



HAL
open science

Evaluating the structure of cognitive tasks with transfer learning

Bruno Aristimunha, Raphael de Camargo, Walter Pinaya, Sylvain Chevallier,
Alexandre Gramfort, Cedric Rommel

► To cite this version:

Bruno Aristimunha, Raphael de Camargo, Walter Pinaya, Sylvain Chevallier, Alexandre Gramfort, et al.. Evaluating the structure of cognitive tasks with transfer learning. 2024. <hal-04677674>

HAL Id: hal-04677674

<https://hal.science/hal-04677674v1>

Preprint submitted on 26 Aug 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

EVALUATING THE STRUCTURE OF COGNITIVE TASKS WITH TRANSFER LEARNING

Bruno Aristimunha^{1,2} Raphael Y. de Camargo² Walter H. Lopez Pinaya³
Sylvain Chevallier¹ Alexandre Gramfort¹ Cédric Rommel^{1,4}

¹Université Paris-Saclay, Inria, France

²Federal University of ABC, Santo Andre, Brazil

³King's College London, London, United Kingdom

⁴Valeo.ai, Paris, France

ABSTRACT

Electroencephalography (EEG) decoding is a challenging task due to the limited availability of labelled data. While transfer learning is a promising technique to address this challenge, it assumes that transferable data domains and task are known, which is not the case in this setting. This study investigates the transferability of deep learning representations between different EEG decoding tasks. We conduct extensive experiments using state-of-the-art decoding models on two recently released EEG datasets, ERP CORE and M³CV, containing over 140 subjects and 11 distinct cognitive tasks. We measure the transferability of learned representations by pre-training deep neural networks on one task and assessing their ability to decode subsequent tasks. Our experiments demonstrate that, even with linear probing transfer, significant improvements in decoding performance can be obtained, with gains of up to 28% compare with the pure supervised approach. Additionally, we discover evidence that certain decoding paradigms elicit specific and narrow brain activities, while others benefit from pre-training on a broad range of representations. By revealing which tasks transfer well and demonstrating the benefits of transfer learning for EEG decoding, our findings have practical implications for mitigating data scarcity in this setting. The transfer maps generated also provide insights into the hierarchical relations between cognitive tasks, hence enhancing our understanding of how these tasks are connected from a neuroscientific standpoint.

Keywords Transfer Learning · Deep Learning · Electroencephalography · Event-Related Potential · Brain-Computer Interfaces

1 Introduction

While brain encoding consists in predicting brain activations given a certain stimulus, brain decoding tackles the inverse problem: translating recorded neural activity into its originating stimulus or behavior [1]. This stimulus or behavior can be a visual or auditory element presented, the subject internal mental state (e.g. sleep stage), or the cognitive or motor task being performed during the experiment. Brain decoding has several important applications, such as the diagnosis of neurological disorders [2–4], the detection of seizures [5, 6], automatic processing of polysomnographic recordings [7] and brain-computer interfaces [8–10], among others. Electroencephalography (EEG) is a common and affordable way to record the neural activity in this context [11]. It has the benefit of being non-invasive, having very high time resolution compared to functional magnetic resonance imaging (fMRI) and not requiring a complex and costly infrastructure such as magnetoencephalography (MEG). In recent years, there has been an increasing interest in using deep learning (DL) models for EEG decoding [12]. As an example, DL has been shown to be the gold standard when it comes to automatic sleep stage classification [13–15] and has also demonstrated impressive performances in brain-computer interfaces [16–19].

Unfortunately, DL models are notorious for being data-hungry to extract generalizable discriminative representations. This characteristic can be a problem when it comes to EEG decoding since the acquisition of labelled EEG data remains




Figure 1: Revealing cognitive connections through transfer learning. This graph depicts the transferability of representations used for EEG decoding, capturing the intricate interplay between cognitive tasks. Each node corresponds to a distinct paradigm within the ERP CORE dataset. Arrow width represents the average transfer performance when using the representations learned from a source task to decode a target task. MMN is a great target task for transfer learning, benefiting from representations learned from all source tasks. On the contrary, LRP is mostly a source task: its representations are useful for a wide range of other paradigms.

a constraint, resulting in datasets of limited size. Indeed, EEG annotation often requires a specialist to run experiments or visually inspect recorded signals for specific patterns. In addition, EEG signals have a very low signal-to-noise ratio, especially when we compare to other fields of application which DL thrived, such as image, speech, and text. This characteristic makes EEG decoding even more challenging for DL methods in the context of small datasets.

A common technique in DL for dealing with scarce data scenarios is transfer learning (TL), which consists in applying what you have learned in one context to another [20]. In other words, it can be used to leverage a large dataset to improve the performance in a related smaller dataset, making it a promising technique to alleviate the lack of EEG decoding data. However, TL assumes the knowledge of transferable data domains and tasks, which is not fully understood when it comes to brain data. Indeed, even beyond EEG decoding, understanding hierarchical relations between cognitive tasks remains a core question in neuroscience.

Inspired by the Taskonomy study [21] from the computer vision field, this work investigates the relations between cognitive tasks in an EEG decoding setting. Specifically, we measure the transferability of representations learned by DL models between cognitive tasks. This transferability is measured by pre-training DL models to decode the EEG signals of subjects carrying certain tasks and assessing how well a classifier can reuse the learned representations to decode a subsequent task. We carry extensive experiments with three state-of-the-art decoding models [16–18] trained and evaluated on two recently released EEG datasets, ERP CORE [22] and M³CV [23], containing in total over 140 subjects with 11 distinct decoding modalities. This enables us to create transfer maps capturing the relationships between pairs of cognitive tasks, as presented in Figure 1. From an EEG processing perspective, our maps can be used to leverage related datasets for alleviating EEG data scarcity with transfer learning. We show that even with a linear probing transfer method, we are able to boost by up to 28% the performance of some tasks. From a neuroscientific standpoint, our results broaden our understanding of the connections between cognitive tasks. We discover evidence that some decoding paradigms elicit very specific and narrow brain activities, since no other paradigm transfer well into them. On the other hand, the decoding of some cognitive tasks benefits from pre-training on all other paradigms, thus demonstrating that they rely on a broad range of representations.

We organize the paper as follows. First, section 2 introduces the relevant related work in transfer learning, EEG decoding and cognitive tasks structure investigation. In section 3, we formalize transfer learning for EEG decoding. Then, section 4 outlines our experimental setting, including a description of the two EEG datasets and their respective decoding modalities, as well as our training and data processing protocols. Subsequently, we present the main results of our experiments in section 5, highlighting the significant improvements in decoding performance achieved through transfer learning. Finally, in section 6, we discuss the broader implications of our findings and offer insights into the hierarchical and asymmetric relations between cognitive tasks.

2 Related Works

2.1 Transfer Learning and Taskonomy

Transfer learning refers to the technique of leveraging knowledge acquired from one source domain and task to enhance the performance in another target domain and task [24, 25]. Domain adaptation, domain generalization and self-supervised learning are hence considered sub-fields of TL under this broad definition [25–27]. The most common TL approach consists in fine-tuning on a target dataset a model that was pre-trained on a source dataset. In this context, the main factor determining the success of TL is the relationship between source and target domains and tasks. TL might indeed hinder learning performance in some situations, such as when the source and target are unrelated, which is known as negative transfer [28, 29].

Numerous studies have examined the relationship between tasks for TL purposes [21, 30–34]. One of the most influential works in this area is Taskonomy [21], which investigates these relationships in the context of computer vision tasks. Although we take great inspiration from it, our works differ primarily in the definition of what a "task" is. While Taskonomy tries to uncover the relation between *visual learning tasks* (i.e. given an image, what is the visual label?), we focus in relating *cognitive tasks* (i.e. given the measured brain activity, what is the subject doing?). In this sense, while Taskonomy works with the same input images, analyzing transfer to different output distributions $P_S(Y) \rightarrow P_T(Y)$, our setting is more challenging as we need to transfer between joint distributions of input EEG signals and output decoded stimuli $P_S(X, Y) \rightarrow P_T(X, Y)$. This distinction is made clearer in section 3.

2.2 Transfer learning with brain data

Most EEG decoding works involving transfer learning focus on cross-subject evaluation within the same decoding task [20, 35–46]. Given the extremely high inter-subject variability of EEG signals, this is a crucial question in the design of real-world systems able to generalize from one subject to another. While we study the transfer between different tasks, previous studies consider the same fixed task in different data domains corresponding to each subject or session. They enter in the more specific sub-category of domain adaptation and generalization.

Most related to our work, [34, 47–49] measure the transferability between cognitive tasks, drawing inspiration from the Taskonomy framework [21]. However, these works are all based on fMRI data and most use *encoding* models for very specific types of visual or language stimuli. In contrast, our work uses EEG *decoding* models to compare a broader range of stimuli in different modalities. By working with decoding instead of encoding models, our results are not only useful to understand the relation between cognitive tasks, but also to improve the performance of automatic EEG processing systems in real-world scenarios where data is scarce.

Similar to our work, Qu et al. [49] also uses decoding models to create a transferability map between a large set of cognitive tasks with fMRI data. By working with EEG data rather than fMRI, our analysis can better detect signal-level patterns at the expense of a smaller spatial precision, making it complementary to this previous study. Moreover, while all experiments in [49] were conducted on a single dataset with a small fully connected network architecture, our experimental analysis encompasses consistent results on two very distinct datasets using three different state-of-the-art deep neural network architectures.

2.3 Uncovering cognitive tasks relations with EEG

Beyond the framework of transfer learning, researchers in the EEG decoding community have been interested in identifying the latent structure between cognitive tasks by other means. For example, Rommel et al. [50] analyzed what data augmentations work best for decoding different cognitive tasks. They showed that some augmentations work well for many different tasks, demonstrating that they share some common invariances, while others only improve the decoding performance in some specific cases. Although some structure arises as a by-product of this study, it focuses on task invariances. Our work differs substantially from [50] since our analysis directly connects one cognitive task to another based on transferred predictive performance instead of invariances.

As another example, studies conducted by Banville et al. [51, 52] have demonstrated that self-supervised representations can encode clinically-relevant structures from EEG data, such as the sleep stages, pathology, age, apnea and gender information, without any access to such labels. A similar outcome was observed in [53, 54]. While these methods are capable of learning representations useful for different downstream tasks, these works do not study transfer to different domains and datasets in practice. Our work focuses not only on discovering the structure between predefined supervised tasks but also quantitatively assesses the transferred performances into new data domains and tasks.

3 Method




Figure 2: **A.** Data splitting and alignment. Source and target tasks correspond to different ERP and BCI paradigms; **B.** EEG decoding models as a representer network and a classification head; **C.** Standard EEG decoding training and evaluation; **D.** Transfer with linear probing. Only the classification head h_ψ is re-trained, while the representer network r_{w_S} trained on the source task is kept intact.

3.1 EEG Decoding

From a machine learning perspective, EEG decoding is defined as a classification problem using a dataset of N pairs of trial recordings and decoding labels $\mathcal{E} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$. Each trial recording $\mathbf{x} \in \mathcal{X} = \mathbb{R}^{m \times c}$ is modelled as a matrix with m rows and c columns, where m is the number of time-steps and c is the number of electrodes. The decoding labels $y \in \mathcal{Y}$ correspond to the class of stimulus to predict, where \mathcal{Y} is a discrete set containing 2 or 3 classes in the decoding scenarios studied in this paper. We assume the dataset is randomly split into a training set $\mathcal{E}_{\text{train}}$ of size N_{train} and test set $\mathcal{E}_{\text{test}}$ of size N_{test} . The decoding problem amounts to using the training set $\mathcal{E}_{\text{train}}$ to learn a model $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ with parameters θ mapping each trial \mathbf{x} to the associated label y . This is achieved by optimizing the model parameters to minimize an average loss function ℓ across the training dataset:

$$\min_{\theta} \frac{1}{N_{\text{train}}} \sum_{\mathcal{E}_{\text{train}}} \ell(f_\theta(\mathbf{x}_i), y_i) . \quad (1)$$

In all our experiments, the loss function ℓ used is the balanced cross-entropy. The generalization performance of the trained model is assessed on the test set $\mathcal{E}_{\text{test}}$ (*c.f.* Figure 2 A, B, C).

3.2 Transfer learning

More generally, we assume that our dataset corresponds to observations of a joint random variable (X, Y) valued on a space $\mathcal{X} \times \mathcal{Y}$, where some decision function f exists such that $Y = f(X)$. Hence, problem (1) amounts to trying to approximate the decision function f by optimizing the model f_θ to fit the training data sampled from $P(X, Y)$.

Following [25], a *domain* is defined by a feature space and inputs distribution $\mathcal{D} = \{\mathcal{X}, P(X)\}$. Likewise, a task consists of an output space and a decision function $\mathcal{T} = \{\mathcal{Y}, f\}$. Transfer learning aims to use the knowledge extracted from a source domain \mathcal{D}_S and task \mathcal{T}_S to improve the performance of a model f_θ on another target domain \mathcal{D}_T and task \mathcal{T}_T . Both domains and tasks are materialized by source and target datasets $\mathcal{E}_S = \{(x_i, y_i) \sim P_S(X, Y); Y = f_S(X)\}$ and $\mathcal{E}_T = \{(x_i, y_i) \sim P_T(X, Y); Y = f_T(X)\}$.

In our context of EEG decoding, the source and target datasets correspond to different ERP and BCI cognitive tasks performed by the same cohort of subjects in comparable experimental settings (*c.f.* section 4). Hence, while source and target domains share the same feature space $\mathcal{X} = \mathbb{R}^{m \times c}$, they differ in terms of marginal distributions of input trial recordings $P_S(X) \neq P_T(X)$. Regarding learning tasks, source and target datasets differ in terms of decision functions $f_S \neq f_T$ and output spaces $\mathcal{Y}_S \neq \mathcal{Y}_T$. As opposed to our setting, Taskonomy [21] works with a common source and target domains $(\mathcal{X}_S, P_S(X)) = (\mathcal{X}_T, P_T(X))$.

3.3 Transferability through linear probing

We evaluate the transferability between source and target datasets through *linear probing* [55, 56], as described below. We assume that f_θ is a neural network made of two parts: a representer model $r_\omega : \mathcal{X} \rightarrow \mathcal{R}$, with parameters ω , and a classifier head $h_\psi : \mathcal{R} \rightarrow \mathcal{Y}$, with parameters ψ

$$f_\theta(x) = h_\psi(r_\omega(\mathbf{x})). \quad (2)$$

The parameters θ of the model are hence the concatenation of ω and ψ . While the representer r_ω is responsible for learning useful representations for the learning task, the classifier head h_ω is just the last linear layer of the network used to deliver the classification decision. When assessing the transferability between a source and target datasets (*c.f.* Figure 2):

1. We first train the whole model $f_{\theta_S} = h_{\psi_S} \circ r_{\omega_S}$ on the training split of the source dataset $\mathcal{E}_{S,\text{train}}$ by solving equation (1) ;
2. Then we freeze the representer parameters ω_S and retrain the classifier head h_{ψ_T} from scratch on the training split of the target dataset $\mathcal{E}_{T,\text{train}}$:

$$\psi_T \in \arg \min_{\psi} \frac{1}{N_{\text{train}}} \sum_{\mathcal{E}_{T,\text{train}}} \ell(h_\psi(r_{\omega_S}(\mathbf{x}_i)), y_i); \quad (3)$$

3. Finally, we evaluate the network obtained $h_{\psi_T} \circ r_{\omega_S}$ on the test split of the target dataset $\mathcal{E}_{T,\text{test}}$ and use this metric to assess the transferability between S and T .

We evaluate the transferability through linear probing since it assesses whether the representations learned from the source dataset allow us to classify the target data. In contrast, fine-tuning (*i.e.* pushing the training of the whole model h and r further with target data) would modify the learned source representations, which would complicate the analysis [57].

4 Experiments

4.1 ERP CORE and M³CVdatasets

During EEG decoding experiments, subjects perform cognitive tasks with stimuli that evoke specific brain signatures. When the nature of the stimuli is similar, they can be categorized into the same paradigm. EEG decoding studies have been interested in a very large and diverse number of paradigms, which can be categorized as exogenous (where an external stimulus is used *e.g.* event related potential) or endogenous (where stimuli are induced by a predetermined mental task or behavior, *e.g.* motor imagery) [8, 58]. Only a small number of EEG datasets contain recordings in a diverse set of paradigms with the same subjects and configurations. Most existing datasets include a limited number of subjects [59, 60] or a limited number of cognitive tasks [61–63].

In our study, we use two of the few EEG datasets that explore a large diversity of paradigms with a single large cohort of subjects in comparable experimental settings. The first dataset, ERP CORE [22], comprises 40 subjects (25 females

and 15 males between 18 and 30 years old). It focuses on exogenous paradigms, featuring seven isolated tasks eliciting specific event-related potential components (ERP), namely Active Visual Oddball (P3b), Word Pair Judgement (N400), Face Perception (N170), Passive Auditory Oddball (MMN), Flankers (LRP and ERN) and Simple Visual Search (N2pc).

The second dataset, M³CV [23], is a large multi-task, multi-session, multi-subject investigation of EEG commonality and variability. It includes 106 subjects who performed specific tasks in six different paradigms. In this work, we only focus on trials for which the subject and task labels were available, namely: Motor Execution (ME), Transient-State Sensory (TSS), Resting-State (RS) and Steady-State Sensory (SSS). This reduces our dataset to 95 subjects (22 females and 73 males between 19 and 24 years old) corresponding to the enrolment and calibration subsets of the original dataset. The different paradigms for both datasets are presented in Figure 3. Other details about these datasets are also listed in Table 1.

Table 1: Dataset, Cognitive task, time-locking event, time window and number of classes for all the sub-datasets in ERP CORE and M³CV.

| Dataset | Cognitive Task (\mathcal{E}) | Time-Locking Event | Window size (ms) (m) | # classes (\mathcal{Y}) |
|-------------------|----------------------------------|--------------------|--------------------------|-----------------------------|
| ERP CORE | N170 | Stimulus-locked | -200 to 800 | 2 |
| ERP CORE | MMN | Stimulus-locked | -200 to 800 | 2 |
| ERP CORE | N2pc | Stimulus-locked | -200 to 800 | 2 |
| ERP CORE | N400 | Stimulus-locked | -200 to 800 | 2 |
| ERP CORE | P3 | Stimulus-locked | -200 to 800 | 2 |
| ERP CORE | LRP | Response-locked | -800 to 200 | 2 |
| ERP CORE | ERN | Response-locked | -600 to 400 | 2 |
| M ³ CV | ME | Stimulus-locked | 0 to 1000 | 3 |
| M ³ CV | RS | Stimulus-locked | 0 to 1000 | 2 |
| M ³ CV | SSS | Stimulus-locked | 0 to 1000 | 3 |
| M ³ CV | TSS | Stimulus-locked | 0 to 1000 | 2 |

4.2 Defining the decoding labels

The two EEG datasets under investigation offer many options to define the decoding labels y . For instance, one could consider the subject’s response accuracy to the stimulus or whether the answer was in the appropriate time frame. Another possible label definition involves predicting the stimulus presented based on the recorded signal or the presence or absence of the task component [22, 64]. We describe hereafter how we defined the decoding labels in our experiments, starting with the ERP CORE dataset.

The *Flankers* task aims to elicit two possible ERP components: *ERN* and *LRN*. ERN characterizes situations in which the subject makes an error, even when they are not consciously aware of them [65]. It is elicited here by asking subjects to indicate whether they saw a left or right trigger. We assigned one label to trials where the subject’s response matched the target stimulus, and another label when it did not. If the trigger and response were both left or right, the trial was labelled as correct; otherwise, it was incorrect [22, 66, 67].

The same Flankers task is used to elicit the *LRN component*, which is related to the side of the response hand. Unlike ERN, this component is associated with preparing the response rather than its execution, as reported in previous studies [66, 67]. Hence, we relied solely on the subject’s response to determine the label for each trial in this case. Specifically, if the subject answered with the left hand, we assigned one label; if they responded with the right hand, we assigned a different label.

The *Word Pair Judgment* task is used to elicit the *N400* component. It consists in presenting participants with a prime word, followed by a target word and asking them to judge whether they are semantically related or not [68]. To define the label for each trial, we use the stimulus definition. Specifically, if the presented words were semantically related, we assigned one label; if they were unrelated, we assigned a different label.

The *N2pc* component is elicited through a *Simple Visual Search* task [69]. In this task, participants had to visually identify a specific target item (in blue or pink) among multiple distractors (in black). Notably, each item is an outlined square with a gap on top or bottom, and subjects had to indicate the position of the gap on the target item (*c.f.* Figure 3). Despite the task assigned to the subjects, what is decoded in this experiment is the position of the target within the screen, *i.e.* whether it appeared on the left or right side, since this is what the N2pc component really captures. As usually done for this experiment [64], we excluded trials with incorrect responses to ensure that the subjects were paying attention to the target.

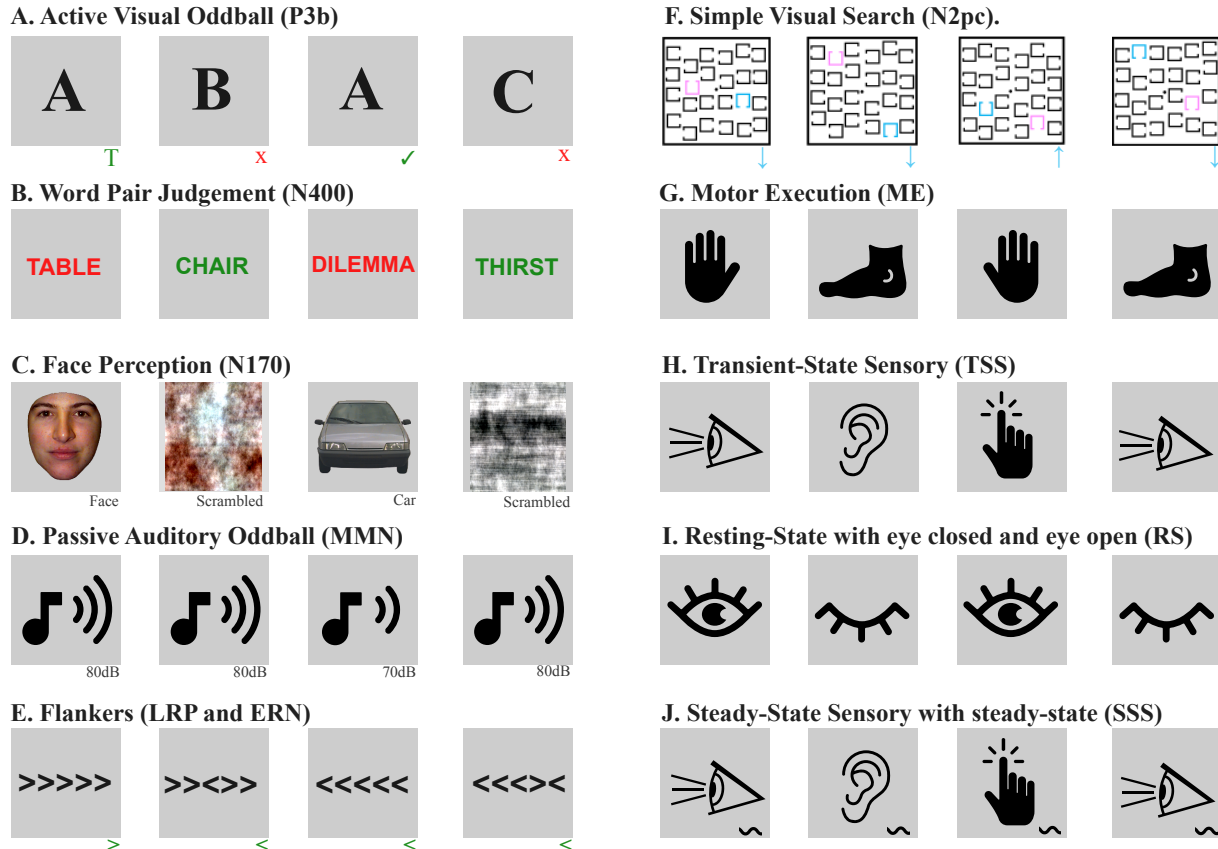


Figure 3: Illustration of experiments recorded in ERP CORE (A to F) and M³CV datasets (G to J). Labels under the stimuli correspond to what subjects are supposed to answer during the trials. They do not necessarily correspond to decoding labels. **A.** In the active visual oddball task P3b, participants viewed random letters and responded to whether each stimulus was matched to the target (T) or not. One of the letters was designated as the target for each block of trials, with a probability of 0.2. **B.** The word pair judgment task N400, involved participants viewing a red prime word followed by a green target word and indicating whether the target word was semantically related or unrelated to the prime word. **C.** During the face perception task N170, participants viewed images of faces, cars, scrambled faces, or scrambled cars and indicated whether the stimulus was an "object" (face or car) or a "texture" (scrambled face or scrambled car). **D.** The passive auditory oddball task MMN involved participants watching a silent video while speakers played standard and deviant tones. Deviant tones had a probability of 0.2. **E.** In the flanker's task, participants indicated the direction of the central arrowhead, which was surrounded by arrowheads pointing in the same (congruent trials) or opposite (incongruent trials) directions. This task elicited the lateralized readiness potential (LRP) and the error-related negativity (ERN). **F.** In the simple visual search task N2pc, either pink or blue was designated the target colour at the beginning of each trial block. Participants had to indicate whether the gap was on the top or bottom of the coloured square. **G.** During the Motor Execution task (ME), participants responded to a visual cue by *executing* movements with their left hand (LH), right hand (RH), or right foot (FT). **H.** Transient-State Sensory (TSS) potentials were elicited by visual, auditory, and somatosensory stimuli. All three categories of stimuli were arranged in random order and lasted 50 milliseconds each. **I.** In the Resting-State task (RS), participants completed two runs with eyes closed (EC) and two with eyes open (EO), each lasting one minute. **J.** The Steady-State Sensory (SSS) task involved separate trains of visual, auditory, and somatosensory stimuli, with different stimulation frequencies and recording times for each run.

The *N170 component* [70] was elicited by presenting stimuli in the form of cars, faces, and deformed versions of these objects, in what is called the *Face Perception* task. Participants had to identify whether the presented stimulus was intact or scrambled. To simplify the analysis, we excluded scrambled stimuli as recommended by the original authors of ERP CORE [64], and labelled the remaining ones according to whether they represented a car or a face.

Concerning the *Active Visual Oddball* task, participants were asked to watch a sequence of random letters among A, B, C, D or E. The first letter in a trial had the role of target, and subjects had to answer whether the following letters

matched the target or not, which should elicit the *P3b component* [71]. To assign trial labels for this task, we considered two factors. First, we checked whether the letter on the screen matched the target letter. Second, we verified whether the participant provided the correct answer or not. If both conditions were satisfied, we assigned one label to the trial. If not, we assigned another label.

Finally, the *Passive Auditory Oddball* task aims to elicit the *MMN component* [72]. In our study, the stimuli were used to define the trials labels. Specifically, we assigned one label if the individual heard a standard tone at 80 dB and another label if they heard a deviant tone at 70 dB.

In the M^3CV dataset, the definition of labels was more straightforward. For instance, decoding labels for the *RS* task simply corresponded to whether the subjects had their eyes open or closed. For the *TSS* and *SSS* tasks, labels were directly defined based on the stimulus presented to the subjects, *i.e.* whether it was a visual, auditory or somatosensory stimulation (3-class classification problem). Likewise, decoding labels of the *ME* task corresponded to the movement being executed by the subjects, *i.e.* either the right foot, the right hand or the left hand (3-class classification problem).

4.3 EEG pre-processing and epoching

Both datasets were pre-processed following the authors’ recommendations in all our experiments. Namely, ERP CORE recordings were filtered between $[0.5 - 40]$ Hz with overlap-add FIR filtering. Electric potentials were referenced on the average of electrodes P9 and P10 for all tasks except N170, for which we used the average of all 33 electrodes as commonly done in the literature [22]. Stimuli events were shifted 26 ms forward in time to account for the LCD monitor delay in the MMN task. We downsampled the data from 1000 Hz to 250 Hz and used ICA [73] to correct artefacts and discard particularly bad trials. Finally, trials were cropped into 1000 ms windows based on the stimulus or response depending on the task, following values reported by [22]. This dataset was entirely pre-processed using the MNE-PYTHON library [74].

For the M^3CV dataset, we used the dataset pre-processed by the authors as described below. Signals were band-pass filtered between $[0.01 - 200]$ Hz using a 4th order Butterworth filter and notch filtered between $[49 - 51]$ Hz. Potentials were referenced on the average of electrodes TP9 and TP10. Visual inspection and ICA were used to remove artefacts, and bad channels were replaced by the average of the three neighbouring channels. Finally, 1000 ms signals were cropped from each trial. More pre-processing details can be found in [23].

4.4 Data splitting

The datasets were randomly split into a training, a validation, and a test set with respective proportions of 56%, 24% and 20%. Each split contains data from different subjects since we want to assess the cross-subject generalization of our models. Trainings and evaluations were repeated with different splits following a 5-fold cross-validation scheme.

In the standard decoding experiments (subsection 5.1), models were trained and evaluated using data from the same cognitive task (*c.f.* Figure 2 C). In the transfer learning experiments (subsection 5.2), models were pre-trained on training subjects of some task *A*, then fine-tuned through linear probing on data from the same subjects performing a different task *B* and finally evaluated on unseen test subjects carrying out this same task *B*, as described in subsection 3.3 and illustrated in Figure 2 D. Thanks to this data splitting alignment across cognitive tasks, we ensure that the test subjects remain the same after transfer, avoiding any leakage between test and training sets.

4.5 Decoding models

We evaluated the transferability of learned tasks using three state-of-the-art deep learning models: ShallowNet [18], EEGNet [16], and EEGInception [17]. ShallowNet is an efficient network inspired by the Filter Bank Common Spatial Pattern method [75]. It contains $36k$ trainable parameters. EEGNet model has a convolutional layer for channel-based EEG data filtering, a depth-wise convolutional layer acting as a spatial filter across channels, a separable convolutional layer for categorical feature extraction, and a fully connected layer as classification head. Finally, EEGInception has $15k$ parameters and extracts low- and high-frequency features in parallel at different scales. These architectures delivered state-of-the-art results in several EEG decoding tasks and datasets [76–81]. In all our experiments, we used the implementations of these models available in the BRAINDECODE library [18] with default hyperparameter values.

As baselines, we also tested Riemannian methods, which are known as the best machine learning techniques for BCI and ERP, together with DL models. To this end, we evaluated both Minimum Distance to Riemannian Mean (MDM) [82, 83] and Tangent Space methods with either logistic regression or support vector machine classifiers. We also evaluated these methods with a varied set of covariance matrices, such as ERPCovariances [84] and Xdawn [85]. These methods were implemented and trained using the MOABB and PYRIEMANN library [86, 87].

4.6 Training setting

All DL models were initialized with the Xavier sampling [88] and trained with the AdamW optimiser [89] using default parameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$ and a weight decay of 5×10^{-4} . The initial learning rate was set to 10^{-4} for ShallowNet and EEGInception, and to 6.25×10^{-4} for EEGNet. Trainings lasted at most 200 epochs and an early stopping [90] procedure with a patience of 50 epochs was used. The training of all deep learning models was carried out using PYTORCH and BRAINDECODE libraries [18, 91] on an Nvidia DGX with 4 A100 boards.

5 Results and Findings

5.1 Decoding performance

We first analyse the decoding performance of studied models on ERP CORE and M³CV datasets in a standard decoding setting, without any transfer. Overall, DL methods deliver superior balanced accuracies than baseline methods, as shown in Figure 4.

Their scores also exhibit lower standard deviations than machine learning methods, demonstrating the stability of learned representations across subject splits. Notably, EEGNet consistently led to the best performance on ERP CORE for all paradigms, whereas ShallowNet outperformed others on three out of four paradigms of the M³CV dataset. Also, note that our results on ERP CORE are consistent with those reported in [92, 93].

To support this analysis, we performed a permutation signed-rank test for each model pair within each paradigm to determine whether observed performance gaps are significant (Figure 5). The resulting p-values were combined using Stouffer’s method, with a weighting given by the square root of the number of subjects, and a Bonferroni correction was applied to account for multiple comparisons, as done in [86]. The standardized mean difference was calculated within each dataset to determine the effect size. Examining Figure 4(b), we observe that EEGNet slightly outperforms ShallowNet, although the difference is not statistically significant, as highlighted in Figure 5(b).

From a paradigms perspective, the most challenging task within the ERP CORE dataset is the MMN paradigm, which is unique in that the subject does not have an active response moment. Conversely, the least difficult tasks correspond to the ERN and LRP.

Overall, the best DL architectures (EEGNet for ERP CORE and ShallowNet for M³CV) lead to similar performances across paradigms, even in scenarios with different numbers of classes (*e.g.* two classes in RS vs. three classes in ME, SSS and TSS). This suggests that the results are robust and confirms that the selected decoding architectures are well-suited to the different experimental paradigms studied in this work.

5.2 Transfer learning performance

5.2.1 ERP CORE

We now analyse the transfer performance between paradigms to quantify how transferable each cognitive task is in relation to one another. Figure 6 shows the obtained balanced accuracy for each pair of source and target cognitive tasks in the ERP CORE dataset. A first intriguing observation is the asymmetry of the transfer matrices obtained, regardless of the model used, indicating that the transferability between tasks is highly directional. We can see that some tasks do not transfer well, leading to accuracies close to chance (50% in this dataset, as all classification problems are binary). Notably, no source task is correctly transferring to either to LRP or N400 paradigms. This means that these cognitive tasks rely on very specific representations that are not elicited by the other tasks and are hence erased by models pre-trained on other paradigms.

Conversely, representations learned on some source tasks seem useful for decoding others, at least to a level of accuracy comparable to a model fully trained and evaluated in a standard fashion (diagonal values in the figures). For example, LRP proves to be a good source task when transferring to ERN and MMN. This suggests that LRP elicits some brain activations common to ERN and MMN, which are learned by the decoding networks into its hidden representations.

More strikingly, pre-training on some source tasks, such as ERN, improves performance on some other target tasks, like N170, even *exceeding* the reference performance of a model fully trained on this target task. The best example of this is the MMN task, whose performance is boosted beyond the reference accuracy when transferring *from all possible source tasks*. With ShallowNet, for example, we observe improvements of 28.3%, 6.1%, 4.6%, 3.7%, 2.7%, and 2.2% for the LRP, N2pc, ERN, N400, P3, and ERN source tasks, respectively. This is also observed for EEGNet and EEGInception, although more discretely.





Figure 4: Cross-subject balanced accuracy across paradigms. Error bars correspond to the standard-deviation across 5-fold cross-validation. DL methods outperform machine learning baselines on most paradigms. EEGNet is consistently better than other architectures in ERP CORE. ShallowNet leads to the best scores across M³CV paradigms, except for TSS.

In order to better visualize these complex connections between tasks, we processed the matrices from Figure 6 into transferability maps, shown in Figure 1. Arrows' widths are proportional to corresponding transfer accuracies, where performances close to chance level were omitted. More precisely, they correspond to transferability scores $s_{S,T}$ for each source and target tasks, obtained by linearly rescaling the corresponding accuracies $a_{S,T}$ so that 0 corresponds to the chance level c_T and 1 corresponds to the reference accuracy without transfer (*i.e.* matrix diagonal $a_{T,T}$):

$$s_{S,T} = \frac{\max(a_{S,T} - c_T, 0)}{a_{T,T} - c_T}. \quad (4)$$

The Scalable Force-Directed Placement (SFD) layout algorithm [94] was used to create the graphs. This figure clearly shows that LRP is a very good source task and that MMN is a great target task. Note that all three different architectures lead to consistent maps overall, evidencing the stability of our findings.

As another way to compare tasks, we also clustered them as done in [21] by representing each task by its row vector in the transferability matrix rescaled through (4). This led to the dendrograms depicted in Figure 7(a). By comparing tasks through their row vectors, we are trying to see whether they transfer in a similar fashion to other tasks. Most tasks cluster together, with N400-N2pc and N170-ERN being the closest pairs. We also see that LRP stands outside of the cluster, as it is a particularly good source task.

Evaluating the structure of cognitive tasks with transfer learning




Figure 5: Significant pairwise standardized mean performance difference between decoding models from Figure 4. Statistical significance was computed using a corrected permutation test, with a cutoff p-value of 0.05.




Figure 6: Transfer balanced accuracy for each pair of sources and target cognitive tasks in the ERP CORE dataset. Cell values correspond to average performance and standard-deviation across 5-fold cross-validation. Diagonal values correspond to standard decoding balanced accuracies, without transfer (*c.f.* subsection 5.1).

5.2.2 M³CV

Moving on to M³CV, we see in Figure 8 that we also obtain very asymmetrical transferability matrices for this dataset. The best example of this asymmetry are the tightly related SSS and TSS tasks, which share the same class labels: visual, auditory or somatosensory stimuli. Surprisingly, while SSS transfers relatively well to TSS, the opposite is not true, showing that steady-state representations have something specific which is not captured when pre-training on transient state data.

Another striking observation is that the EEGInception matrix differs from ShallowNet and EEGNet ones. This is probably due to the lower pre-training performance of this model seen on Figure 4(b), evidencing that it learned weaker representations. Interestingly, we see on both ERP CORE and M³CV results that transfer performances across models (Figures 6 and 8) rank consistently with pre-training performances (Figure 4).

From a paradigms perspective, we can see that no transfer accuracies exceed the reference diagonal values in this dataset. Nonetheless, RS appears as a great target task, benefiting from representations learned in all source tasks. To a smaller extent, TSS also exhibits good performances for all source tasks when using the ShallowNet architecture. We hypothesize that this is only visible for this architecture because it is the only one capable of extracting the information shared by all 4 paradigms, given its superior decoding performance for this dataset (*c.f.* 4(b)). On the contrary, it appears that it is very difficult to transfer onto the SSS paradigm, the extreme case being obtained with EEGInception, for which the chance accuracy is attained for all sources (33% in this case, since it is a 3-class classification problem).

Evaluating the structure of cognitive tasks with transfer learning




Figure 7: Hierarchical clustering dendrogram, showing the resemblance between source tasks in both datasets. Each task is represented by its row in the transferability matrices (Figures 6 and 8) rescaled through Equation 4. Pairwise distances were then computed using the Euclidean distance, before clustering the tasks with the UPGMA algorithm [95].

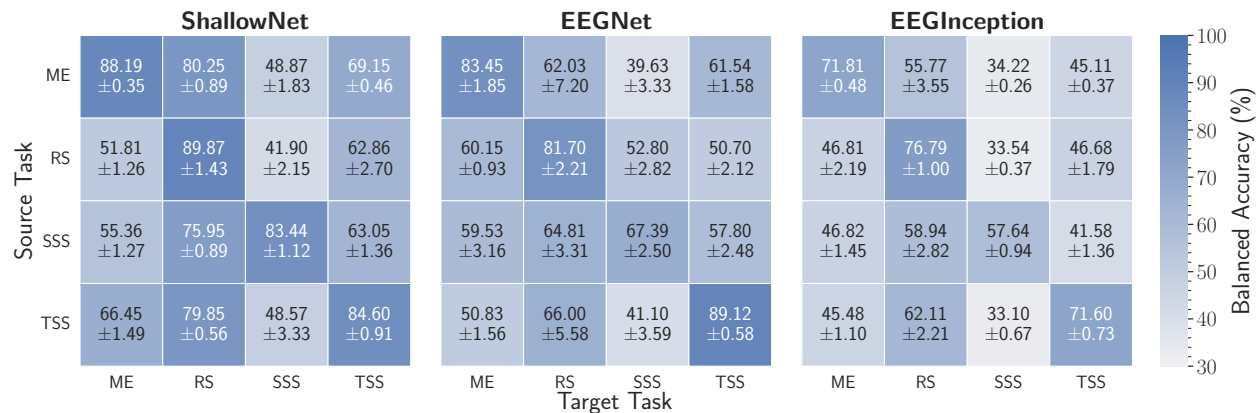


Figure 8: Average performance of the three deep learning models for the transfer learning setting in the M³CV dataset.

As for ERP CORE, we computed the transferability maps for M³CV in Figure 9. This visualization allows to clearly see the consistent connections from TSS, ME and SSS towards RS, regardless of the model employed. The same pattern is visible for TSS with ShallowNet, which receives broad arrows from all neighbouring tasks. We also note a persistent bidirectional connection between TSS and ME, which is quite surprising given that they are very different tasks.

The similarity between the best target tasks, RS and TSS, becomes even more apparent on the clustering dendrogram from Figure 7(b). Indeed, these paradigms are grouped in their own cluster with EEGNet and appear as the closest paradigms with ShallowNet and EEGInception. This figure also confirms that SSS is a very particular task, as it is outside the main cluster for two out of three dendrograms.




Figure 9: Transferability map of the M^3CV dataset. Each node corresponds to a distinct paradigm, while arrows width represent the average transfer performance when using the representations learned from a source task to decode a target task. All tasks transfer well to RS. While SSS transfers to TSS, the opposite is not true. Surprisingly, ME and TSS transfer well to each other in both ways.

6 Discussion and impact

Our analyses have yielded a reliable cognitive map through the use of transfer learning and its successful generalization across various cognitive paradigms. The resulting taxonomy maps, which encapsulate the knowledge of model parameters, open up diverse possibilities for practical applications and further research. In this discussion, we will delve into some potential applications of cognitive mapping from the perspective of machine learning models, emphasizing the benefits and implications they offer.

First and foremost, these cognitive maps can significantly enhance and optimize the training of predictive models. As exemplified in this study, one compelling approach is to leverage a task that exhibits strong transferability as a template for decoding challenging cognitive tasks, as shown for MMN or N170 ERPs. This aspect proves especially crucial in psychophysical experiments, where subject numbers are limited while experimental variables are abundant. These factors seem to be observed in other areas, such as computer vision [96].

Additionally, the cognitive taxonomy maps can play a pivotal role in the design of improved Human Machine Interfaces (HMIs) for Brain-Computer Interface (BCI) applications. Our focus particularly centers around integrated BCIs [8], which extend beyond mere control functions and delve into scenarios involving error detection and negative potential. By incorporating insights gained from cognitive mapping, we can significantly enhance the performance and effectiveness of HMIs in such contexts. For example, the detection of ERN could help to mitigate situations where the BCI system selects unwanted commands, resulting in user frustration.

Furthermore, the clinical applications of cognitive mapping hold immense potential, especially in aiding patients with schizophrenia who encounter difficulties in emotion detection and facial recognition [97]. Leveraging the insights derived from the N170 task [98], we can devise interventions that boost performance and facilitate improved outcomes for individuals grappling with these challenges. This is especially useful in the context of remediation of emotion recognition, where schizophrenic patients have difficulties to recognise faces and to extract relevant cues (like eyebrows or mouth movements) for interpreting social interaction adequately. Objective information indicating that the patient has recognised a face during an interaction, instead of focusing on irrelevant details for instance, could help therapists for rehabilitation purposes.

Moreover, our findings shed light on the functional networks underlying related tasks. Notably, closely related tasks often engage similar functional networks, suggesting the presence of shared evoked components [66, 67]. This implies that activations observed in one task can enhance and refine the performance of another task. While these tasks may exhibit distinct temporal dynamics and spatial orientations, they ultimately serve the same overarching purpose and engage overlapping cortical networks. Unraveling the sources responsible for generating event-related potentials (ERPs) becomes more feasible when considering these perspective-dependent variations in activation patterns. This understanding illuminates the intricate interplay among different components of the cortical network.

One compelling avenue for applying the insights gained from our study lies in the realm of source localization, a classic inverse problem in neuroscience [74]. Cognitive mapping can serve as a potent constraint in reconstructing brain activity, leveraging functional regularization. Source localization methods rely on physical and anatomical constraints to yield plausible solutions. Information regarding shared evoked components could help to define functional constraints on estimated solutions. By leveraging the connections uncovered through cognitive mapping, we can construct robust source localization models with enhanced accuracy and interpretability.

Finally, we can address the issue of BCI illiteracy, which refers to the challenges faced by individuals to effectively operate BCIs [99, 100]. Our transfer learning approach, which could be applied at the subject level, holds promise in mitigating this challenge. BCI illiteracy is mostly observed on few specific tasks for a given user while accuracy on other tasks is correct [8]. Based on our results, a first possible approach is to apply task transfer, training on an effective decoding task for a user to transfer the results to an inefficient task. A second and more exploratory approach is to investigate user transfer, that is transfer the most effective decoding task from one subject and apply it to another with poor decoding abilities. This area presents an exciting and promising avenue for further exploration and improvement.

7 Conclusion

Our study investigates the transferability of deep learning representations through extensive experiments on two EEG datasets. To our knowledge, our work is the first to construct representation transfer maps for EEG decoding. These maps reveal a complex and asymmetric hierarchical relationship between cognitive tasks, enhancing our understanding of brain decoding and neural representations. Our findings also have very practical implications for mitigating data scarcity, demonstrating performance improvements in real world data.

One limitation of our work is that we focus on ERP and BCI datasets. Future work could hence extend our analysis to a broader range of cognitive tasks, or even investigate the transferability across datasets that do not share the same cohort of subjects. Other interesting avenues for investigation could be the study of other types of transferability, for example by carrying out a similar analysis in a self-supervised setting.

Acknowledgments

We would like to thank Dr André Cravo for his valuable feedback on this manuscript. This work was supported by the ANR BrAIN (ANR-20-CHIA-0016), ANR AI-Cog grants (ANR-20-IADJ-0002), DATA IA (project YARN). The work of BA was supported in part by the CAPES under Grant 001 and Data IA mobilité internationale. The work of WHLP was supported by Wellcome Innovations under Grant [WT213038/Z/18/Z].

References

- [1] Jean-Rémi King, Laura Gwilliams, Chris Holdgraf, Jona Sassenhagen, Alexandre Barachant, Denis Engemann, Eric Larson, and Alexandre Gramfort. Encoding and Decoding Framework to Uncover the Algorithms of Cognition. In *The Cognitive Neurosciences*. The MIT Press, 05 2020. ISBN 9780262356176.
- [2] Subhrajit Roy, Isabell Kiral-Kornek, and Stefan Harrer. Chrononet: A deep recurrent neural network for abnormal EEG identification. In *Proceedings of the Artificial Intelligence in Medicine*, pages 47–56, Cham, 2019. Springer International Publishing.
- [3] Giulio Ruffini, David Ibañez, Marta Castellano, Laura Dubreuil-Vall, Aureli Soria-Frisch, Ron Postuma, Jean-François Gagnon, and Jacques Montplaisir. Deep learning with EEG spectrograms in rapid eye movement behavior disorder. *Frontiers in neurology*, 10:806, 2019.
- [4] Lukas A.W. Gemein, Robin T. Schirrmeyer, Patryk Chrabąszcz, Daniel Wilson, Joschka Boedecker, Andreas Schulze-Bonhage, Frank Hutter, and Tonio Ball. Machine-learning-based diagnostics of EEG pathology. *NeuroImage*, 220:117021, 2020. ISSN 1053-8119.
- [5] U Rajendra Acharya, S Vinitha Sree, G Swapna, Roshan Joy Martis, and Jasjit S Suri. Automated EEG analysis of epilepsy: a review. *Knowledge-Based Systems*, 45:147–165, 2013.
- [6] U Rajendra Acharya, Shu Lih Oh, Yuki Hagiwara, Jen Hong Tan, and Hojjat Adeli. Deep convolutional neural network for the automated detection and diagnosis of seizure using EEG signals. *Computers in biology and medicine*, 100:270–278, 2018.
- [7] Khalid Ali I Aboalayon, Miad Faezipour, Wafaa S Almuhammadi, and Saeid Moslehpour. Sleep stage classification using EEG signal analysis: a comprehensive survey and new investigation. *Entropy*, 18(9):272, 2016.

- [8] Fabien Lotte, Laurent Bougrain, Andrzej Cichocki, Maureen Clerc, Marco Congedo, Alain Rakotomamonjy, and Florian Yger. A review of classification algorithms for EEG-based brain–computer interfaces: a 10 year update. *Journal of neural engineering*, 15(3):031005, 2018.
- [9] Jan Van Erp, Fabien Lotte, and Michael Tangermann. Brain-computer interfaces: beyond medical applications. *Computer*, 45(4):26–34, 2012.
- [10] Anatole Lécuyer, Fabien Lotte, Richard B Reilly, Robert Leeb, Michitaka Hirose, and Mel Slater. Brain-computer interfaces, virtual reality, and videogames. *Computer*, 41(10):66–72, 2008.
- [11] Alexander J Casson. Wearable EEG and beyond. *Biomedical engineering letters*, 9(1):53–71, 2019.
- [12] Yannick Roy, Hubert Banville, Isabela Albuquerque, Alexandre Gramfort, Tiago H Falk, and Jocelyn Faubert. Deep learning-based electroencephalography analysis: a systematic review. *Journal of neural engineering*, 16(5): 051001, 2019. doi: 10.1088/1741-2552/ab260c.
- [13] Stanislas Chambon, Mathieu N Galtier, Pierrick J Arnal, Gilles Wainrib, and Alexandre Gramfort. A deep learning architecture for temporal sleep stage classification using multivariate and multimodal time series. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 26(4):758–769, 2018.
- [14] Mathias Perslev, Sune Darkner, Lykke Kempfner, Miki Nikolic, Poul Jørgen Jennum, and Christian Igel. U-Sleep: resilient high-frequency sleep staging. *NPJ digital medicine*, 4(1):72, 2021.
- [15] H. Phan, O. Y. Chen, M. C. Tran, P. Koch, A. Mertins, and M. De Vos. XSleepNet: Multi-view sequential model for automatic sleep staging. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 44(09):5903–5915, 9 2022. ISSN 1939-3539.
- [16] Vernon J Lawhern, Amelia J Solon, Nicholas R Waytowich, Stephen M Gordon, Chou P Hung, and Brent J Lance. EEGnet: a compact convolutional neural network for EEG-based brain–computer interfaces. *Journal of Neural Engineering*, 15(5):056013, 7 2018. doi: 10.1088/1741-2552/aace8c.
- [17] Eduardo Santamaria-Vazquez, Victor Martinez-Cagigal, Fernando Vaquerizo-Villar, and Roberto Hornero. EEG-inception: A novel deep convolutional neural network for assistive ERP-based brain-computer interfaces. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 28(12):2773–2782, 2020. doi: 10.1109/TNSRE.2020.3048106.
- [18] Robin Tibor Schirrmester, Jost Tobias Springenberg, Lukas Dominique Josef Fiederer, Martin Glasstetter, Katharina Eggenberger, Michael Tangermann, Frank Hutter, Wolfram Burgard, and Tonio Ball. Deep learning with convolutional neural networks for EEG decoding and visualization. *Human Brain Mapping*, 8 2017. ISSN 1097-0193. doi: 10.1002/hbm.23730.
- [19] Daniel Wilson, Robin Tibor Schirrmester, Lukas Alexander Wilhelm Gemein, and Tonio Ball. Deep Riemannian networks for EEG decoding, 2023.
- [20] Dongrui Wu, Xue Jiang, and Ruimin Peng. Transfer learning for motor imagery based brain–computer interfaces: A tutorial. *Neural Networks*, 153:235–253, 2022. ISSN 0893-6080. doi: 10.1016/j.neunet.2022.06.008.
- [21] Amir R. Zamir, Alexander Sax, William B. Shen, Leonidas J. Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition - CVPR*, pages 3712–3722. Computer Vision Foundation / IEEE Computer Society, 6 2018. doi: 10.1109/CVPR.2018.00391.
- [22] Emily S Kappenman, Jaelyn L Farrens, Wendy Zhang, Andrew X Stewart, and Steven J Luck. ERP CORE: An open resource for human event-related potential research. *NeuroImage*, 225:117465, 2021. doi: 10.1016/j.neuroimage.2020.117465.
- [23] Gan Huang, Zhenxing Hu, Weize Chen, Shaorong Zhang, Zhen Liang, Linling Li, Li Zhang, and Zhiguo Zhang. M3CV: A multi-subject, multi-session, and multi-task database for EEG-based biometrics challenge. *NeuroImage*, 264:119666, 2022. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2022.119666.
- [24] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2010. doi: 10.1109/TKDE.2009.191.
- [25] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1):43–76, 2021. doi: 10.1109/JPROC.2020.3004555.
- [26] Garrett Wilson and Diane J. Cook. A survey of unsupervised deep domain adaptation. *ACM Trans. Intell. Syst. Technol.*, 11(5), jul 2020. ISSN 2157-6904.

- [27] Antoine de Mathelin, Francois Deheeger, Mathilde Mougeot, and Nicolas Vayatis. *From Theoretical to Practical Transfer Learning: The ADAPT Library*, pages 283–306. Springer International Publishing, Cham, 2023. ISBN 978-3-031-11748-0. doi: 10.1007/978-3-031-11748-0_12.
- [28] Michael T Rosenstein, Zvika Marx, Leslie Pack Kaelbling, and Thomas G Dietterich. To transfer or not to transfer. In *NIPS 2005 workshop on transfer learning*, 2005.
- [29] Zirui Wang, Zihang Dai, Barnabás Póczos, and Jaime Carbonell. Characterizing and avoiding negative transfer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11293–11302, 2019.
- [30] Trevor Standley, Amir Zamir, Dawn Chen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Which tasks should be learned together in multi-task learning? In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 9120–9132. PMLR, 13–18 Jul 2020.
- [31] Arghya Pal and Vineeth N Balasubramanian. Zero-shot task transfer. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2184–2193, 2019. doi: 10.1109/CVPR.2019.00229.
- [32] Kshitij Dwivedi and Gemma Roig. Representation similarity analysis for efficient task taxonomy & transfer learning. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12379–12388, 2019. doi: 10.1109/CVPR.2019.01267.
- [33] Alessandro Achille, Michael Lam, Rahul Tewari, Avinash Ravichandran, Subhransu Maji, Charless Fowlkes, Stefano Soatto, and Pietro Perona. Task2vec: Task embedding for meta-learning. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6429–6438, 2019. doi: 10.1109/ICCV.2019.00653.
- [34] Aria Wang, Michael Tarr, and Leila Wehbe. Neural taskonomy: Inferring the similarity of task-derived representations from brain activity. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [35] Jingcong Li, Fei Wang, Haiyun Huang, Feifei Qi, and Jiahui Pan. A novel semi-supervised meta learning method for subject-transfer brain–computer interface. *Neural Networks*, 163:195–204, 2023. ISSN 0893-6080.
- [36] Fabien Lotte and Cuntai Guan. Learning from other subjects helps reducing brain-computer interface calibration time. In *2010 IEEE International conference on acoustics, speech and signal processing*, pages 614–617. IEEE, 2010.
- [37] Dieter Devlaminck, Bart Wyns, Moritz Grosse-Wentrup, Georges Otte, and Patrick Santens. Multisubject learning for common spatial patterns in motor-imagery BCI. *Computational intelligence and neuroscience*, 2011: 8–8, 2011.
- [38] Vinay Jayaram, Morteza Alamgir, Yasemin Altun, Bernhard Scholkopf, and Moritz Grosse-Wentrup. Transfer learning in brain-computer interfaces. *IEEE Computational Intelligence Magazine*, 11(1):20–31, 2016.
- [39] Jin Feng, Yunde Li, Chengliang Jiang, Yu Liu, Mingxin Li, and Qinghui Hu. Classification of motor imagery electroencephalogram signals by using adaptive cross-subject transfer learning. *Frontiers in Human Neuroscience*, 16, 2022. ISSN 1662-5161. doi: 10.3389/fnhum.2022.1068165.
- [40] Fulin Wei, Xueyuan Xu, Tianyuan Jia, Daoqiang Zhang, and Xia Wu. A multi-source transfer joint matching method for inter-subject motor imagery decoding. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 31:1258–1267, 2023. doi: 10.1109/TNSRE.2023.3243257.
- [41] Yunyuan Gao, Mengting Li, Yun Peng, Feng Fang, and Yingchun Zhang. Double stage transfer learning for brain–computer interfaces. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 31:1128–1136, 2023.
- [42] Yunyuan Gao, Yici Liu, Qingshan She, and Jianhai Zhang. Domain adaptive algorithm based on multi-manifold embedded distributed alignment for brain-computer interfaces. *IEEE Journal of Biomedical and Health Informatics*, 27(1):296–307, 2023. doi: 10.1109/JBHI.2022.3218453.
- [43] Qingshan She, Yinhao Cai, Shengzhi Du, and Yun Chen. Multi-source manifold feature transfer learning with domain selection for brain-computer interfaces. *Neurocomputing*, 514:313–327, 2022. ISSN 0925-2312. doi: <https://doi.org/10.1016/j.neucom.2022.09.124>.
- [44] Emmanuel K Kalunga, Sylvain Chevallier, and Quentin Barthélemy. Transfer learning for SSVEP-based BCI using Riemannian similarities between users. In *2018 26th European Signal Processing Conference (EUSIPCO)*, pages 1685–1689. IEEE, 2018.
- [45] Qingshan She, Tie Chen, Feng Fang, Jianhai Zhang, Yunyuan Gao, and Yingchun Zhang. Improved domain adaptation network based on Wasserstein distance for motor imagery EEG classification. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 31:1137–1148, 2023. doi: 10.1109/TNSRE.2023.3241846.

- [46] Andrea Apicella, Francesco Isgrò, Andrea Pollastro, and Roberto Prevede. On the effects of data normalisation for domain adaptation on EEG data. *ArXiv*, abs/2210.01081, 2022.
- [47] Richard Antonello, Javier S Turek, Vy Vo, and Alexander Huth. Low-dimensional structure in the space of language representations is reflected in brain responses. In *Advances in Neural Information Processing Systems*, volume 34, pages 8332–8344. Curran Associates, Inc., 2021.
- [48] Subba Reddy Oota, Jashn Arora, Veeral Agarwal, Mounika Marreddy, Manish Gupta, and Bapi Surampudi. Neural language taskonomy which NLP tasks are the most predictive of fMRI brain activity? In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics Human Language Technologies*, pages 3220–3237, Seattle, United States, July 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.naacl-main.235.
- [49] Youzhi Qu, Xinyao Jian, Wenxin Che, Penghui Du, Kai Fu, and Quanying Liu. Transfer learning to decode brain states reflecting the relationship between cognitive tasks. In *Human Brain and Artificial Intelligence*, pages 110–122, Singapore, 2023. Springer Nature Singapore. ISBN 978-981-19-8222-4.
- [50] Cédric Rommel, Joseph Paillard, Thomas Moreau, and Alexandre Gramfort. Data augmentation for learning predictive models on EEG: a systematic comparison. *Journal of Neural Engineering*, 19(6):066020, 11 2022. doi: 10.1088/1741-2552/aca220.
- [51] Hubert Banville, Omar Chehab, Aapo Hyvärinen, Denis-Alexander Engemann, and Alexandre Gramfort. Uncovering the structure of clinical EEG signals with self-supervised learning. *Journal of Neural Engineering*, 18(4): 046020, 3 2021. doi: 10.1088/1741-2552/abca18.
- [52] Hubert Banville, Isabela Albuquerque, Aapo Hyvärinen, Graeme Moffat, Denis-Alexander Engemann, and Alexandre Gramfort. Self-supervised representation learning from electroencephalography signals. In *2019 IEEE 29th International Workshop on Machine Learning for Signal Processing (MLSP)*, pages 1–6, 2019. doi: 10.1109/MLSP.2019.8918693.
- [53] Mehdi Azabou, Max Dabagia, Ran Liu, Chi-Heng Lin, Keith B Hengen, WashU-St Louis, and Eva L Dyer. Using self-supervision and augmentations to build insights into neural coding. In *Proceedings of Workshop on Self-Supervised Learning - Theory and Practice on NeurIPS*, Proceedings of Machine Learning Research, pages 1–7. NeurIPS, 12 2021.
- [54] Mostafa Neo Mohsenvand, Mohammad Rasool Izadi, and Pattie Maes. Contrastive representation learning for electroencephalogram classification. In *Proceedings of the Machine Learning for Health NeurIPS Workshop*, volume 136 of *Proceedings of Machine Learning Research*, pages 238–253. PMLR, 11 Dec 2020.
- [55] Guillaume Alain and Yoshua Bengio. Understanding intermediate layers using linear classifier probes. *arXiv preprint arXiv:1610.01644*, 2016.
- [56] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9640–9649, 2021.
- [57] Ananya Kumar, Aditi Raghunathan, Robbie Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pretrained features and underperform out-of-distribution. *arXiv preprint arXiv:2202.10054*, 2022.
- [58] Jonathan R Wolpaw, Niels Birbaumer, Dennis J McFarland, Gert Pfurtscheller, and Theresa M Vaughan. Brain–computer interfaces for communication and control. *Clinical Neurophysiology*, 113(6):767–791, 2002. ISSN 1388-2457. doi: 10.1016/S1388-2457(02)00057-3.
- [59] Brendan Z. Allison and Christa Neuper. *Could Anyone Use a BCI?*, pages 35–54. Springer London, London, 2010. ISBN 978-1-84996-272-8.
- [60] Wei-Long Zheng, Wei Liu, Yifei Lu, Bao-Liang Lu, and Andrzej Cichocki. Emotionmeter: A multimodal framework for recognizing human emotions. *IEEE Transactions on Cybernetics*, 49(3):1110–1122, 2019. doi: 10.1109/TCYB.2018.2797176.
- [61] Pablo Arnau-González, Stamos Katsigiannis, Miguel Arevalillo-Herráez, and Naeem Ramzan. Bed: A new data set for EEG-based biometrics. *IEEE Internet of Things Journal*, 8(15):12219–12230, 2021. doi: 10.1109/JIOT.2021.3061727.
- [62] Hohyun Cho, Minkyu Ahn, Sangtae Ahn, Moonyoung Kwon, and Sung Chan Jun. EEG datasets for motor imagery brain–computer interface. *GigaScience*, 6(7), 05 2017. ISSN 2047-217X. doi: 10.1093/gigascience/gix034. gix034.
- [63] Ary L. Goldberger, Luis A. N. Amaral, Leon Glass, Jeffrey M. Hausdorff, Plamen Ch. Ivanov, Roger G. Mark, Joseph E. Mietus, George B. Moody, Chung-Kang Peng, and H. Eugene Stanley. Physiobank, physiotoolkit, and physionet. *Circulation*, 101(23):e215–e220, 2000. doi: 10.1161/01.CIR.101.23.e215.

- [64] Guanghui Zhang and Steven J. Luck. Variations in ERP data quality across paradigms, participants, and scoring procedures. *Psychophysiology*, n/a(n/a):e14264, 2023. doi: 10.1111/psyp.14264.
- [65] William J Gehring, Yanni Liu, Joseph M Orr, and Joshua Carp. The error-related negativity (ERN/Ne). *Oxford handbook of event-related potential components*, pages 231–291, 2012.
- [66] Steven J. Luck. *An introduction to the event-related potential technique*. The MIT Press, 2014.
- [67] Steven J. Luck. *Applied Event-Related Potential Data Analysis*, volume 1. LibreTexts, 2022.
- [68] Marta Kutas and Kara D Federmeier. Thirty years and counting: finding meaning in the N400 component of the event-related brain potential (ERP). *Annual review of psychology*, 62:621–647, 2011.
- [69] Steven J Luck et al. Electrophysiological correlates of the focusing of attention within complex visual scenes: N2pc and related ERP components. *The Oxford handbook of event-related potential components*, pages 329–360, 2012.
- [70] Bruno Rossion and Stéphanie Caharel. ERP evidence for the speed of face categorization in the human brain: Disentangling the contribution of low-level visual cues from face perception. *Vision research*, 51(12):1297–1311, 2011.
- [71] Walter Ritter and Herbert G. Vaughan. Averaged Evoked Responses in Vigilance and Discrimination: A Reassessment. *Science*, 164(3877):326–328, 1969.
- [72] Risto Näätänen, Satu Pakarinen, Teemu Rinne, and Rika Takegata. The mismatch negativity (MMN): towards the optimal paradigm. *Clinical neurophysiology*, 115(1):140–144, 2004.
- [73] Pierre Ablin, Jean-François Cardoso, and Alexandre Gramfort. Faster independent component analysis by preconditioning with hessian approximations. *IEEE Transactions on Signal Processing*, 66(15):4040–4049, 2018.
- [74] Alexandre Gramfort, Martin Luessi, Eric Larson, Denis Engemann, Daniel Strohmeier, Christian Brodbeck, Roman Goj, Mainak Jas, Teon Brooks, Lauri Parkkonen, and Matti Hämäläinen. MEG and EEG data analysis with MNE-Python. *Frontiers in Neuroscience*, 7:267, 2013. ISSN 1662-453X. doi: 10.3389/fnins.2013.00267.
- [75] Kai Keng Ang, Zheng Yang Chin, Haihong Zhang, and Cuntai Guan. Filter bank common spatial pattern (fbcsp) in brain-computer interface. In *2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*, pages 2390–2397, 2008. doi: 10.1109/IJCNN.2008.4634130.
- [76] Kostas M Tsiouris, Vasileios C Pezoulas, Michalis Zervakis, Spiros Konitsiotis, Dimitrios D Koutsouris, and Dimitrios I Fotiadis. A long short-term memory deep learning network for the prediction of epileptic seizures using EEG signals. *Computers in biology and medicine*, 99:24–37, 2018.
- [77] Davide Borra, Silvia Fantozzi, and Elisa Magosso. Interpretable and lightweight convolutional neural network for EEG decoding: Application to movement execution and imagination. *Neural Networks*, 129:55–74, 2020. ISSN 0893-6080. doi: 10.1016/j.neunet.2020.05.032.
- [78] Alberto Zancanaro, Giulia Cisotto, João Ruivo Paulo, Gabriel Pires, and Urbano J Nunes. CNN-based approaches for cross-subject classification in motor imagery: From the state-of-the-art to dynamicnet. In *Proceedings of the IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology - CIBCB*, pages 1–7. IEEE, 2021.
- [79] Huan Cai, Min Xia, Li Nie, Yihan Wu, and Yangsong Zhang. Deep learning models with time delay embedding for EEG-based attentive state classification. In *Proceedings of the Neural Information Processing*, pages 307–314, Cham, 2021. Springer International Publishing.
- [80] Stylianos Bakas, Siegfried Ludwig, Konstantinos Barmpas, Mehdi Bahri, Yannis Panagakis, Nikolaos Laskaris, Dimitrios A Adamos, and Stefanos Zafeiriou. Team cogitat at neurips 2021: Benchmarks for EEG transfer learning competition. *arXiv preprint arXiv:2202.03267*, 2022.
- [81] Xiaoxi Wei, A. Aldo Faisal, Moritz Grosse-Wentrup, Alexandre Gramfort, Sylvain Chevallier, Vinay Jayaram, Camille Jeunet, Stylianos Bakas, Siegfried Ludwig, Konstantinos Barmpas, Mehdi Bahri, Yannis Panagakis, Nikolaos Laskaris, Dimitrios A. Adamos, Stefanos Zafeiriou, William C. Duong, Stephen M. Gordon, Vernon J. Lawhern, Maciej Śliwowski, Vincent Rouanne, and Piotr Tempczyk. 2021 beetl competition: Advancing transfer learning for subject independence and heterogenous EEG data sets. In *Proceedings of the NeurIPS 2021 Competitions and Demonstrations Track*, volume 176 of *Proceedings of Machine Learning Research*, pages 205–219. PMLR, 12 2022.
- [82] Sylvain Chevallier, Guillaume Bao, Mayssa Hammami, Fabienne Marlats, Louis Mayaud, Djillali Annane, Frédéric Lofaso, and Eric Azabou. Brain-machine interface for mechanical ventilation using respiratory-related evoked potential. In *Artificial Neural Networks and Machine Learning—ICANN 2018: 27th International*

- Conference on Artificial Neural Networks, Rhodes, Greece, October 4-7, 2018, Proceedings, Part III 27*, pages 662–671. Springer, 2018.
- [83] Alexandre Barachant, Stéphane Bonnet, Marco Congedo, and Christian Jutten. Riemannian geometry applied to BCI classification. *LVA/ICA*, 10:629–636, 2010.
- [84] Alexandre Barachant and Marco Congedo. A plug&play P300 BCI using information geometry. *arXiv preprint arXiv:1409.0107*, 2014.
- [85] Salim Khazem, Sylvain Chevallier, Quentin Barthélemy, Karim Haroun, and Camille Noûs. Minimizing subject-dependent calibration for BCI with Riemannian transfer learning. In *2021 10th International IEEE/EMBS Conference on Neural Engineering (NER)*, pages 523–526. IEEE, 2021.
- [86] Vinay Jayaram and Alexandre Barachant. MOABB: trustworthy algorithm benchmarking for BCIs. *Journal of neural engineering*, 15(6):066011, 2018.
- [87] Alexandre Barachant, Quentin Barthélemy, Jean-Rémi King, Alexandre Gramfort, Sylvain Chevallier, Pedro L. C. Rodrigues, Emanuele Olivetti, Vladislav Goncharenko, Gabriel Wagner vom Berg, Ghiles Reguig, Arthur Lebourrier, Erik Bjäreholt, Maria Sayu Yamamoto, Pierre Clisson, and Marie-Constance Corsi. pyRiemann/pyRiemann: v0.3, July 2022.
- [88] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics - AISTATS*, volume 9 of *JMLR Proceedings*, pages 249–256. JMLR.org, 2010.
- [89] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *Proceedings of the International Conference on Learning Representations - ICLR*. OpenReview.net, 2019.
- [90] Lutz Prechelt. Early stopping - but when? In *Neural Networks: Tricks of the Trade - Second Edition*, volume 7700 of *Lecture Notes in Computer Science*, pages 53–67. Springer, 2012. doi: 10.1007/978-3-642-35289-8_5.
- [91] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [92] Hannes Bonasch and Benedikt V Ehinger. Decoding accuracies as well as ERP amplitudes do not show between-task correlations. *bioRxiv*, pages 2023–05, 2023.
- [93] Timothy Trammel, Natalia Khodayari, Steven J Luck, Matthew J Traxler, and Tamara Y Swaab. Decoding semantic relatedness and prediction from EEG: A classification method comparison. *NeuroImage*, page 120268, 2023.
- [94] Y. F. Hu. Efficient and high quality force-directed graph drawing. *The Mathematica Journal*, 10:37–71, 2005.
- [95] Robert R. Sokal and Charles Duncan Michener. A statistical method for evaluating systematic relationships. *University of Kansas science bulletin*, 38:1409–1438, 1958.
- [96] Thomas Mensink, Jasper Uijlings, Alina Kuznetsova, Michael Gygli, and Vittorio Ferrari. Factors of Influence for Transfer Learning Across Diverse Appearance Domains and Task Types. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):9298–9314, 2022. doi: 10.1109/TPAMI.2021.3129870.
- [97] Tamara Watson. Implications of holistic face processing in autism and schizophrenia. *Frontiers in Psychology*, 4, 2013. ISSN 1664-1078. doi: 10.3389/fpsyg.2013.00414.
- [98] Yu-Fang Yang, Eric Brunet-Gouet, Mariana Burca, Emmanuel K. Kalunga, and Michel-Ange Amorim. Brain Processes While Struggling With Evidence Accumulation During Facial Emotion Recognition: An ERP Study. *Frontiers in Human Neuroscience*, 14, 2020. ISSN 1662-5161. doi: 10.3389/fnhum.2020.00340.
- [99] Min-Ho Lee, O-Yeon Kwon, Yong-Jeong Kim, Hong-Kyung Kim, Young-Eun Lee, John Williamson, Siamac Fazli, and Seong-Whan Lee. EEG dataset and OpenBMI toolbox for three BCI paradigms: an investigation into BCI illiteracy. *GigaScience*, 8(5), 01 2019. ISSN 2047-217X. doi: 10.1093/gigascience/giz002. giz002.
- [100] Biao Sun, Zexu Wu, Yong Hu, and Ting Li. Golden subject is everyone: A subject transfer neural network for motor imagery-based brain computer interfaces. *Neural Networks*, 151:111–120, 2022. ISSN 0893-6080.