



# Enhancing Representation Learning of EEG Data with Masked Autoencoders

Yifei Zhou<sup>(✉)</sup> and Sitong Liu

George Washington University, Washington, DC 20052, USA  
{yzhou87, sitong.liu}@gwmail.gwu.edu

**Abstract.** Self-supervised learning has been a powerful training paradigm to facilitate representation learning. In this study, we design a masked autoencoder (MAE) to guide deep learning models to learn electroencephalography (EEG) signal representation. Our MAE includes an encoder and a decoder. A certain proportion of input EEG signals are randomly masked and sent to our MAE. The goal is to recover these masked signals. After this self-supervised pre-training, the encoder is fine-tuned on downstream tasks. We evaluate our MAE on EEGEyeNet gaze estimation task. We find that the MAE is an effective brain signal learner. It also significantly improves learning efficiency. Compared to the model without MAE pre-training, the pre-trained one achieves equal performance with 1/3 the time of training and outperforms it in half the training time. Our study shows that self-supervised learning is a promising research direction for EEG-based applications as other fields (natural language processing, computer vision, robotics, etc.), and thus we expect foundation models to be successful in EEG domain.

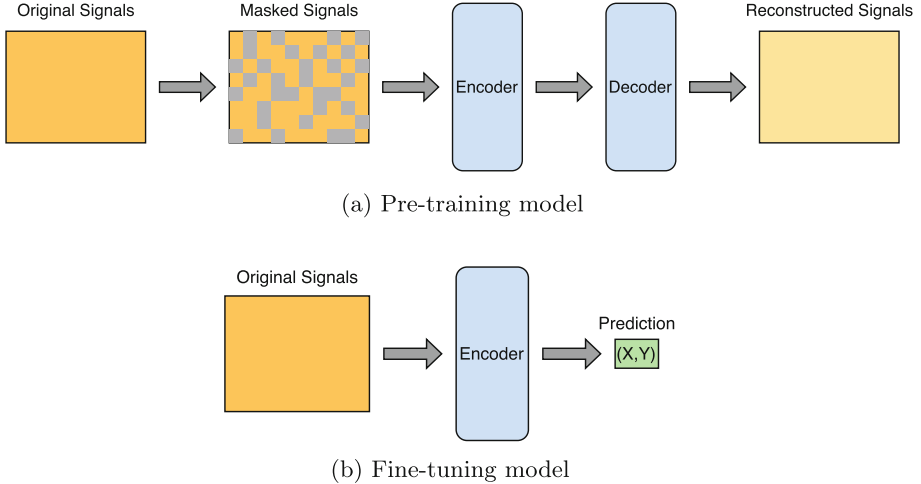
**Keywords:** EEG · Gaze estimation · Self-supervised pre-training · Masked autoencoders

## 1 Introduction

Electroencephalogram (EEG) data, with its rich multidimensional structure, offers unique insights into various neurological phenomena [20]. Understanding the complexities of human brain activity through EEG signals has long been a focal point in neuroscience. EEG-based research holds immense potential of decoding cognitive processes, mental states, and various spatial and temporal aspects of brain functioning. The EEGEyeNet dataset [13], specifically designed for diverse neurocognitive studies, presents a valuable repository for exploring and analyzing EEG data in the context of predictive modeling tasks.

Among the numerous EEG-based tasks, gaze position estimation is a significant challenge due to its relevance in spatial cognition. This task is performed based on the *Large Grid Paradigm* where participants are instructed to focus on a succession of dots that appear one after another, with each dot appearing at one of 25 distinct positions on the screen [13]. The task is to

predict the XY-coordinate of the participant’s gaze position. Accurate decoding of absolute positions from EEG signals holds implications for neurorehabilitation, brain-computer interfaces, and understanding fundamental aspects of spatial awareness.



**Fig. 1. Pre-training and fine-tuning model architectures.** EEG signals collected from multiple channels are arranged into a matrix. **(a)** We mask random elements from the input EEG signal matrix. Our MAE learns to recover these missing signals. **(b)** Our main purpose is to measure the encoder’s performance change after MAE pre-training, so we remove the decoder and fine-tune the encoder to predict gaze positions.

Deep learning methodologies have shown remarkable promise in unraveling intricate patterns within EEG data [1, 26]. Recently, the widely-used Vision Transformer (ViT) model [10] has been proven to be able to significantly improve the accuracy of absolute position prediction [29]. The model proposed by this study, EEGViT, provides further evidence that EEG-based tasks could benefit from computer vision models. EEGViT leverages ViT model weights pre-trained on the ImageNet dataset [9] to achieve state-of-the-art performance, demonstrating that pre-training can contribute to the success of the model in addition to the model architecture [29]. Our study further explores the potential of pre-training to boost the model performance without data augmentation or modifying the model architecture.

Self-supervised pre-training is a prevailing practice to facilitate the representation learning of deep learning models. It helps the models learn useful patterns and representation from the data and thus the models achieve better performance on downstream tasks. In natural language processing (NLP), self-supervised pre-training has been employed to guide large language models to learn contextual information from text corpora [5, 14, 21, 24, 25]. Inspired by

BERT [14], masked autoencoder (MAE) is applied to computer vision models and shown to be successful and scalable vision learners [2, 10, 12].

As a self-supervised pre-training technique, MAE removes certain ratios of content from inputs and tries to reconstruct them. When it is applied in ViT, a certain ratio of input image patches are masked, and the goal is to recover these masked patches [12]. Since EEGViT has shown the capability of the ViT on EEG data, the applicability of MAE on EEG data is worth studying as well. Therefore, our research question is: are MAEs effective brain signal learners? We attempt to answer this question by employing a MAE design that is similar to the one used for ViT pre-training. Our MAE masks random signals from the input EEG signal matrix and reconstructs the missing signals. It has an encoder-decoder architecture (Fig. 1a). The encoder operates on masked EEG signals and learns meaningful latent representations. The decoder then reconstructs the input signals from these latent representations. After pre-training with our MAE, the decoder is removed and the encoder is applied to unmasked EEG signals for gaze position prediction (Fig. 1b).

We compare the performance of the encoder pre-trained with our MAE to the encoder trained from scratch. Experiment results show that MAE pre-training boosts the encoder’s performance on EEGEyeNet gaze estimation task without data augmentation or modifying the encoder architecture. Compared to the encoder without MAE pre-training, the pre-trained one achieves equal performance with 1/3 the time of training and outperforms it in half the training time (Fig. 3). We anticipate that EEG-based applications will benefit more from self-supervised pre-trained deep learning models just as other fields (NLP, computer vision, robotics, etc.), and this even suggests the promising research on foundation models [4, 11, 18, 30, 32] in the EEG domain.

## 2 Related Work

### 2.1 Masked Modeling in Language and Vision

Self-supervised pre-training by masked modeling has brought huge progress to natural language processing (NLP). The masking mechanism in BERT [14] is to randomly mask a certain percentage of the input tokens, and train the model to predict the original token that has been masked out. GPT [5, 21, 24, 25] adopts an autoregressive training approach that predicts the next word in a sentence given all the previous words, which means that during training, the model looks at a part of a sentence and learns to predict the word that comes next. Inspired by the practices in NLP, masked encoding has been applied to visual representation learning [2, 6, 10, 12].

### 2.2 Masked Autoencoder for EEG Data

Various deep learning models such as convolutional neural network (CNN), recurrent neural network (RNN) and Transformer have been applied to EEG

data [3, 8, 16, 19, 26–29, 31]. While supervised learning has been a dominant paradigm of training large deep learning models for a decade, in recent years, self-supervised pre-training by masked modeling has been a great performance booster. A deep learning model pre-trained with masked autoencoders (MAE) often outperforms the same model solely trained with supervised learning. The success of MAE in NLP and computer vision suggests that it is an effective representation learner for both temporal and spatial data. Therefore, it is a natural idea to apply MAE to EEG data.

Previous work has demonstrated the advantage of MAE on EEG-based sleep stage classification [7], seizure sub-type classification [22] and cognitive load classification [23]. The MAEs in these studies reconstruct original features or raw signals from masked *features*. Our study, however, employs a simple approach that reconstructs original EEG signals from masked *signals*. The input EEG signals are directly masked and fed to our MAE without further preprocessing and feature extraction. Experiments have shown that this simple design can still guide our MAE to learn signal representation that is useful for downstream tasks.

### 2.3 EEG-Based Gaze Estimation

EEG-based gaze estimation aims at combining EEG signals with computational techniques to predict the direction or position of a person’s gaze. This approach leverages the fact that certain patterns in brain activity, as captured by EEG, correlate with where a person is looking.

The EEGEyeNet dataset [13] is a comprehensive collection of high-density, 128-channel EEG data synchronized with eye-tracking recordings from 356 healthy adults. This dataset is unique due to its large scale and precise annotation, encompassing over 47h of recording. The third task in the associated benchmark involves determining the absolute position of the subject’s gaze on a screen, described in terms of XY-coordinates. This task is performed using data from the Large Grid paradigm, where participants fixate on a series of dots at different screen positions. It is the most challenging task in the benchmark, aiming to simulate a purely EEG-based eye-tracker. The performance is measured as the euclidean distance in millimeters between the actual and the estimated gaze position. Current performance of deep learning models on this task is presented in Table 4 of [29].

## 3 Methods

We design a masked autoencoder (MAE) that randomly masks signals from the input EEG signal matrix and recovers these missing signals. As shown in Fig. 1, our MAE has an encoder-decoder architecture. The encoder operates on masked EEG signals and learns meaningful latent representations. The decoder then reconstructs the input signals from these latent representations. As the overall goal is to enhance the encoder’s capability to learn useful signal representations,

after MAE pre-training, the decoder is removed and the encoder is applied to unmasked EEG signals to perform downstream tasks. By doing so, we are able to measure the encoder’s performance change after MAE pre-training.

### 3.1 Masking Mechanism

The masking is applied based on the matrix representation of EEG signals. Raw EEG signals are collected from multiple channels. The signals from each channel can be stacked row by row to form a matrix that is suitable for being neural network input [29].

Before an EEG signal matrix is sent to our MAE encoder, a certain proportion of its elements are randomly selected to be set to zero. We implement a simple random selection. Suppose the dimension of EEG signal matrices is  $m \times n$  and the masking ratio is  $r$ . First we generate a random permutation of integers from 0 to  $m \times n - 1$ . Then we select the first  $m \times n \times r$  integers from this permutation as the indices to be masked. Next these selected indices are converted into 2D indices corresponding to the row and column dimensions of the EEG signal matrix. For index  $i$  in the selected indices, its corresponding row index is  $\lfloor \frac{i}{n} \rfloor$  and column index is  $i \bmod n$ . The corresponding elements in the EEG signal matrix will be set to 0.

During training, a mask is generated for each batch and epoch, which means that none of the previously used masks is directly reapplied to the current batch. This will avoid overfitting by ensuring that our MAE can learn as rich local and global patterns as possible. The MAE cannot solve the reconstruction task by simply memorizing the signal values.

### 3.2 Encoder Design

Our MAE encoder is EEGViT [29], a hybrid Vision Transformer (ViT) architecture designed for EEG data. It combines a two-step convolution block [17] with the ViT layers. When the ViT layers are initialized with the model weights pre-trained on ImageNet dataset [9], EEGViT achieves state-of-the-art performance (Table 4 of [29]).

The visual knowledge that ViT learns from large image datasets is beneficial to EEG data as well. However, EEGViT utilizes pre-trained ViT model weights directly for supervised training. We believe that the ViT model can first learn some general EEG signal knowledge before it is applied to a specific task at hand, by which the model can experience a milder transfer from vision domain to EEG. We bridge this gap by using pre-trained ViT weights for MAE pre-training. The ViT layers in our encoder are initialized with the model weights pre-trained on ImageNet dataset. After the encoder learns general EEG signal representation, it will be fine-tuned on downstream tasks.

### 3.3 Decoder Design

Following the MAE for ViT [12], our MAE decoder is a series of Transformer blocks. The reason for this choice resembles the one for vision MAE. Our

reconstruction task is at *signal* level. It requires a low-level understanding of EEG raw signals. A low-level reconstruction task like pixels, or in our case, signals, needs a non-trivial decoder architecture. As described in [12], the decoder design determines the semantic level of learned information. Different decoder structures drive the encoder to extract different levels of signal patterns.

As introduced before, in the fine-tuning stage, only the encoder is kept for supervised training. The MAE decoder assists the encoder with efficient signal encoding, but since our main purpose is to compare the encoder’s performance before and after MAE pre-training, the decoder is not used for downstream tasks.

### 3.4 Reconstruction Task

Our MAE takes in masked EEG raw signals and outputs reconstructed signals. Note that we aim to recover the missing signals, but for implementation simplicity the unmasked signals are also “reconstructed”. That is, our MAE output has the same dimension as the input. Since we only care about the recovery of missing signals, the reconstruction loss is computed on the masked elements of an EEG signal matrix. This practice is similar to previous work [12,14].

Following MAEEG [7], we adopt a similarity loss function<sup>1</sup>:

$$\mathcal{L} = 1 - \frac{\hat{\mathbf{x}} \cdot \mathbf{x}}{\|\hat{\mathbf{x}}\|\|\mathbf{x}\|} \quad (1)$$

where  $\mathbf{x}$  is the original signals and  $\hat{\mathbf{x}}$  is the reconstructed signals.  $\frac{\hat{\mathbf{x}} \cdot \mathbf{x}}{\|\hat{\mathbf{x}}\|\|\mathbf{x}\|}$  computes cosine similarity. Subtracting it from 1 ensures that our MAE learns to minimize the reconstruction loss. Cosine similarity encourages our MAE to capture the intrinsic characteristics of EEG signals. We apply a reversed mask to both the MAE output and full input, so that previously masked positions are now retained and unmasked positions are now set to zero. Then we flatten these two matrices to compute the loss.

## 4 Experiment Setting

We use the EEGEyeNet dataset [13] for MAE pre-training. Then we fine-tune all layers of the MAE encoder on the same dataset.

### 4.1 EEG Data

The EEG data for training our model are from “Large Grid Paradigm” in EEGEyeNet dataset which involves participants fixating on 25 different positions on a screen [13]. EEGEyeNet provides both minimally and maximally pre-processed data. We focus on the minimally pre-processed data. This data includes trials from 27 participants and a total of 21464 samples. Following EEGViT [29], we split 70% of these samples into the training set, 15% into the validation set, and 15% into the test set.

<sup>1</sup> We also experiment with mean squared error (MSE) loss function, the performance increase brought by it is not obvious.

## 4.2 Training

We train our models on Google Colaboratory with 1 NVIDIA A100 GPU. Table 1 shows our training settings. For pre-training, we employ a larger learning rate decay step size and train for more epochs than during fine-tuning. This is because the reconstruction task is more complicated than the downstream gaze estimation task. For fine-tuning, our settings are consistent with EEGViT. The reason is that we use EEGViT model as our MAE encoder, and the goal is to evaluate the encoder’s performance increase solely brought by MAE pre-training. This consistent approach ensures that we are making a fair comparison.

**Table 1.** Pre-training and fine-tuning settings.

	Pre-training	Fine-tuning
optimizer	Adam [15]	Adam
base learning rate (lr)	1e-4	1e-4
batch size	64	64
lr decay step size	10	6
lr decay factor	0.1	0.1
epochs	30	15

## 5 Results

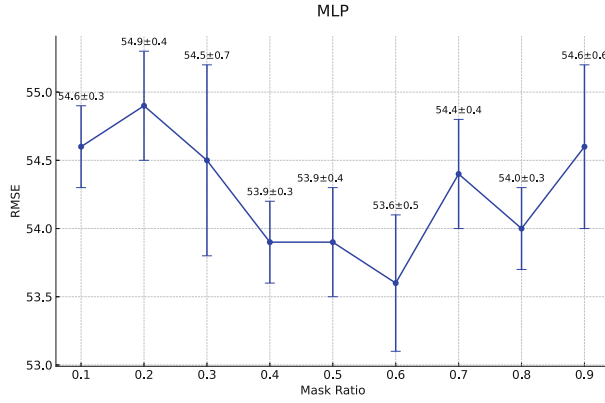
We study the effects of masking ratio and decoder architecture, and report the root mean squared error (RMSE) on the test set. The RMSE is in millimeters (mm). See Sect. 2.3 for details of the gaze estimation task.

Each pre-training epoch takes approximately 2.4 to 2.6 min. A higher masking ratio takes slightly more time. Each fine-tuning epoch takes approximately 2 min.

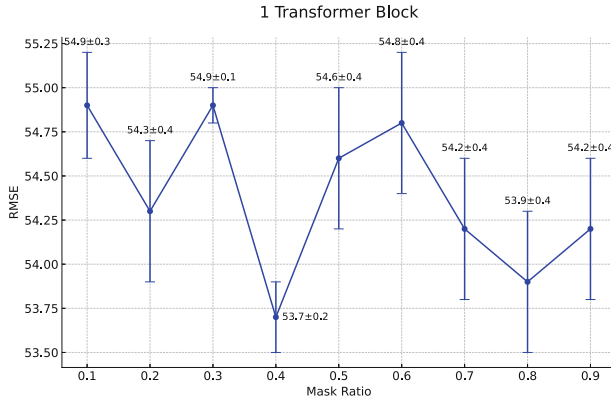
### 5.1 Encoder’s Performance

For MAE pre-training, we experiment with different masking ratios (10%–90%). The MAE decoder has 1 or 2 Transformer blocks. In Sect. 3.3, we hypothesize that our reconstruction task needs a non-trivial decoder architecture. Here, we also use a simple multilayer perceptron (MLP) decoder as a baseline. Table 2 shows the mean and standard deviation over 5 fine-tuning runs. EEGViT’s result is from our experiment<sup>2</sup>. For each decoder architecture, the best result among all the masking ratios is presented in the table. See Fig. 2 for the full results.


<sup>2</sup> Here “EEGViT” is equivalent to “EEGViT Pre-trained” in Table 4 of [29]. This applies to the following mentions as well.



(a) MLP decoder



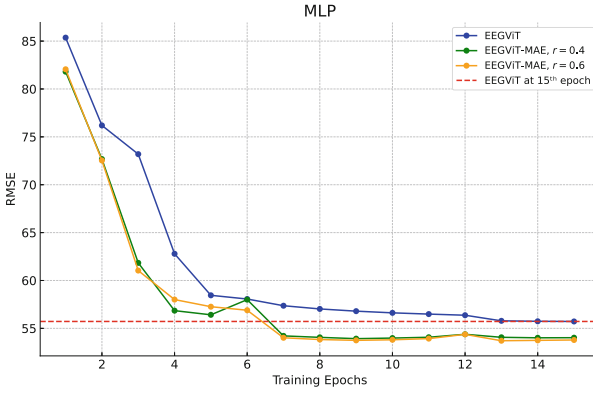
(b) 1 Transformer block decoder




(c) 2 Transformer blocks decoder

**Fig. 2.** Fine-tuning results under different settings.






(a) MLP decoder



(b) 1 Transformer block decoder



(c) 2 Transformer blocks decoder

**Fig. 3. Fine-tuning loss curves.** For each decoder setting, top two results among all the masking ratios ( $r$ ) are presented.

We find that MAE pre-training reduces the encoder’s prediction error without extra hyperparameter tuning. MAE decoder with 2 Transformer blocks achieves the lowest average RMSE. However, the best results of these three decoder architectures are fairly close. From Fig. 2, we see that the encoder’s variance on the gaze estimation task tends to be lower when pre-trained along with more complex decoders, indicating that non-trivial decoder architectures help stabilize the fine-tuning. We also notice that masking 40% of the input signal gives relatively good results in all these three decoder settings. We infer that a masking ratio between 40% and 50% is the optimal choice for our MAE.

**Table 2.** Results from 5 fine-tuning runs.

Model	RMSE (mm)
EEGViT	55.9 $\pm$ 0.7
EEGViT-MAE, MLP	53.6 $\pm$ 0.5
EEGViT-MAE, 1 Transformer Block	53.7 $\pm$ 0.2
EEGViT-MAE, 2 Transformer Blocks	<b>53.5 <math>\pm</math> 0.3</b>

## 5.2 Encoder’s Efficiency

As discussed in Sect. 4.2, our supervised fine-tuning setting is consistent with EEGViT supervised training. We have shown that EEGViT pre-trained with our MAE achieves better results within the same training epochs. This suggests that it adapts faster to the gaze estimation task after MAE pre-training. Figure 3 shows the fine-tuning loss curves. For each decoder setting, top two results among all the masking ratios are presented.

We find that after MAE pre-training, EEGViT achieves better performance with half the training epochs. For masking ratio  $r = 0.4$  in the 1 Transformer block setting and  $r = 0.5$  in the 2 Transformer blocks setting, EEGViT achieves equal performance with 1/3 of training epochs. This demonstrates a significant improvement in learning efficiency. We also observe mild overfitting in EEGViT-MAE models, but it is mitigated in the 2 Transformer blocks setting.

## 6 Discussion and Conclusion

Visual knowledge that is learned from large image datasets like ImageNet can be transferred to the EEG domain, which indicates that these two different signals share some common underlying characteristics. Masked autoencoders (MAEs) are capable of learning useful visual representations. We show that MAEs are effective brain signal learners as well. MAE pre-training is beneficial to downstream tasks in terms of prediction precision and learning efficiency. In this work,

we use the EEGViT model as the MAE encoder. However, we expect MAE pre-training to be a generalizable approach to learn EEG signal representations. The encoder model's choice is flexible. In our future work, we plan to explore alternative encoder models beyond EEGViT to evaluate the generalizability of MAE pre-training. Additionally, we plan to extend the experiments to cover more EEG datasets.

Self-supervised pre-training has been widely explored in NLP and computer vision. Similarly, EEG signal research could take this path by building large and diverse EEG datasets to pre-train deep learning models. These pre-trained models can serve as foundation models [4, 11, 18, 30, 32] for EEG-based applications. They can be fine-tuned on downstream tasks and are expected to obtain superior performance and efficiency compared to models trained solely with supervised learning.

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

## References

1. Altaheri, H., et al.: Deep learning techniques for classification of electroencephalogram (eeg) motor imagery (mi) signals: a review. *Neural Comput. Appl.* **35**(20), 14681–14722 (2023)
2. Bao, H., Dong, L., Piao, S., Wei, F.: Beit: bert pre-training of image transformers. arXiv preprint [arXiv:2106.08254](https://arxiv.org/abs/2106.08254) (2021)
3. Bashivan, P., Rish, I., Yeasin, M., Codella, N.: Learning representations from EEG with deep recurrent-convolutional neural networks. arXiv preprint [arXiv:1511.06448](https://arxiv.org/abs/1511.06448) (2015)
4. Bommasani, R., et al.: On the opportunities and risks of foundation models. arXiv preprint [arXiv:2108.07258](https://arxiv.org/abs/2108.07258) (2021)
5. Brown, T., et al.: Language models are few-shot learners. *Adv. Neural. Inf. Process. Syst.* **33**, 1877–1901 (2020)
6. Chen, M., et al.: Generative pretraining from pixels. In: *International Conference on Machine Learning*. pp. 1691–1703. PMLR (2020)
7. Chien, H.Y.S., Goh, H., Sandino, C.M., Cheng, J.Y.: Maeeg: masked auto-encoder for EEG representation learning. arXiv preprint [arXiv:2211.02625](https://arxiv.org/abs/2211.02625) (2022)
8. Craik, A., He, Y., Contreras-Vidal, J.L.: Deep learning for electroencephalogram (EEG) classification tasks: a review. *J. Neural Eng.* **16**(3), 031001 (2019)
9. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: a large-scale hierarchical image database. In: *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255. IEEE (2009)
10. Dosovitskiy, A., et al.: An image is worth 16×16 words: transformers for image recognition at scale. arXiv preprint [arXiv:2010.11929](https://arxiv.org/abs/2010.11929) (2020)
11. Firoozi, R., et al.: Foundation models in robotics: applications, challenges, and the future. arXiv preprint [arXiv:2312.07843](https://arxiv.org/abs/2312.07843) (2023)
12. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16000–16009 (2022)

13. Kastrati, A., et al.: EEGEyeNet: a simultaneous electroencephalography and eye-tracking dataset and benchmark for eye movement prediction. In: Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1) (2021)
14. Kenton, J.D.M.W.C., Toutanova, L.K.: Bert: pre-training of deep bidirectional transformers for language understanding. In: Proceedings of naacl-HLT, vol. 1, p. 2 (2019)
15. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
16. Kostas, D., Aroca-Ouellette, S., Rudzicz, F.: Bendr: using transformers and a contrastive self-supervised learning task to learn from massive amounts of eeg data. *Front. Hum. Neurosci.* **15**, 653659 (2021)
17. Lawhern, V.J., Solon, A.J., Waytowich, N.R., Gordon, S.M., Hung, C.P., Lance, B.J.: Eegnet: a compact convolutional neural network for EEG-based brain-computer interfaces. *J. Neural Eng.* **15**(5), 056013 (2018)
18. Li, C., et al.: Multimodal foundation models: from specialists to general-purpose assistants, vol. 1, no. 2, p. 2 (2023). arXiv preprint [arXiv:2309.10020](https://arxiv.org/abs/2309.10020)
19. Mao, W., Fathurrahman, H., Lee, Y., Chang, T.: EEG dataset classification using CNN method. In: Journal of Physics: Conference Series, vol. 1456, p. 012017. IOP Publishing (2020)
20. Murungi, N.K., Pham, M.V., Dai, X.C., Qu, X.: Empowering computer science students in electroencephalography (EEG) analysis: a review of machine learning algorithms for EEG datasets (2023)
21. OpenAI, R.: Gpt-4 technical report. arXiv, pp. 2303–08774 (2023)
22. Peng, R., et al.: Wavelet2vec: a filter bank masked autoencoder for EEG-based seizure subtype classification. In: ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 1–5. IEEE (2023)
23. Pulver, D., Angkan, P., Hungler, P., Etemad, A.: EEG-based cognitive load classification using feature masked autoencoding and emotion transfer learning. In: Proceedings of the 25th International Conference on Multimodal Interaction, pp. 190–197 (2023)
24. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., et al.: Improving language understanding by generative pre-training (2018)
25. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. *OpenAI blog* **1**(8), 9 (2019)
26. Roy, Y., Banville, H., Albuquerque, I., Gramfort, A., Falk, T.H., Faubert, J.: Deep learning-based electroencephalography analysis: a systematic review. *J. Neural Eng.* **16**(5), 051001 (2019)
27. Weng, N., Płomecka, M.B., Kaufmann, M., Kastrati, A., Wattenhofer, R., Langer, N.: An interpretable attention-based method for gaze estimation using electroencephalography (2023)
28. Xiao, G., Shi, M., Ye, M., Xu, B., Chen, Z., Ren, Q.: 4d attention-based neural network for EEG emotion recognition. *Cogn. Neurodyn.* 1–14 (2022)
29. Yang, R., Modesitt, E.: Vit2eeg: leveraging hybrid pretrained vision transformers for eeg data. arXiv preprint [arXiv:2308.00454](https://arxiv.org/abs/2308.00454) (2023)
30. Yang, S., Nachum, O., Du, Y., Wei, J., Abbeel, P., Schuurmans, D.: Foundation models for decision making: problems, methods, and opportunities. arXiv preprint [arXiv:2303.04129](https://arxiv.org/abs/2303.04129) (2023)

31. Yi, L., Qu, X.: Attention-based CNN capturing EEG recording's average voltage and local change. In: Degen, H., Ntoa, S. (eds.) HCII 2022. LNCS, vol. 13336, pp. 448–459. Springer, Heidelberg (2022). [https://doi.org/10.1007/978-3-031-05643-7\\_29](https://doi.org/10.1007/978-3-031-05643-7_29)
32. Zhou, C., et al.: A comprehensive survey on pretrained foundation models: a history from bert to chatGPT. arXiv preprint [arXiv:2302.09419](https://arxiv.org/abs/2302.09419) (2023)