEGNet [23].
- DeepConvNet [24].
- ShallowConvNet [24].

The k-Nearest Neighbors (k-NN) algorithm [25] was selected to investigate the ability of a simple classical machine learning algorithm to learn exemplar representations. The neural network models were implemented using the PyTorch library [26] with some modifications to the original implementations to

account for the lower temporal resolution of our data relative to the dataset these models were originally trained on. The scikit-learn library [27] was used to implement the SVM, LDA and kNN models. A detailed description of the model architectures and training procedures can be found in the supplementary material.

IV. LEARNING EXEMPLAR REPRESENTATIONS FROM CATEGORIES

A. Model Implementations and Training/Testing Procedures

Within subject pseudocategory decoding was then performed for each of the models on both datasets. A stratified k-fold cross-validation procedure was used to train and test the models. The data was stratified using the exemplar labels to ensure that the distribution of exemplars in each fold was representative of the distribution in the dataset. The number of folds was set to 12 for the Kaneshiro dataset [2] and 10 for the Gifford dataset [15]. The accuracy of the models was then calculated as the proportion of correctly classified trials. This resulted in 120 and 100 accuracy values for the Kaneshiro dataset [2] and Gifford datasets [15] respectively.

V. STATISTICAL ANALYSIS

To determine if the models were learning exemplar representations we compared the decoding accuracy of the models on the pseudocategory labels to chance accuracy using a one-sample t-test. The significance level was set at $\alpha = 0.05$, and then Bonferroni correction was applied to account for multiple comparisons. Since we performed tests for 6 models over 2 datasets this resulted in the significance level being adjusted to $\alpha = \frac{0.05}{6 \cdot 2} \approx 0.0042$.

VI. RESULTS

The results of the statistical analysis are shown in Table I. Starred values indicate that the accuracy was above chance with statistical significance (i.e. $p < 0.001$, given a one-sample t-test with Bonferroni correction). Given that all models achieved decoding accuracies above chance with statistical significance on both datasets, our results show that classification algorithms are indeed able to learn exemplar

representations from categories. While the effect may have been more dramatic for highly complex models, even models such as kNN with only a single parameter can exploit exemplar representations to some degree. Moreover, it is of particular interest that the DeepConvNet had an accuracy of over twice the chance accuracy on the Kaneshiro dataset [2]. This is interesting because, as mentioned previously, the Kaneshiro dataset [2] has been used for category decoding in 12 additional studies following its publication; 10 of which used deep learning models. This suggests that the true accuracy of models affected by the leakage within the literature may differ dramatically from the published results. Box plots of the pseudocategory decoding accuracy are shown for both of the datasets in Fig. 3 and Fig 4. Moreover, it is of interest that all the models performed above chance on the Gifford dataset [15], even though it featured 23 exemplars per pseudocategory. This demonstrates that the classification algorithms are able to exploit exemplar representations even when the number of exemplars per category is much higher than in the Kaneshiro dataset [2], and it is not a flaw which can be easily mitigated by increasing the number of stimuli used in experiments.

VII. FUTURE WORK

While the results of this study demonstrate the ability of classification algorithms to learn exemplar representations given only category labels, further work is needed in recreating a sample of published models with reported decoding accuracies which are suspected to be inflated due to repeated exemplar leakage. This would allow a more comprehensive quantitative analysis of the impact of the leakage and clarify the true state-of-the-art in single-trial category decoding. Moreover, while we know that the models are capable of learning exemplar representations in a category decoding task, it has not yet been demonstrated that the same form of leakage can occur in other forms of EEG analysis. In particular if multiple categories are allowed to appear in a single image then instead of decoding a class we would be detecting the presence of classes. While it seems likely that the same form of leakage would occur in this case, it has not yet been demonstrated, though datasets featuring repeated stimuli and class presence labels are available [28].

VIII. CONCLUSION

The results of this study demonstrate that given the multiple trials per stimulus, and improper model evaluation methods, almost any classification algorithm is capable of learning exemplar representations given only category labels. This has significant implications not only for the fields of neuroimaging and machine-learning, but also for any brain computer interface (BCI) applications which might be developed on the basis of results affected by such leakage. We achieved this using two datasets, the Kaneshiro dataset [2] and the Gifford dataset [15], and a variety of prominent EEG classification algorithms. Use of the former dataset demonstrated that this leakage likely affects several published works. And use of the latter dataset established both the potential of this leakage to

TABLE I
PSEUDOCATEGORY DECODING ACCURACIES

Model	Accuracy (%)
Kaneshiro et al. (2015)	
kNN	10.6809*
SVM	14.5370*
LDA	14.2477*
EEGNet	14.4483*
DeepConvNet	16.6532*
ShallowConvNet	12.8106*
Gifford et al. (2022)	
kNN	22.1076*
SVM	30.1250*
LDA	30.3076*
EEGNet	30.1033*
DeepConvNet	30.5511*
ShallowConvNet	27.8163*




Fig. 3. A box plot of the pseudocategory decoding accuracies for each model trained using the Kaneshiro dataset [2].




Fig. 4. A box plot of the pseudocategory decoding accuracies for each model trained using the Gifford dataset [15].

occur in other datasets, and that classification algorithms may still learn exemplar specific features, even when the number of exemplars per category is much higher than in the Kaneshiro dataset [2]. The significance of our findings is not only in highlighting the potential for leakage in single-trial category decoding studies which use repeated trials per stimulus, but

also in revealing that the true state-of-the-art in single-trial category decoding is uncertain. It seems intuitive that as highest reported performances of models increase on this dataset, the extent to which they are likely affected by repeated exemplar leakage also increases. Therefore, which models have performed best at learning category representations is

currently unclear.

REFERENCES

- [1] P. A. Davis, "Effects of acoustic stimuli on the waking human brain," *Journal of neurophysiology*, vol. 2, no. 6, pp. 494–499, 1939.
- [2] B. Kaneshiro, M. Perreau Guimaraes, H.-S. Kim, A. M. Norcia, and P. Suppes, "A representational similarity analysis of the dynamics of object processing using single-trial EEG classification," *PLOS ONE*, vol. 10, no. 8, p. e0135697, Aug. 2015.
- [3] H. Ahmadiéh, F. Gassemi, and M. H. Moradi, "A hybrid deep learning framework for automated visual image classification using EEG signals," *Neural Computing and Applications*, vol. 35, no. 28, pp. 20 989–21 005, Oct. 2023.
- [4] A. S. Bobe, A. S. Alekseev, M. V. Komarova, and D. Fastovets, "Single-trial ERP feature extraction and classification for visual object recognition task," in *2018 Engineering and Telecommunication (EnT-MIPT)*. Moscow, Russia: IEEE, Nov. 2018, pp. 188–192.
- [5] Z. Jiao, H. You, F. Yang, X. Li, H. Zhang, and D. Shen, "Decoding EEG by visual-guided deep neural networks," in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, Macao, China, Aug. 2019, pp. 1387–1393.
- [6] S. Bagchi and D. R. Bathula, "Adequately wide 1D CNN facilitates improved EEG based visual object recognition," in *2021 29th European Signal Processing Conference (EUSIPCO)*. Dublin, Ireland: IEEE, Aug. 2021, pp. 1276–1280.
- [7] —, "EEG-ConvTransformer for single-trial EEG-based visual stimulus classification," *Pattern Recognition*, vol. 129, p. 108757, Sep. 2022.
- [8] Y. Deng, S. Ding, W. Li, Q. Lai, and L. Cao, "EEG-based visual stimuli classification via reusable LSTM," *Biomedical Signal Processing and Control*, vol. 82, p. 104588, Apr. 2023.
- [9] A. Fares, S.-h. Zhong, and J. Jiang, "Brain-media: A dual conditioned and lateralization supported GAN (DCLS-GAN) towards visualization of image-evoked brain activities," in *Proceedings of the 28th ACM International Conference on Multimedia*. Seattle WA USA: ACM, Oct. 2020, pp. 1764–1772.
- [10] J. Kalafatovich, M. Lee, and S.-W. Lee, "Decoding visual recognition of objects from EEG signals based on attention-driven convolutional neural network," in *2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, Toronto, ON, Canada, Oct. 2020, pp. 2985–2990.
- [11] —, "Learning spatiotemporal graph representations for visual perception using EEG signals," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 31, pp. 97–108, 2023.
- [12] J. Kalafatovich and M. Lee, "Subject-independent object classification based on convolutional neural network from EEG signals," in *2021 9th International Winter Conference on Brain-Computer Interface (BCI)*. Gangwon, Korea (South): IEEE, Feb. 2021, pp. 1–4.
- [13] J. Luo, W. Cui, S. Xu, L. Wang, X. Li, X. Liao, and Y. Li, "A dual-branch spatio-temporal-spectral transformer feature fusion network for EEG-based visual recognition," *IEEE Transactions on Industrial Informatics*, pp. 1–11, 2023.
- [14] M. Yavandhasani and F. Ghaderi, "Visual object recognition from single-trial EEG signals using machine learning wrapper techniques," *IEEE Transactions on Biomedical Engineering*, vol. 69, no. 7, pp. 2176–2183, Jul. 2022.
- [15] A. T. Gifford, K. Dwivedi, G. Roig, and R. M. Cichy, "A large and rich eeg dataset for modeling human visual object recognition," *NeuroImage*, vol. 264, p. 119754, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1053811922008758>
- [16] Princeton Center for Statistics and Machine Learning, "The Reproducibility Crisis in ML-based Science," 28 Jul. 2022, workshop. [Online]. Available: <https://sites.google.com/princeton.edu/rep-workshop/>
- [17] S. Kapoor and A. Narayanan, "Leakage and the reproducibility crisis in machine-learning-based science," *Patterns*, vol. 4, no. 9, p. 100804, Sep. 2023.
- [18] R. Li, J. S. Johansen, H. Ahmed, T. V. Ilyevsky, R. B. Wilbur, H. M. Bharadwaj, and J. M. Siskind, "The perils and pitfalls of block design for EEG classification experiments," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 1, pp. 316–333, 2021.
- [19] A. J. Bell and T. J. Sejnowski, "An information-maximization approach to blind separation and blind deconvolution," *Neural Computation*, vol. 7, no. 6, pp. 1129–59, 1995.
- [20] M. N. Hebart, A. H. Dickter, A. Kidder, W. Y. Kwok, A. Coriveau, C. Van Wicklin, and C. I. Baker, "Things: A database of 1,854 object concepts and more than 26,000 naturalistic object images," *PLOS ONE*, vol. 14, no. 10, pp. 1–24, 10 2019. [Online]. Available: <https://doi.org/10.1371/journal.pone.0223792>
- [21] C. Cortes and V. Vapnik, "Support-vector networks," *Machine learning*, vol. 20, no. 3, pp. 273–297, 1995.
- [22] R. A. Fisher, "The use of multiple measurements in taxonomic problems," *Annals of Eugenics*, vol. 7, no. 2, pp. 179–188, 1936. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1469-1809.1936.tb02137.x>
- [23] V. J. Lawhern, A. J. Solon, N. R. Waytowich, S. M. Gordon, C. P. Hung, and B. J. Lance, "EEGNet: a compact convolutional neural network for EEG-based brain-computer interfaces," vol. 15, no. 5, p. 056013.
- [24] R. T. Schirmmeister, J. T. Springenberg, L. D. J. Fiederer, M. Glasstetter, K. Eggenesperger, M. Tangermann, F. Hutter, W. Burgard, and T. Ball, "Deep learning with convolutional neural networks for eeg decoding and visualization," *Human Brain Mapping*, vol. 38, no. 11, pp. 5391–5420, 2017. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/hbm.23730>
- [25] T. Cover and P. Hart, "Nearest neighbor pattern classification," *IEEE Transactions on Information Theory*, vol. 13, no. 1, pp. 21–27, 1967.
- [26] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019, pp. 8024–8035. [Online]. Available: <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- [27] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [28] J. Xu, B. Aristimunha, M. E. Feucht, E. Qian, C. Liu, T. Shahjahan, M. Spyra, S. Z. Zhang, N. Short, J. Kim, P. Perdomo, R. R. Mao, Y. Sabharwal, M. A. M. Shoura, and A. Nestor, "Alljoined1 – a dataset for eeg-to-image decoding," 2024.