

# S-JEPA: TOWARDS SEAMLESS CROSS-DATASET TRANSFER THROUGH DYNAMIC SPATIAL ATTENTION

Pierre Guetschel<sup>1</sup>, Thomas Moreau<sup>2</sup>, Michael Tangermann<sup>1</sup>

<sup>1</sup>Donders Institute for Brain, Cognition and Behaviour, Radboud University, Nijmegen, Netherlands

<sup>2</sup>Université Paris-Saclay, Inria, CEA, Palaiseau 91120, France

E-mail: pierre.guetschel@donders.ru.nl

**ABSTRACT:** Motivated by the challenge of seamless cross-dataset transfer in EEG signal processing, this article presents an exploratory study on the use of Joint Embedding Predictive Architectures (JEPAs). In recent years, self-supervised learning has emerged as a promising approach for transfer learning in various domains. However, its application to EEG signals remains largely unexplored. In this article, we introduce Signal-JEPA for representing EEG recordings which includes a novel domain-specific spatial block masking strategy and three novel architectures for downstream classification. The study is conducted on a 54 subjects dataset and the downstream performance of the models is evaluated on three different BCI paradigms: motor imagery, ERP and SSVEP. Our study provides preliminary evidence for the potential of JEPAs in EEG signal encoding. Notably, our results highlight the importance of spatial filtering for accurate downstream classification and reveal an influence of the length of the pre-training examples but not of the mask size on the downstream performance.

## INTRODUCTION

Electroencephalography (EEG) allows capturing neural activity directly from the scalp, offering a high temporal resolution signal to investigate brain functions. The interpretation of EEG signals with machine learning methods opens the door to build brain-computer interfaces (BCIs). Despite the potential of BCIs, their practical application is hindered by the intensive requirement for calibration data, which is both time-consuming and demanding for participants. Recently, transfer learning has been explored to mitigate the constraints imposed by data scarcity and calibration demands. Self-supervised learning (SSL) presents itself as a pivotal strategy to tackle transfer learning, enabling models to learn rich representations from unlabeled data that can be used to efficiently solve downstream tasks. A particular approach to SSL consists in strategically masking parts of the input data and training a model to predict these masked elements. A key advantage of masking-based SSL methods is that they can be applied to virtually any type of data.

*Masking Strategies* proved a determining factor for the success of SSL methods in various domains, including image [2–4], speech [2, 5], text [2], and video [1] pro-

cessing. For instance, *random masking* strategies, where the regions to be masked are sparsely selected, typically deliver inferior results compared to *block masking* strategies, where larger continuous regions are masked, requiring the model to gain a deeper understanding of the data distribution [2].

*Masked Autoencoders (MAEs)* stand as the quintessential entry point to masking-based SSL [3]. These models aim to reconstruct the masked sections of the input directly, and the training objective is computed by comparing the original input sections to the reconstructed ones. Unfortunately, reconstructing the input and comparing elements in the original space is not without challenges. When the original input space has a high dimensionality, the reconstruction can be computationally expensive, and necessitate the use of domain-specific constraints to produce valid signals. The original signals can also be noisy, increasing the difficulty of encoding meaningful parts of the signal. Moreover, the reconstruction's difficulty can vary significantly across different areas of the input data: with images, reconstructing a monochromatic, non-structured sky is less difficult than reconstructing a structured object like a hand. Such structural disparities are one of the reasons, why conventional metrics like mean-square error often fail to assess the reconstruction quality. These challenges question the scalability and adaptability of MAEs as a universal SSL approach.

*Joint-Embedding Predictive Architectures (JEPAs)* offer a promising alternative to address the limitations associated with direct reconstruction. JEPA-like methods avoid reconstructing the input in its original space and focus on predicting latent representations, or *embeddings*, of the data [4]. This approach confers two major benefits: first, it is computationally efficient, especially with high-dimensional input spaces as the embeddings can reduce the dimensionality; second, the metric's selection in the embedding space is less critical, as the embeddings are learned adaptively to the chosen metric. However, as the "ideal" embedding vectors are unknown, the reconstruction objective is undefined a priori. This challenge is addressed by constructing target embedding vectors during the training by using a bootstrapping procedure which will be further explained in the S-JEPA Framework section. The potential of JEPA-like frameworks has been highlighted by their promising results with images [2, 4],

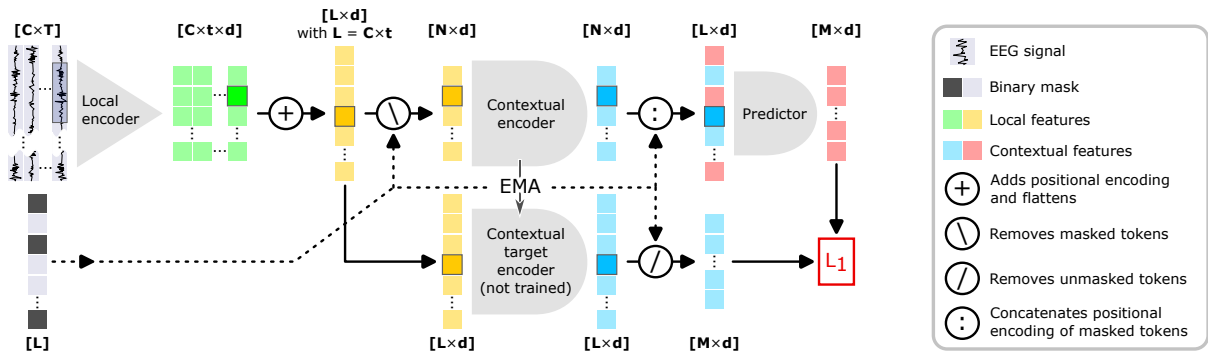


Figure 1: **S-JEPA** training procedure. The framework takes as input EEG recordings with  $C$  channels and  $T$  time samples, and binary masks of length  $L$ . First, the *Local encoder* independently transforms  $t$  windows from each channel into  $C \times t = L$  embedding vectors, called *tokens*, of dimensionality  $d$ . Then, the tokens are marked according to their originating channel and temporal position, and are flattened into a sequence of length  $L$ . Subsequently, only the unmasked tokens are passed to the *Contextual encoder*, while the full tokens sequence is given to the *Contextual target encoder* to generate training targets. Finally, the *Predictor* attempts to reconstruct the masked tokens and its predictions are compared with the targets using an  $L_1$  loss. During the optimisation, the parameters of the Contextual target encoder are not trained via gradient backpropagation but follow those of the Contextual encoder by Exponential Moving Average (EMA). Figure inspired from [1].

speech [2, 5], text [2], and videos [1].

*Applications to the EEG Domain* of masking-based SSL techniques have started to emerge. Pérez-Velasco and colleagues used a masked autoencoder approach, with random masking over the spatial and temporal dimensions [6]. Chien and colleagues experimented with MAEs with block masking over the temporal dimension [7]. Kostas et al. and Foumani et al. also used block masking over the temporal dimension but coupled it with JEPA-like training strategies [8, 9].

*Motivation.* Despite the advancements in applying SSL to EEG data, the exploration of block masking strategies over EEG channels remains uncharted territory. Such an approach holds the potential for developing robust channel attention mechanisms, and thus could facilitate dynamic spatial filtering. This capability could prove instrumental in adapting to recordings with varying channel sets, thereby facilitating cross-dataset transfer learning or tackling corrupted channels. This paper seeks to bridge this gap by investigating the implications of a channel-based block masking strategy on SSL efficacy in EEG signal processing.

The main application of models trained with SSL is the following fine-tuning on the actual task of interest, the so-called downstream task. While most research within the EEG domain has focused on fine-tuning for sleep stage classification tasks, their application to a BCI context remains largely untapped, with only two studies by Kostas et al. and Pérez-Velasco et al. exploring the impact of SSL on BCI tasks [6, 8]. As BCI systems suffer from data scarcity there is a consistent goal to minimize the amount of calibration data required before each online session, specifically as the calibration phase requires sustained attention from the participant and thus is tiring. This highlights a significant opportunity to explore the effectiveness of SSL models across various BCI paradigms, including but not limited to motor imagery protocols, thereby contributing to the broader understanding and ap-

plication of SSL in enhancing BCI performance.

Furthermore, the application of pre-trained SSL models for solving downstream tasks often involves the addition of a linear layer atop the embedding dimension. This practice, however, may not be optimal in high-dimensional embedding spaces. Through a comparative analysis of six different strategies for leveraging pre-trained architectures in downstream tasks, this work aims to support our understanding and application of SSL in EEG data processing.

*Research Questions and Plan.* This manuscript is an explorative study investigating what approaches should be adopted for training SSL algorithms on EEG signals, and what domain-specific considerations warrant attention. We approach this through three research questions: 1) What constitutes the most efficacious masking strategy for SSL when applied to EEG data? 2) How does the temporal length of examples used influence the SSL training process? 3) What fine-tuning strategies lead to the best downstream performance?

To answer these, we first propose a novel masking strategy as part of the Signal-based Joint-Embedding Predictive Architecture (S-JEPA) framework in the S-JEPA Framework section. Then, we introduce fine-tuning strategies tailored for S-JEPA in the Downstream Evaluation section. The datasets used for pre-training and for downstream evaluation are presented in the Datasets section. Finally, the Results and Discussion sections will report and critically discuss the outcomes and implications of the experiments we conducted.

## S-JEPA FRAMEWORK

The Signal-based Joint-Embedding Predictive Architecture (S-JEPA) framework is illustrated in Figure 1. It is used to pre-train models. Its architecture is inspired by BENDR and MAEEG, introduced in the pioneering studies by respectively Kostas et. al. [8] and Chien et. al. [7].

A key modification is introduced in the design of the local encoder to enable our novel masking strategy. This section details the architecture’s components, the spatial masking strategy, and the training process.

*Local encoder.* The local encoder is implemented as a convolutional neural network (CNN) with five layers, each formed of a convolution with a Gaussian error linear unit (GELU) non-linearity. Contrary to the encoders in BENDR and MAEEG, which accept multi-channel input windows, our encoder processes windows from a singular channel. Each window is encoded into a 64-dimensional embedding vector, hereafter referred to as a *token*. The windows are 1.19 s long, with a stride of 1.0 s. The first convolutional kernel covers 0.25 s while the following ones simply combine feature vectors in pairs, i.e., both the kernel temporal lengths and strides are 2.

*Contextual encoder.* This encoder consists of a transformer architecture with eight layers, as introduced by Vaswani and colleagues [10]. It processes the unordered sequence of tokens generated by the local encoder, necessitating the addition of position-encoding information to localize them temporally and spatially. The temporal positioning of each token is defined using a cosine encoding [10] over the first 34 dimensions, whereas spatial positioning is achieved through trainable embeddings for each channel, initialized using cosine encoding based on their three-dimensional coordinates. The contextual encoder receives tokens that only contain local information; its role is to establish relationships between them.

*Spatial Block Masking strategy.* Unique to our methodology is the independent encoding of each channel by the local encoder. While it avoids the learning of spatial filters at early stages, it paves the way for innovative spatial dimension masking strategies. Literature in both, image processing [2] and EEG signal analysis [9] suggests that block masking yields superior results compared to random masking. Motivated by these findings, our novel approach extends block masking to the spatial domain of EEG channels.

Given the irregular distribution of EEG channels, the concept of a contiguous block of tokens does not translate di-




Figure 2: Visualisation of the **spatial block masking** strategy for three example mask centres (red electrodes). The dark to light green spheres represent masks of diameters 40 %, 60 % and 80 % of the head size, as used in our experiments. Assuming a top-down view upon the scalp, the depth of the electrodes is denoted by their intensity (black: close, grey: distant). For a given mask, all electrodes within the corresponding sphere are hidden from the contextual encoder and must be predicted by the predictor.

rectly from its application in images or temporal signals, where the pixels and time samples are regularly spaced. Thus our approach masks all channels within a predetermined radius of a randomly chosen central channel, as illustrated in Figure 2. In this work, we compare three mask sizes with diameters approx. 40 %, 60 % and 80 % of the head size. This strategy inherently introduces variability in the number of masked tokens.

*S-JEPA pre-training.* With the operational principles of both the local and contextual encoders established, we introduce two ancillary components exclusively utilized during the training phase. The first, termed the *Contextual target encoder*, is a non-trainable duplicate of the contextual encoder which serves to generate the training targets. Its parameters are updated via exponential moving average (EMA). The second component, the *Predictor*, is a transformer decoder architecture with four layers as delineated by Vaswani et al. [10]. The comprehensive training methodology is illustrated in Figure 1.

The models are trained until no improvement of the validation loss is observed for 10 epochs (i.e., complete pass through the entire dataset), at which point the best model is saved for subsequent fine-tuning.

## DOWNSTREAM EVALUATION

Upon completion of the network’s pre-training through the SSL task, which holds no intrinsic value beyond training purposes, we proceed to evaluate its efficacy on practical *downstream* classification tasks which, in our case, are BCI tasks. This step is crucial for determining the real-world applicability of the pre-trained model.

*Downstream classification architectures.* Using pre-trained models for downstream classification tasks necessitates altering their architecture to allow predicting class probabilities. The most widely adopted modification, linear probing, consists of adding a linear classification layer directly above the embedding space [4]. However, the individual tokenization of each channel in our approach leads to a high-dimensional latent space, close to the dimensionality of the raw input examples, which would make linear probing inefficient. In response to this challenge, we enrich the architecture with two layers instead which are *Spatial aggregation* and *Fully-Connected*, as explained in Figure 3.

The integration of these layers is explored in three distinct configurations. (a) The **Contextual** downstream architecture places both layers after the contextual encoder as depicted in Figure 3a. (b) The **Post-local** downstream architecture discards the pre-trained contextual encoder and adds the novel layers atop the local encoder as shown in Figure 3b. (c) The **Pre-local** downstream architecture also discards the pre-trained contextual encoder but then places the spatial averaging layer *before* the local encoder as illustrated in Figure 3c. This third alternative allows the network to perform a spatial EEG filtering step, as commonly present in BCI architectures.

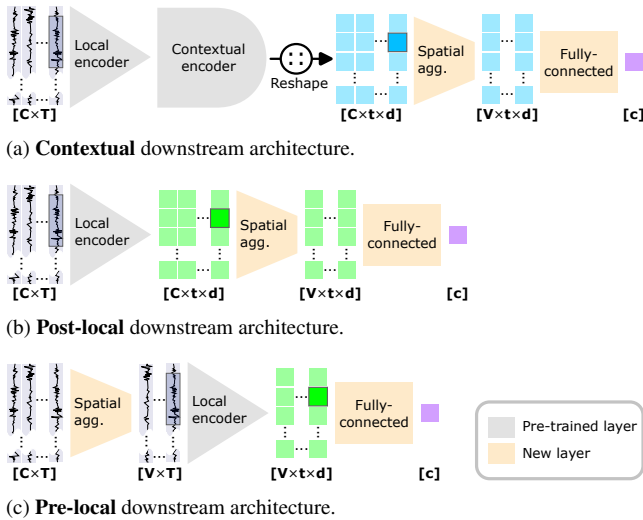


Figure 3: **Downstream classification architectures.** In each of the three alternative alterations of the pre-trained networks, two new layers are added. 1) *Spatial aggregation* is a convolutional layer that realizes weighted combinations of the elements in the channels dimension into  $V \ll C$  "virtual" channels. 2) *Fully-connected* is a linear layer that predicts  $c$  class probabilities.

**Fine-tuning.** We examine two distinct fine-tuning strategies for these downstream architectures. The first method, indicated by the prefix *new-*, involves exclusively training the newly introduced layers, keeping the pre-trained components frozen. The alternative strategy, denoted with the prefix *full-*, consists in fine-tuning the entire network. This second strategy starts with a warm-up phase of 10 epochs where only the newly added layers undergo training, preventing the deterioration of the pre-trained layers' performance due to irrelevant feedback [11], before the previously existing layers are included into the training.

For both strategies, the model is fine-tuned until no improvement of the validation loss is observed for 50 epochs, at which point the best model is restored for testing. Additionally, it should be noted that the temporal length of examples used during the fine-tuning phase is determined by the requirements of the downstream task, which is independent of the length of examples utilized during the SSL pre-training phase.

## DATASETS

For the exploratory investigation in this work, we used the dataset introduced by Lee and colleagues [12], subsequently referred to as the *lee2019* dataset. It contains EEG recordings from 54 subjects, each undergoing three distinct BCI paradigms: steady-state visual evoked potentials (SSVEP) with four classes, visual event-related potentials (ERP), and left vs. right hand motor imagery (MI). The recordings are from  $C = 62$  EEG channels, the spatial distribution of which is detailed in Figure 2. The dataset was loaded from the MOABB framework [13], bandpass filtered at 0.5 - 40 Hz and downsampled to 128 Hz.

We used the first 40 subjects to pre-train any model. The subsequent 7 subjects were used for the validation during this pre-training phase. The remaining 7 subjects were reserved for the downstream performance evaluation.

**Pre-training data.** SSL methods do not necessitate labels. As such, the training and validation examples are slices of the continuous recordings taken at a fixed interval of 16.9 seconds. This study compared three example lengths  $T$  approximately distributed on a logarithmic scale, namely 1, 4, and 16 seconds. All allocated subjects and all paradigms were used collectively during the pre-training phase which yielded a total of 36,576 training examples and 6,528 validation examples. No artefact rejection method was applied.

**Downstream evaluation data.** For each subject allocated for downstream evaluation and each paradigm, the fine-tuning performance of the different pre-trained models is assessed using a 5-folds within-subject stratified cross-validation procedure. The examples used are 4.19 seconds long for MI and SSVEP, and 1.19 seconds long for ERP. These lengths are chosen to respect the durations defined in the original dataset [12] while maintaining compatibility with our tokenization process.

## RESULTS

**Experimental details.** For pre-training, our setup compares all combinations of signal durations (1s, 4s, and 16s) and mask sizes (40 %, 60 %, and 80 % of head size), along with a *no-pre-training* baseline.

For downstream performance evaluation, we assess each pre-trained model using the three downstream architectures—*contextual*, *post-local*, and *pre-local*—and the two fine-tuning approaches—*full* and *new*. The *no-pre-training* is only assessed with the *full* fine-tuning approach. A given pairing of pre-training and fine-tuning configurations is referred to as a *pipeline*.

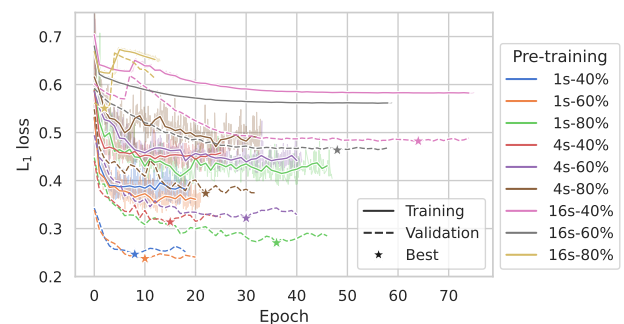


Figure 4: **Pre-training curves** of the different configurations tested. The solid and dashed lines indicate the loss on the training and validation sets. While the validation loss was tested once per epoch only, the training loss was logged after every optimisation step. The train loss on individual optimisation steps is visible in the background, corresponding epoch-wise averages are outlined in white. A star marks the lowest validation loss per curve, the early stopping time point and consequently the checkpoint from which any fine-tuning started.

*Pre-training dynamics.* The pre-training phase, see Figure 4, provides insights into the training process under the S-JEPA framework. It reveals that training curves under the 16s condition are significantly smoother than the other configurations. Additionally, the early stopping mechanism concludes the training of the 16s-80% configuration prematurely at 12 epochs due to an early trough in the loss curve. Similarly, the 16s-40% setup also encounters an early trough yet manages to recover and complete a longer training. The longest training durations are observed in the 16s-40% and 16s-60% configurations, enduring for 74 and 58 epochs, respectively, translating to approximately 12 and 10 hours of training.

*Pipelines ranking on downstream performances.* The comprehensive ranking of all tested pipelines is detailed in Figure 5, offering a comparative overview of the performance of the different combinations over the experimental protocols. Notably, the top-performing pipelines in the downstream tasks are the 16s-60% and 16s-40%

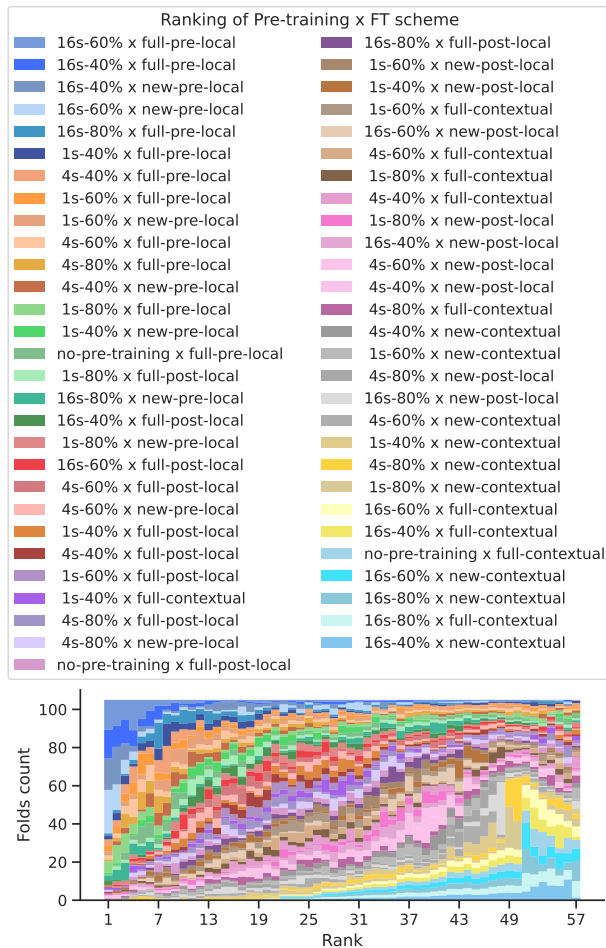


Figure 5: Global downstream classification ranking of all the combinations of pre-training configurations and fine-tuning schemes. Each of the three test datasets has 7 subjects and 5 folds per subject, which makes a total of 105 folds. In the legend, the combinations are ordered according to their average rank over all folds. The vertical span of a coloured "pixel" in the plot represents the number of folds in which this configuration has obtained the rank indicated by the x-axis.

models, especially when paired with any pre-local fine-tuning strategy. They notably occupy the first rank in two-thirds of the cases. In contrast, lower-performing pipelines feature new-contextual or full-contextual fine-tuning, particularly when combined with 16s and no-pre-training configurations.

*Paradigm-wise downstream performances.* on the individual paradigms is reported in Figure 6. For the subsequent analysis, we discern several key observations: 1) The pipelines obtaining the best score on the ERP, SSVEP and MI tasks are respectively 16s-40%-full-pre-local with a 97% AUC, 16s-60%-new-pre-local with a 94% accuracy, and 16s-40%-new-pre-local with a 65% accuracy. 2) Pipelines combining 16s or no-pre-training with contextual architectures frequently result at chance level on average. 3) Pipelines combining 1s or 4s with contextual architectures also perform at chance level on the SSVEP paradigm but above chance level on the MI and ERP ones. 4) Most pipelines manage to achieve respectable scores on the ERP dataset. 5) Only a select subset of pipelines excel on the SSVEP dataset. 6) The MI paradigm scores exhibit notable variability, as indicated by the considerable standard deviation.

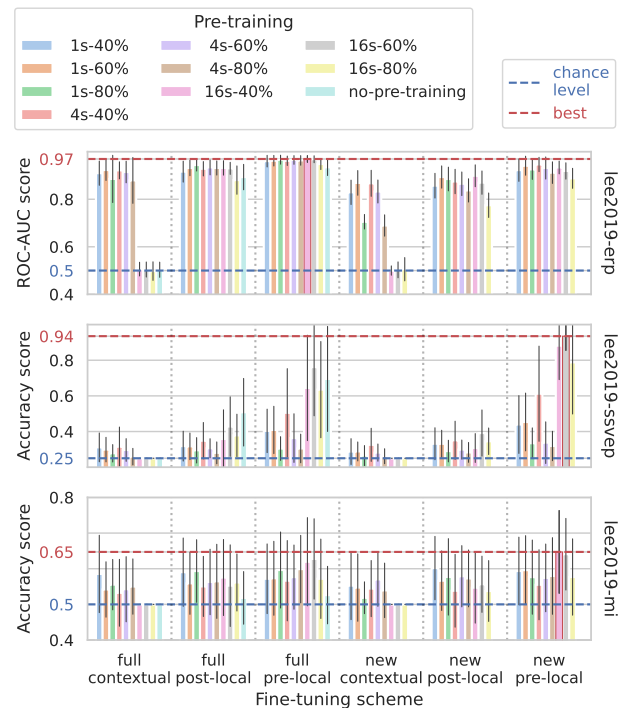


Figure 6: Downstream classification scores of all the pipelines on the three test datasets. The height of the coloured bars corresponds to the average classification score over all the test subjects and cross-validation folds, while the thin black bars correspond to their standard deviation.

## DISCUSSION

In the light of observations made in the Results section, we aim to answer the research questions posed, draw conclusions, and provide guidelines for future research.

*Examples' length largely influences downstream performance.* In particular, the 16s pipelines consistently show the best performance, highlighting the advantage of a longer context during pre-training. Conversely, when considering solely the results using the contextual downstream architecture, the 16s pipelines often perform at chance level. On the other hand, the 1s and 4s pipelines yield better results, which may be attributed to their better alignment of the attention mechanism's training with the short signal windows of the downstream tasks.

*Mask radius' impact on downstream performance uncertain,* as our results do not demonstrate a clear trend between the masks compared. It is possible that the range of masks we compared is not optimal or that other factors are influencing the results. Future work should consider comparing our spatial masking strategy with temporal masking to better understand the relative strengths and weaknesses of these approaches.

*The best fine-tuning strategy implements spatial filtering* The pre-local architecture emerges as the best for downstream classification. This architecture's approach to linearly combine channels before computing features enables effective spatial filtering, thereby enhancing the signal-to-noise ratio. This finding underscores the critical role of spatial filtering in boosting model performance by leveraging the inherent spatial properties of EEG data.

*State-of-the-art comparison.* According to Chevallier and colleagues [14] who benchmarked numerous decoding algorithms across all the datasets available in MOABB [13], the current state-of-the-art (SOTA) performances for within-session classification on the Lee2019 dataset are: ERP at  $98.41 \pm 2.03\%$ , SSVEP at  $89.44 \pm 13.84\%$ , and MI at  $84.74 \pm 13.19\%$ . Notably, all pipelines establishing the SOTA utilize Riemannian geometry. Our approach matches the SOTA for ERP, enhances it for SSVEP, but falls short on MI. A critical difference in our evaluation methodologies should be noted: our downstream evaluation only focuses on the last 7 subjects, unlike their analysis on all 54 subjects. Specifically, Lee and colleagues have identified 6 out of these 7 subjects as hard to classify on the MI task [12]. We believe this exceptionally high rate of challenging subjects might explain our low MI performance.

*Choice of Dataset.* The need for large datasets is paramount when training transformers, potentially explaining the underperformance of the contextual downstream strategy. Although this exploratory study on the Lee2019 dataset provides valuable insights, future research should pivot towards larger datasets to fully harness the capabilities of contextual architectures.

*Conclusion.* This exploratory work introduces a novel masking strategy and three fine-tuning approaches, positioning our method competitively within the realm of BCI tasks. We achieve SOTA performance on two out of three evaluated downstream tasks. Our findings suggest that long pre-training windows favor the local features encoder, while short windows benefit the contextual encoder. Therefore, future research should aim at suc-

cessfully training both the local and contextual encoders. However, no influence of the mask radius on the downstream performance was found. Finally, the best downstream architecture includes a spatial filtering step and discards the contextual encoder.

#### ACKNOWLEDGEMENTS

This work is in part supported by the Donders Center for Cognition (DCC) and is part of the project Dutch Brain Interface Initiative (DBI2) with project number 024.005.022 of the research programme Gravitation which is (partly) financed by the Dutch Research Council (NWO).

#### REFERENCES

- [1] Bardes A *et al.* Revisiting Feature Prediction for Learning Visual Representations from Video. Feb. 2024. arXiv: 2404.08471 [cs].
- [2] Baevski A, Babu A, Hsu WN, Auli M. Efficient Self-supervised Learning with Contextualized Target Representations for Vision, Speech and Language. Dec. 2022. arXiv: 2212.07525 [cs, eess].
- [3] He K, Chen X, Xie S, Li Y, Dollár P, Girshick R. Masked Autoencoders Are Scalable Vision Learners. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE: New Orleans, LA, USA, Jun. 2022, 15979–15988.
- [4] Assran M *et al.* Self-Supervised Learning from Images with a Joint-Embedding Predictive Architecture. Apr. 2023. arXiv: 2301.08243 [cs, eess].
- [5] Baevski A, Zhou Y, Mohamed A, Auli M. Wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations. In: Advances in Neural Information Processing Systems. Curran Associates, Inc., 2020, 12449–12460.
- [6] Pérez Velasco S, Marcos Martínez D, Santamaría Vázquez E, Ruiz Gávez R, Hornero R. Evaluación del Impacto del Aprendizaje Auto-Supervisado en la Precisión de Interfaces Cerebro-Ordenador basadas en Imagenación Motora. Universidad Politécnica de Cartagena, Nov. 2023.
- [7] Chien HYS, Goh H, Sandino CM, Cheng JY. MAEEG: Masked Auto-encoder for EEG Representation Learning. Oct. 2022. arXiv: 2211.02625 [cs, eess].
- [8] 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. Frontiers in Human Neuroscience. 2021;15:653659.
- [9] Foumani NM, Mackellar G, Ghane S, Irtza S, Nguyen N, Salehi M. EEG2Rep: Enhancing Self-supervised EEG Representation Through Informative Masked Inputs. Feb. 2024. arXiv: 2402.17772 [cs, eess].
- [10] Vaswani A *et al.* Attention is All you Need. In: Advances in Neural Information Processing Systems. Curran Associates, Inc., 2017.
- [11] Kumar A, Raghunathan A, Jones R, Ma T, Liang P. Fine-Tuning can Distort Pretrained Features and Underperform Out-of-Distribution. In: ICLR. arXiv, Feb. 2022.
- [12] Lee MH *et al.* EEG dataset and OpenBMI toolbox for three BCI paradigms: An investigation into BCI illiteracy. GigaScience. 2019;8(5):giz002.
- [13] Aristimunha B *et al.* Mother of all BCI Benchmarks. Zenodo. Oct. 2023.
- [14] Chevallier S *et al.* The largest EEG-based BCI reproducibility study for open science: The MOABB benchmark. Apr. 2024. arXiv: 2404.15319 [cs, eess, q-bio].