Journal of Neural Engineering

PAPER • OPEN ACCESS

A meta-learning BCI for estimating decision confidence

To cite this article: Christoph Tremmel *et al* 2022 *J. Neural Eng.* **19** 046009

View the article online for updates and enhancements.

You may also like

- Imbalanced fault identification via embedding-augmented Gaussian prototype network with meta-learning perspective
Rujie Hou, Zhenyi Chen, Jinglong Chen et al.
- Short-Term Wind Speed Forecasting Using Meta Learning-based Elman Neural Network
M. Madhiarasan
- Few-shot bearing fault diagnosis based on meta-learning with discriminant space optimization
Dengming Zhang, Kai Zheng, Bai Yin et al.

This content was downloaded from IP address 193.60.238.99 on 04/08/2022 at 16:43

IOP Publishing J. Neural Eng. 19 (2022) 046009 https://doi.org/10.1088/1741-2552/ac7ba8

Journal of Neural Engineering

CrossMark

OPEN ACCESS

RECEIVED
11 February 2022

REVISED
20 June 2022

ACCEPTED FOR PUBLICATION
23 June 2022

PUBLISHED
11 July 2022

Original Content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence.

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

PAPER

A meta-learning BCI for estimating decision confidence

Christoph Tremmel[1], **Jacobo Fernandez-Vargas**[1], **Dimitris Stamos**[2], **Caterina Cinel**[1], **Massimiliano Pontil**[2,3], **Luca Citi**[1] **and Riccardo Poli**[1,*]

[1] Brain-Computer Interfaces and Neural Engineering Laboratory, School of Computer Science and Electronic Engineering, University of Essex, Colchester, Essex, United Kingdom
[2] Department of Computer Science, University College of London, London, United Kingdom
[3] Computational Statistics and Machine Learning, Italian Institute of Technoloy, Genoa, Italy
[*] Author to whom any correspondence should be addressed.

E-mail: rpoli@essex.ac.uk

Keywords: brain–computer interfaces, EEG, meta learning, decision confidence prediction, decision making

Abstract

Objective. We investigated whether a recently introduced transfer-learning technique based on meta-learning could improve the performance of brain–computer interfaces (BCIs) for decision-confidence prediction with respect to more traditional machine learning methods. *Approach.* We adapted the meta-learning by biased regularisation algorithm to the problem of predicting decision confidence from electroencephalography (EEG) and electro-oculogram (EOG) data on a decision-by-decision basis in a difficult target discrimination task based on video feeds. The method exploits previous participants' data to produce a prediction algorithm that is then quickly tuned to new participants. We compared it with with the traditional single-subject training almost universally adopted in BCIs, a state-of-the-art transfer learning technique called domain adversarial neural networks, a transfer-learning adaptation of a zero-training method we used recently for a similar task, and with a simple baseline algorithm. *Main results.* The meta-learning approach was significantly better than other approaches in most conditions, and much better in situations where limited data from a new participant are available for training/tuning. Meta-learning by biased regularisation allowed our BCI to seamlessly integrate information from past participants with data from a specific user to produce high-performance predictors. Its robustness in the presence of small training sets is a real-plus in BCI applications, as new users need to train the BCI for a much shorter period. *Significance.* Due to the variability and noise of EEG/EOG data, BCIs need to be normally trained with data from a specific participant. This work shows that even better performance can be obtained using our version of meta-learning by biased regularisation.

1. Introduction

1.1. Decision making

Accurate decision making is an essential aspect in human lives, especially in complex environments such as police force, fire brigade, air traffic controller etc where wrong decisions can endanger the well-being of others. Making a decision is a deliberative process in which the brain weighs contextual cues and pre-existing knowledge that results in the commitment to a categorical choice to achieve a particular goal [1–4]. During the decision-making process, the human brain accumulates and integrates all sources of priors, evidence, and values to produce a choice. From this information the decision maker derives also an evaluation of the likelihood of the decision being correct, which is known as the *confidence* [5–8]. Even though confidence is a subjective measure, studies have revealed that self-reports of confidence are generally good predictors of task performance and that both confidence and performance also correlate with secondary behavioural and physiological measures, such as the response time (RT) [5, 8–11]. Confidence estimation processes are also triggered when judging the performance of decision tasks performed by other people or even AIs [12].

Since real-life decisions are often not immediately followed up by objective feedback on their

© 2022 The Author(s). Published by IOP Publishing Ltd

quality, decision confidence is one option to assess the outcome, at least until more reliable information becomes available.

1.2. Neural correlates of decision making and brain–computer interfaces

Neuro-imaging techniques like electroencephalography (EEG) have identified multiple brain areas that are relevant for the decision making process [13–15]. In particular, activity in the pre-frontal and parietal brain regions has been shown to correlate with the reported confidence [7, 16–19]. Such neural correlates, however, can only be measured on collections of tens or hundreds of trials/decisions, so they do not allow a decision-by-decision analysis.

The field of brain–computer interfaces (BCIs) has accumulated significant experience in detecting, predicting and classifying brain responses on a trial-by-trial basis. BCIs establish a communication pathway between the human central nervous system and electronic devices, such as computers [20]. In the more widespread forms of BCI, which are called *active BCIs*, normally the purpose is to classify the brain activity recorded in a trial into a command for the computer. However, limitations in sensing technologies and variability in brain signals result in active BCIs being able to reliably recognise only a small number of brain-activity patterns and to issue the corresponding commands at a slow pace [21–23]. Other forms of BCI, namely *passive BCIs*, do not directly map brain activity into commands, instead focusing on assessing the mental state of the user for other practical purposes. Such BCIs have been used to predict mental workload, valence, and arousal [24–26] but also the probability of decisions being correct which was then used to aid group decision making [27–33].

The prediction of decision confidence from EEG has also recently gained popularity. However, such a prediction is often not very precise as it is mostly treated as a binary classification problem with the two classes being 'confident' and 'non-confident' [17, 34, 35], which is a severe limitation, as confidence admits of degree. Recently, we found that the event-related potentials (ERPs) produced during the decision also vary in shape and magnitude as a function of the reported confidence level and such variations allow for a trial-by-trial confidence prediction with four different confidence levels [36], which is a significant improvement over the state of the art, but still not precise enough.

A challenge for BCIs is the variability of EEG signals across participants, sessions and experiments. Given that modern BCIs are virtually all based on machine learning algorithms, this results in a misalignment between the data used to train algorithms and new data, e.g. from a new participant, which is referred to as a *domain shift* [37]. The issue increases the time a participant needs to wait before using the BCI (as training or retraining of the BCI is required before use) and significantly limits BCI applicability in many domains [38].

1.3. Transfer learning

In order to tackle the domain shift problem, *transfer learning* techniques have been introduced, where algorithms are trained with data from different subjects to reduce the training needs [39, 40]. Essentially transfer learning techniques learn which features are transferable between different domains to avoid degradation of algorithm performance. Transfer learning techniques have been successfully applied in classical BCIs paradigms such as ERP detection [41–43], motor imagery [42, 44, 45] and steady-state visual evoked potentials [42, 46]. Furthermore they also have been applied to a more recent topic, the detection and recognition of participants' emotional affective states [47–51]. In our own research, we used a zero-training transfer-learning approach that could successfully predict decision-making confidence for each subject individually, across multiple subjects, and even across different experiments [36]. However, while transfer learning techniques generally perform well, they still cannot achieve the performance of BCIs trained for each individual participant.

A good compromise between the generalizability of transfer learning and fast adaptation to individual participants could be provided by *meta-learning techniques*. From a theoretical perspective, the problem attacked by such techniques is the following. The use of data mostly from other participants, as ordinary transfer-learning does, introduces a possibly large bias (i.e. the BCI learns to decode a hypothetical average participant rather than specialising to the target participant). On the other hand, using only data from target participants is likely to produce a large variance (the degradation in performance due to training on insufficient amounts of data), unless a large enough amount of data is collected from the new participant, which is often undesirable because it is time consuming, tiring, etc. Meta-learning tries to achieve the best of both worlds by using as much past data as possible while also specialising to the target participant. In other words, meta-learning algorithms learn certain properties of a general problem to allow fast adaptation to new problems or environments [52]. This flexibility is achieved by introducing multiple learning episodes covering additional distributions of different tasks or participants [53].

Meta-learning has recently been tested in the field of BCIs for motor imagery tasks [54, 55] and drowsiness prediction [56], but not in decision making tasks.

2. Related work and contributions

2.1. Meta-learning by biased regularisation

The meta-learning technique that was used in this work belongs to the family of inner stochastic gradient descent algorithms with biased regularization

[57, 58] and was introduced in [59]. Its online algorithm has the advantage that it processes data sequentially and does not have to keep previous training sets in memory compared to other online techniques [60, 61]. The resulting low memory footprint and complexity make this method particularly suitable for BCI applications. In this section we provide a high-level description of the technique.

In the standard supervised learning setting, the goal is to learn a functional relation between an input space $\mathcal{X}$ and an output space $\mathcal{Y}$ from a finite number of training examples. Given a loss function $\ell: \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$ measuring prediction errors and given a distribution μ on the joint data space $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$, the goal is to find a function $f: \mathcal{X} \to \mathcal{Y}$ that minimizes the expected risk, i.e. the expectation of ℓ over the unknown distribution μ. When moving to a meta-learning setting, the addition is that we first sample the distribution μ of each dataset from a meta-distribution ρ, which is called *environment*.

In our experiment, participants were invited to attend three experimental sessions held on different days where they had to carry out a difficult decision task and report their confidence after each decision. The three sessions allowed us to capture session-to-session variability both the physiological signals (EEG and electro-oculogram or EOG) and subject mental state. The objective was to obtain BCIs capable of predicting the participants confidence from the physiological data.

Fitting the general framework of meta-learning described above to the confidence-prediction problem to be tackled by the BCI, we assume that for each day t we observe a dataset (data sequence) $Z_t \equiv (z_{t,1}, \ldots, z_{t,n})$, whose elements are sampled independently from the probability distribution μ_t for day t. In other words we consider each separate session with a participant as a new task to be learned by the meta-learning algorithm. We assume that the probability distributions $\mu_1, \ldots, \mu_T$ are sampled independently from a 'meta-distribution' ρ, that is a probability measure over probability measures on $\mathcal{Z}$. The meta-distribution ρ can then be seen as the general distribution of all possible subjects and days. The main assumption we make in the current setting is that the multiple datasets $Z_1, \ldots, Z_T$ sampled via the above process are captured by linear regression models. More importantly, we assume that the regression vectors are perturbations of an underlying regression model (common mean)[4].

Let w_t be the regression vector for task t and θ the common mean[5]. We now outline the meta-learning procedure to estimate the common mean. The advantage of this is a faster adaptation to new datasets (tasks/sessions) in the environment. Given such a generic dataset $Z = (x_i, y_i)_{i=1}^n$ the corresponding regression vectors are estimated by biased regularization (more precisely biased ridge regression):

$$\boldsymbol{L}_Z(\theta) = \min_{w \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^{n} \ell\big(\langle x_i, w \rangle, y_i\big) + \frac{\lambda}{2} \|w - \theta\|_2^2, \quad (1)$$

where $\lambda > 0$ is the regularization parameter and $\ell: \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ is a loss function, e.g. the squared loss, $\ell(y, y') = (y - y')^2$. We will also refer to the problem in equation (1) as *fine-tuning*, as the meta-parameter θ can be interpreted as a starting model which is fine-tuned to the task at hand.

In turn, the common mean θ is updated as the tasks are observed sequentially, with the ultimate goal of solving the problem

$$\min_{\theta \in \mathbb{R}^d} \mathbb{E}_{\mu \sim \rho} \mathbb{E}_{Z \sim \mu^n} \boldsymbol{L}_Z(\theta). \quad (2)$$

The overall process consists of a bi-level optimization problem in which equation (1) is the inner (or lower level) problem and equation (2) the outer (or upper level) problem. In the inner problem, the regularization term is used to encourage the corresponding test weight vector w to be close to θ. This reduces the number of training points needed to learn a good model. The degree of the regularization λ determines how much the regression algorithm generalizes (i.e. transfers knowledge) across tasks vs adapting to a participant's specific dataset. A higher λ means that the algorithm regularizes towards the meta-parameter relying more on source data to train the regression coefficients. This leads to more transfer from source data sets. A lower λ implies less regularization and so the algorithm relies more on data from the target, resulting in more adaptation.

The solution of the problem is implemented through two stages: *meta-training* where the goal is to solve problem equation (2) with respect to the mean meta-parameter θ, and *fine-tuning*, which tunes the optimal θ thus found in order to adapt it to a new task. More details on the specific adaptations we made will be provided in section 3.5.

2.2. Contributions of this study

In this study we investigate whether a suitably tailored adaptation of a recently introduced transfer-learning technique based on meta-learning—meta-learning by biased regularisation [59]—could improve the

[4] Meta-learning works also when such assumptions do not hold. Indeed, previous theoretical work [59] has shown that when the assumptions are valid, then meta-learning brings a substantial advantage over learning in isolation, when they are not, meta-learning performs as independent learning.

[5] In order to improve readability we assume that all datasets consist of the same number of points n.

performance of BCIs for decision-confidence prediction with respect to other learning methods. We chose a linear algorithm for this problem since the amplitude of the ERPs associated with varying levels of confidence is approximately linearly related to confidence [16, 36, 62]. The algorithm has two phases: first, it iteratively trains a BCI with data from previous participants to learn domain-invariant features and then it uses data from a new participant to quickly fine-tune the BCI.

We compared the performance and behaviour of this meta-learning algorithm with those of: (a) individually-trained BCIs, (b) a simple form of transfer learning obtained by adapting the zero-training approach reported in our previous work [36], and (c) a simple baseline algorithm. We found that our meta-learning approach provides better performance, while requiring smaller training sets, than individually-trained BCIs and the other approaches tested.

Additional contributions of the study are the increased precision of the confidence reporting system, allowing for ten different levels, and the continuous estimation of confidence, setting our setup apart from the state of the art in previous work.

3. Materials and methods

3.1. Participants

Prior to taking part, participants were informed about the experiment and gave their written consent. The research received ethical approval by the UK Ministry of Defence and the University of Essex in June 2017 (Application Number: 832/MoDREC/17). All experiments were performed in accordance with the relevant guidelines and regulations.

All subjects had normal or corrected to normal vision. The experiment consisted of three sessions for each participant that needed to be completed in separated days within one week. Participants were required to schedule the experiments so that the starting time of the experiment across the three sessions did not differ by more than 3 h. They were also required to reach a minimum of 60% in decision accuracy in each session to ensure that the task was taken seriously. Sessions lasted just under 2.5 h excluding preparation time. Participants were paid 15 GPB for each of the first two days and a further 20 GBP after the completion on the third day.

Twenty participants were enrolled in the study, of which 12 attended all 3 days. Two of these had to be discarded because one participant changed their confidence reporting strategy during the experiment while the other did not understand the experiment correctly while still maintaining the required performance. So, in total ten participants (five males, mean age 26.6, SD 5.29) were included in this study.

3.2. Physiological signal acquisition

Participants sat comfortably during the whole experiment approximately 80 cm in front of an LCD screen while being connected to 64-electrode EEG, electrocardiogram (ECG), galvanic skin resistance (GSR), respiration belt, and four-electrode EOG. All signals were sampled at 2048 Hz and synchronized using a Bio Semi Active Two system. In this study, only EEG and EOG signals were considered since the changes in ECG, GSR and respiratory signals proved to be too slow to capture meaningful correlates in our fast-paced recognition task[6].

The EEG recording system used Biosemi's pin-type wet active electrodes. The 64 electrodes were positioned according to the international 10–20 system. Additionally, two electrodes were placed on each of the participants' earlobes to serve as references. Two EOG electrodes were placed above and below the right eye and two near the other canthus of each eye. These provided the vertical and horizontal oculogram, respectively.

3.3. Decision task

Participants performed a visual object recognition task (more details below). They reported decisions and confidence values using a mouse. The mouse was located within the reach of their preferred hand. The experiment started by showing participants how to correctly perform the task and allowing them to familiarise with the two different stimuli that they would be presented with. Then a practice block with 24 trials/decisions was started where feedback was provided right after each decision.

After that, the main experiment started. This consisted of 36 blocks of 24 trials each, for a total of 864 trials. After each block, feedback on the accuracy (the fraction of correct decisions) in that block was provided and the participants could rest for a short time. The timeline of each trial is shown in figure 1. Each trial included the presentation of a sequence of computer generated images that represented a corridor with doors at both sides. The images were updated at a rate of ten frames per second to create the feeling of movement along the corridor. At certain times, three characters in uniform would cross the corridor coming from different doors at either side of the corridor. The characters walked at the same speed, but there was a delay of approximately 70 ms between the movement onset of each character to reduce occlusions. The time between the first character starting to cross the corridor and the last character completing the crossing was 1 s. Two of the characters wore a cap, while the third wore either a

[6] EOG has been recently applied to detect mental state like vigilance [63] or drowsiness [64] and has increasingly been used in combination with EEG to improve the robustness of BCIs [65, 66].

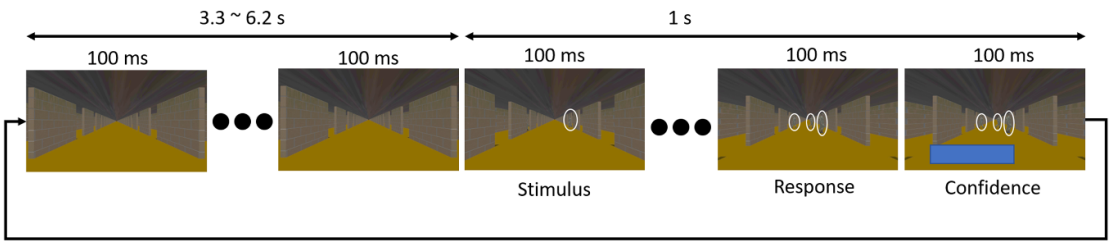

Figure 1. Timeline of the experiment used in this study. Before each trial there is a variable waiting time. Stimuli are presented for one second. When a decision is made by the participant by clicking the right or left mouse button, a bar representing the confidence is displayed. Confidence is reported by using the mouse wheel.

helmet or a cap, with a 50/50 chance. Participants had to decide, as quickly as possible, whether any of the characters they had seen wore a helmet (left click) or not (right click). Immediately after responding, they had to report the confidence on the decision using the scroll wheel of the mouse. Confidence was represented as a blue slider bar at the bottom of the screen. Participants were told to use the lowest confidence level (0%) to indicate that they knew that they had made a mistake, or that they had miss-clicked. Corresponding trials were removed from the analysis. While making decisions and reporting confidences, the simulated movement along the corridors continued (not stopping even if a participant did not provide any response).

3.4. Data pre-processing and feature extraction

The EEG channels were referenced to the average of the ear electrodes, while horizontal and vertical EOG were acquired by subtracting the corresponding electrodes above and below or beside the eyes. Afterwards, EEG and EOG were band-pass filtered between 0.05 and 64 Hz using a 4th order Butterworth filter and down-sampled to 128 Hz for computational efficiency. EOG artifacts in EEG were removed by first creating a linear regression model to predict vertical and horizontal EOG from each EEG channel. Then, these predictions were subtracted from the corresponding channels [67]. These data were split into epochs starting 2 s before, and ending 0.5 s after, the participants' response. We then executed an artifact rejection process inspired by [68]: (a) each EEG epoch was baseline corrected using the average voltage from between 0.1 s before to 0.1 s after the response; (b) for each epoch and each channel a set of parameters, namely variance, amplitude range, channel deviation (the deviation from the channel average for each epoch), Hurst exponent and average correlation between such a channel and all other channels were calculated; (c) Z-scores for each epoch and parameter were calculated across all channels; (d) artifactual channel/epoch combinations were identified as those where the Z-score of one or more parameters deviated from zero by more than ± 3 and were removed from the data set. The remaining epochs were then individually referenced to the mean of all channels using a common average filter, Z-score normalised and, again, baseline corrected as explained above.

Finally, EEG channels were averaged over four key brain regions: occipito-parietal (electrodes P9, P7, P5, P3, P4, P6, P8, P10, PO7, PO3, POz, PO4, PO8, O1, Oz, O2 and Iz), central (electrodes FC1, FCz, FC2, C3, C1, Cz, C2, C4, CP3, CP1, CPz, CP2, CP4, P1, Pz and P2), left fronto-temporal (electrodes Fp1, Fpz, AF7, AF3, AFz, F7, F5, F3, F1, Fz, FT7, FC5, FC3, T7, C5, TP7 and CP5) and right fronto-temporal (electrodes Fp2, FPz, AF6, AF4, Afz, F8, F6, F4, F2, Fz, FT8, FC6, FC4, T8, T6, TP8 and TP6). We chose these regions because typically brain activity associated with confidence evaluation is in the frontal and parietal regions [7, 16–19]. EOG epochs were Z-score normalised based on the average of both the horizontal and vertical EOG. In order to further reduce computational load, both EEG and EOG epochs were then low-pass filtered to 4 Hz and down-sampled to 8 Hz. This is appropriate as our findings in eight other decision making and confidence reporting experiments showed that the ERPs that varied as a function of confidence were relatively low-frequency P300-like potentials [36]. Finally, for each epoch, the average EEG samples for the four regions and the EOG samples in the epoch were concatenated to form a feature vector to be used by the machine learning component of the BCI. This feature vector had 120 elements: 6 (4 EEG + 2 EOG features per sample) $\times$ 8 (sampling frequency in Hz) $\times$ 2.5 (epoch duration in seconds).

3.5. Meta-learning algorithm adaptations

Following the formal framework provided in section 2.1, below we provide details on our implementation and specialisation of meta-learning by biased regularisation for confidence prediction.

An important adaptation is that we have added a validation step called *validation fine-tuning* that uses equation (1) but optimizes the hyperparameter λ with unseen data (see below). This was added as it was hoped it would improve the performance of meta-learning when only little data from the target evaluated participant was available. Furthermore, due to the inter-session variability and non-stationarity of EEG, as already explained section 2.1, in this study we

Algorithm 1. General optimization framework.

1: **requires:** hyper-parameter λ
2: **initialization:** θ_0 model meta-parameters
3: **for** $t = 1$ to T **do**
4: Sample $\mu_t \sim \rho$ task, $Z_t \sim \mu_t^n$ dataset
5: Solve inner problem (e.g. gradient descent)
6: Update $\theta_t = \theta_{t-1} - \frac{1}{\sqrt{2\,t}} \nabla \boldsymbol{L}_{Z_t}(\theta_{t-1})$
7: **end forreturn** $\frac{\sum_{t=1}^{T} \theta_t}{T}$ estimate $\theta^{(\lambda)}$ of the solution of problem equation (2)

treat the data collected in each experimental session as a individual dataset.

3.5.1. Meta-training

As explained in section 2.1, during Meta-Training the goal is to solve the problem in equation (2) with respect to the mean meta-parameter θ. This optimization is performed in an online fashion on one dataset at each iteration using algorithm 1. For each regularization parameter λ we estimate the best meta-parameter $\theta^{(\lambda)}$ through this process. This is then repeated for a range of values of the hyper-parameter λ, and the best value is selected by a validation procedure described in the next section.

3.5.2. Validation fine-tuning

Having recovered a meta-parameter $\theta^{(\lambda)}$, we now want to move on to unseen datasets and optimize with respect to w for each dataset while evaluating the performance. The idea is that the meta-parameter will transfer some information from the training tasks to the unseen tasks. In more detail, for the validation step we used previously unseen tasks. Using each fixed $\theta^{(\lambda)}$ meta-parameter we solve the ridge regression problem in equation (1) on a subset of datapoints within each validation dataset, thereby computing a weight vector w. For the same validation dataset, we then compute the test error of w on the hold-out subset. Finally, we average that performance measurement across all validation tasks and select the best hyper-parameter λ^* and corresponding mean vector $\theta^{(\lambda^*)}$ that offers the best validation performance.

3.5.3. Target fine-tuning

For the test step of the process, naturally we follow the exact same pipeline as we did for validation. We use the dataset of the to-be-evaluated participant and optimise equation (1) using the fixed $\theta^{(\lambda^*)}$ meta-parameter.

3.6. Confidence prediction approaches

We performed an evaluation of multiple alternative systems (including the meta-learning one) for each participant independently. In each evaluation, the participant whose confidence was predicted was referred to as *target*, while the remaining nine participants were referred to as *source*. The number of samples used from each participant was defined by two values: the *source proportion* and the *target proportion*, which represent the proportion of source data and target data used for training, respectively. We performed different analyses by modifying these values (this will be explained in section 4).

3.6.1. Datasets

The data was divided into five sets (the subsection architecture details how each one of these datasets was used by each approach):

- *Meta-training.* This set comprised eight out of the nine source participants, randomly selected. The proportion of trials used from this dataset was determined by the *source proportion*.
- *Validation fine-tuning.* Random samples from the remaining source participant were selected for this set. The number of samples was proportional to the *target proportion*
- *Validation assessment.* The remaining samples for the participant used in the previous set were used.
- *Target train.* Random samples from the target participant were selected for this set. The number of data samples was proportional to the *target proportion*.
- *Target assessment.* The remaining samples from the target participant were used.

Figure 2 shows an example of how the data was divided for a given target participant for the meta-learning algorithm.

Prior to use the data for training, an outlier detection algorithm based on Tukey's fences was used. Specifically, for each feature, the first and third quartiles of the training data were calculated. Two thresholds were then defined: the third quartile plus 1.5 times the interquartile distance, and the first quartile minus 1.5 times the interquartile distance. A sample was deemed an outlier if it was above the former or below the latter thresholds. Outliers were replaced with the value of the closest threshold.

3.6.2. Prediction methods

We considered four confidence-prediction approaches, most of which used *ridge regression*, and which differed in how the five datasets defined above were used to train, validate and test the models (please refer to table 1 for a summary):

- *Naive.* This predictor, which is a baseline for the other approaches, used the average reported confidence from the target train set as prediction for the target assessment set.
- *Single-subject (SS).* We trained a ridge regression model for *each day* of the target participant using the target train set. Testing was done with the target assessment set.

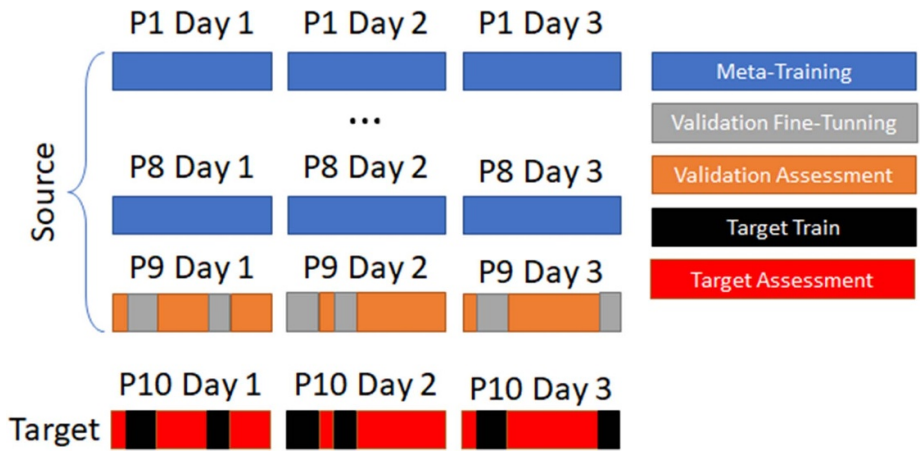

Figure 2. Example data split in this study. Here participant 10 is the target and participants 1 to 9 are the source. Participant 9 is used for validation fine-tuning. In this example, 100% of the *source proportion* is used to meta-train the algorithm.

Table 1. Datasets used by each approach.

Dataset name for meta-learning	Naive	Single-subject	Multi-subject
Meta-training			Training
Validation fine-tuning			Training
Validation assessment			Training
Target train	Training	Training	Training
Target assessment	Test	Test	Test

- *Multi-subject (MS).* The meta-training, validation fine-tuning, validation assessment and target train datasets were concatenated for training, which included the optimization of the regularization parameter λ. The model was trained with data from all 3 days. The use of the target train dataset is a departure from the zero-training approach reported in [36] where only data from source participants were used for training[7].
- *Meta-learning (Meta).* This method learns a meta-parameter $\theta^{(\lambda)}$ from the meta-training set, as outlined in section 3.5. Then the validation fine-tuning set is used to determine the best hyper-parameter λ (as described previously in section 3.5.2) while the validation assessment set is used evaluate the overall performance. Then, for each day, a model was fine-tuned using the target train set and tested using the corresponding target assessment set.

Two of the four described methods, MS and SS, can be considered as extreme cases of Meta. SS works like the implementation of Meta when only fine-tuning is performed, i.e. the algorithm only uses data from the current target. Conversely, MS can be seen as the implementation of Meta where no fine-tuning is performed and knowledge is only transferred from the source data.

For each training and validation step as well as hyperparameter optimization we used the Normalized mean squared error (NMSE) as a performance measure (see section 3.7).

The regularization parameter λ used for the regression models used in SS, MS and Meta were selected from 100 values evenly spaced on a log scale between 10^{-12} and 10^{4}.

For SS and MS the best λ was selected by splitting randomly the training set into 70%–30% for training and validation, respectively. This splitting was the same for all the values of λ. After selecting the best regularization parameter as the one with the lowest validation error, the model was trained again using the full training data with the selected regularization parameter to, then, be tested on the corresponding test set.

For the Meta approach, three regularization parameters, all using the same range of values previously described, needed to be selected λ_{train}, $\lambda_{\text{validation}}$ and λ_{test}. First, for each possible value of λ_{train}, a model was trained using the meta-training data. Each of these trained models, was then fine-tuned with the validation fine-tuning data. To do this, the best $\lambda_{\text{validation}}$ was selected using the aforementioned random 70%/30% split. Each fine-tuned model (one for each possible λ_{train}) was evaluated using the validation assessment data. Then, the λ_{train} that lead to the best performance, $\lambda^{*}_{\text{train}}$, was selected. Finally, the

[7] MS is an extremely simple form of transfer learning, which should be expected to behave quite similarly to its zero-training ancestor. Thus the models identified via MS will be general but are expected to have a worse performance than those produced by Meta or SS.

fine-tuning procedure of the model corresponding to λ^*_{train} was repeated using the target train data. With this last step, the optimal λ_{test} was obtained in the same way as the $\lambda_{\text{validation}}$.

3.7. Performance evaluation metrics

To evaluate the performance of each system we used three metrics: the NMSE, the Fisher discriminant ratio (FDR) and the Pearson correlation coefficient.

The NMSE is defined as:

$$\text{NMSE} = \frac{\text{MSE}(c^{\text{p}}, c^{\text{r}})}{\sum_{i=1}^{n}(\mu^{\text{r}} - c_i^{\text{r}})^2}, \tag{3}$$

where n is the number of samples, c^{p} is the array of predicted confidence values, c^{r} is the corresponding array of reported confidence values, μ^{r} is the mean reported confidence, and MSE is the mean squared error defined as:

$$\text{MSE}(c^{\text{p}}, c^{\text{r}}) = \sum_{i=1}^{n}(c_i^{\text{p}} - c_i^{\text{r}})^2. \tag{4}$$

This metric allows us to estimate how close the prediction of the confidence is to the reported one. NMSE is used here because it normalises the error across different datasets, thus making them comparable.

As discussed in [36], just considering the error of the prediction is not enough to evaluate the quality of the confidence prediction. In particular, given that confidence is a (subjective) evaluation of task performance, we need to evaluate how well the confidence is modulated by decision accuracy. To measure this, we calculated the FDR which evaluates how separable two distributions are and is defined as:

$$\text{FDR}(x_{\text{c}}, x_{\text{in}}) = \frac{n_{\text{c}}(\mu_{\text{c}} - \mu)^2 + n_{\text{in}}(\mu_{\text{in}} - \mu)^2}{n_{\text{c}}\sigma_{\text{c}}^2 + n_{\text{in}}\sigma_{\text{i}}^2}, \tag{5}$$

where x_{c} and x_{in} are the confidence predictions for the correct and incorrect trials, n_{c} and n_{in} are the number of correct and incorrect trials, μ_{c} and μ_{in} are the mean prediction for correct and incorrect trials, and σ_{c} and σ_{i} are the standard variation of correct and incorrect trials, respectively, while μ is the mean of all the confidence predictions. Note that this value can be calculated for both the predicted confidence (splitting c^{p} between correct and incorrect as incorrect trials) and the reported confidence (splitting c^{r} between correct and incorrect trials). However, in all our analysis we used c^{p}. FDR takes into consideration not only the average and dispersion of each class but also the imbalance in the number of samples in each class, which is important as in our dataset there is a significant imbalance between the correct and incorrect decision classes.

People can be systematically over or under confident. Thus in some situations it may be important for a predictor to capture the variations in confidence relative to the participant's average confidence, more than the precise confidence values. Analyse this this we used the Pearson correlation coefficient r, which essentially describes the linear relationship between predicted and reported confidence. This is given by

$$r = \frac{\text{cov}(c^{\text{p}}, c^{\text{r}})}{\sigma_{c^{\text{p}}}\sigma_{c^{\text{r}}}}. \tag{6}$$

3.8. State of the art

Naturally, Meta is not the only way of implementing transfer learning. For instance, transfer component analysis (TCA) [69] is a well known transfer learning algorithm that uses knowledge from a source domain to adapt target data. Methods like TCA, that focus on domain adaption, are not well suited for our application as they need to have all the data in advance or require retraining the whole algorithm for each new data point, which is not practical as the algorithms are too computational demanding. Another way of transferring knowledge are methods based on adversarial neural networks [37, 70–72]. Most of these methods, such as domain adversarial neural networks (DANNs), were originally created for image recognition but have since adapted for EEG classification as well.

To compare Meta with a state-of-the-art algorithm, we decided to implement the DANN method using the configuration as given in [73][8]. For a fairer comparison, in [73] we used features extracted from the EEG as input for the network instead of using the network itself for feature extraction so that the same features were used for the all methods. In [73] a three class classification problem was used and, so, the output layer was a soft-max layer with three neurons. Naturally, this is not appropriate for the regression problem tackled in this paper. So, we replaced the soft-max layer with a single neuron with reLU activation.

4. Results

Most of the results reported in this section will compare the four algorithms described in section 3.6.2 in conditions where all of the data from the source participants are available to the machine learning, but we vary the amount of data available from the target participant as this directly impacts how quickly the participant can start using the BCI. More precisely, we varied the *target proportion* from 2.5% to 80% in steps of 2.5% while keeping the *source proportion* constant at 100%. However, at the end of the section we will also look at how performance changes as the *source proportion* varies, as this would be important for routine use, where data for training might be

[8] We used this implementation even if the problem was different because we were unable to find any paper using a domain transfer learning approach using neural networks from EEG data for a regression problem.

acquired only in the initial phases of participants' sessions. Given that, as indicated in section 3.1, sessions lasted approximately 2.5 h, a (training or source) proportion p translates into a training-set acquisition time of $150 \times p$ min. So, for instance, a 30% *target proportion* requires participants to spend 45 min providing training data for the BCI.

In the analyses reported in this section, the data were randomly sampled from the full source and target datasets without replacement until the desired percentage was reached. To reduce stochastic effects, we averaged the results over 30 different repetitions of this random sampling where the same 30 randomizer seeds were used when testing alternative confidence-prediction methods.

4.1. Performance evaluation

Performance evaluation results are reported in figure 3. This shows how the performance of the four prediction approaches and the DANN method varies as a function of the *target proportion* for the three performance measures—NMSE, FDR and correlation coefficient—described in section 3.7, averaged across all participants (grand averages)[9]. The shaded area around each line indicates the standard error across the 30 independent randomisations of the data.

Looking at our primary performance measure, NMSE (left plot) for which lower values are better than higher ones, we see that Meta is superior to all other approaches except for very small *target proportions* where it is comparable to Naive (which returns the average reported confidence from the target train set). SS is second best but it needs at least 20% of target data to do any better than Naive, beyond which its performance asymptotically approaches that of Meta. SS can achieve similar performance to Meta, but it needs more data. For instance, to reach the performance that Meta has at 20% of *target proportion* (which corresponds to approximately 30 min of training-data collection), SS needs double the data (approximately 40%, i.e. one hour of data collection). After a short transient required for the average to stabilize, Naive's performance does not improve with the amount of target data and, so, it cannot compete with Meta and SS. As expected, MS is significantly worse than other algorithms since it learns primarily from other participants, being only marginally different from a zero-training approach. Interestingly, we see that MS is also worse than Naive at all *target proportions* considered.

Finally, DANN, the only method that used neural networks, was unable to beat Naive in terms of NMSE, but it obtained better results than MS.

Looking at the FDR (middle plot) and correlation (right plot) performance measures, for which higher values are better than lower ones, we see a slightly different picture for very low values of *target proportion*: MS does better than the other methods up until a certain proportion, after which the picture becomes the same as for NMSE: Meta dominating and SS being second but asymptotically approaching the performance of Meta. For FDR, DANN shows worse values than MS until 20% of *target proportion* after which they are the same. The correlation values for DANN never reach those of MS, getting closer as the *target proportion* grows. Naturally, Naive offers the worst performance as it always return the same value and, so, both FDR and correlation are zero irrespective of the *target proportion*.

4.2. Statistical analyses

To verify the statistical veridicity of these findings we compared the results obtained with Meta in the 30 randomisations of the dataset for each level of *target proportion* with the corresponding results obtained with SS, MS and Naive using one-tailed Wilcoxon rank sum tests. We accounted for the multiple comparison problem by applying a Benjamini–Hochberg correction (which controls the false discovery rate). The corresponding p-values for the three performance metrics are shown in the three plots in figure 4.

For NMSE (left plot), Meta is statistically superior to all other prediction methods most of the time, thereby confirming the findings of figure 3. There are only two exception: (a) at *target proportions* above 70% Meta is not statistically better than SS (as one would have expected seeing the asymptotic convergence of the NMSEs for the two methods in figure 3), (b) at extremely low *target proportions* ($\leqslant 2.5\%$), Meta is not statistically better than MS and Naive.

Also the plots for FDR (middle) and correlation (right) in figure 4 confirm the observations from figure 3. Meta performs better than MS for *target proportion* above 30% for FDR and 22% for correlation. Also, Meta performs significantly better than SS below *target proportions* of 27% for FDR and 60% for correlation. Also, as expected, Naive is always significantly worse than Meta.

To assess to what degree the relative performance of Meta might vary across participants, we performed an additional analysis at individual participant level and recorded the number of participants for which Meta was significantly better than other confidence-prediction approaches. Once again, we used a one-tailed Wilcoxon rank sum test ($p < 0.05$) with Benjamini–Hochberg correction.

Figure 5 reports the results for NMSE as this is our key performance measure. Interestingly, Meta is significantly better than all other methods for most participants with the exception of very small and very large *target proportions* where it is not statistically superior to SS, Naive and MS.

Finally, the mean absolute error (MAE) of the prediction is featured in figure 6 to give the reader

[9] Due to the mini-batch size (96), we were not able to train the DANN model with the two smallest target proportions (2.5% and 5%) as there were not enough samples for the batch training.

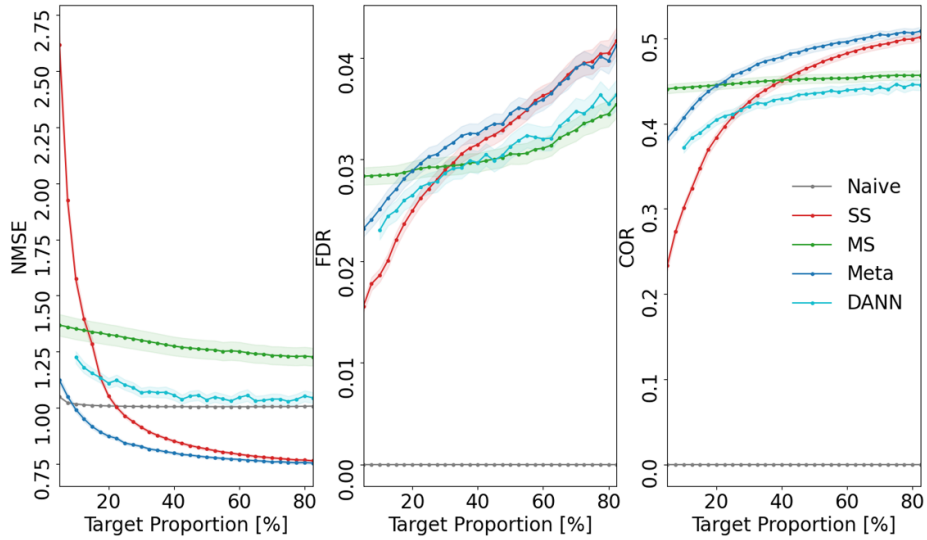

Figure 3. Performance of NMSE, FDR and correlation for the four different approaches and the DANN method with a *source proportion* of 100% and increasing *target proportion*.

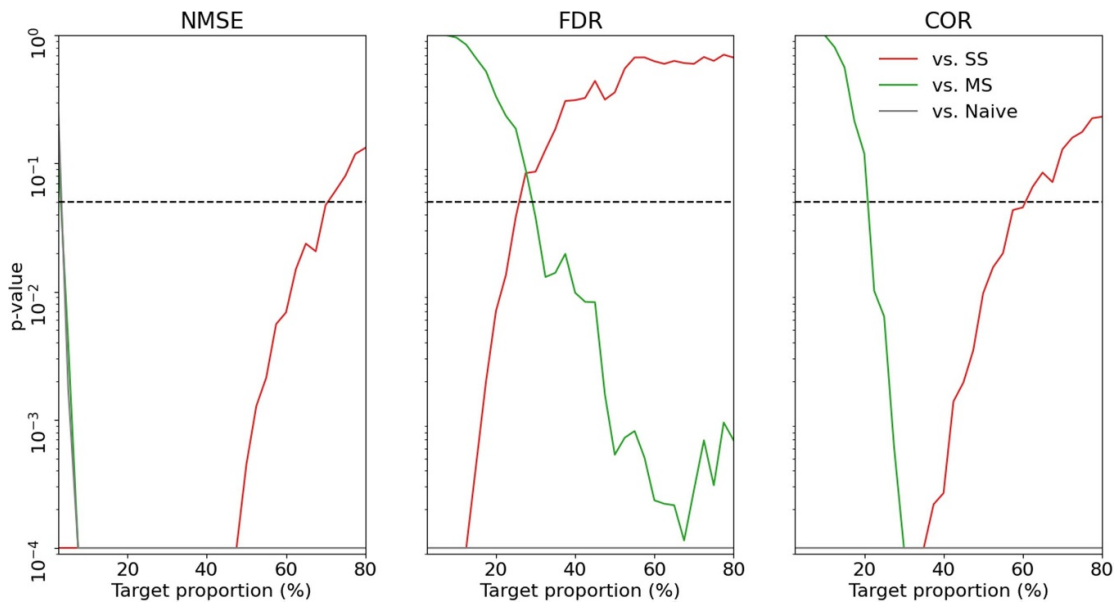

Figure 4. Benjamini–Hochberg corrected *p*-values for the performance measures NMSE, FDR and correlation, when comparing Meta with MS, SS and Naive different *target proportions*. Any *p*-values lower than 10^{-4} were clipped to 10^{-4} to improve readability. Values below the dashed line indicate statistical significance.

a more intuitive reference to interpret the algorithm performance. MAE for the four approaches is presented in a similar fashion as NMSE and denotes to:

$$\mathrm{MAE}(c^{\mathrm{p}}, c^{\mathrm{r}}) = \sum_{i=1}^{n} |c_i^{\mathrm{p}} - c_i^{\mathrm{r}}|. \tag{7}$$

It also allows comparison with our previous publication [36]. The most comparable conditions are, SS at 80% *target proportion* which has a MAE of 0.14, and MS at 0.025% *target proportion* which has a MAE of 0.2. Both values are nearly identical to those obtained in our previous study, further validating these new results.

4.3. Comparison of learned models

The asymptotic convergence of all the performance metrics for Meta and SS as the *target proportion* increases seen in figure 3 suggests that the coefficients of the corresponding ridge regression models might also be converging. Conversely one would expect the regression models of Meta and MS to be significantly different at all target proportions. To test these hypotheses, we measured the Euclidean

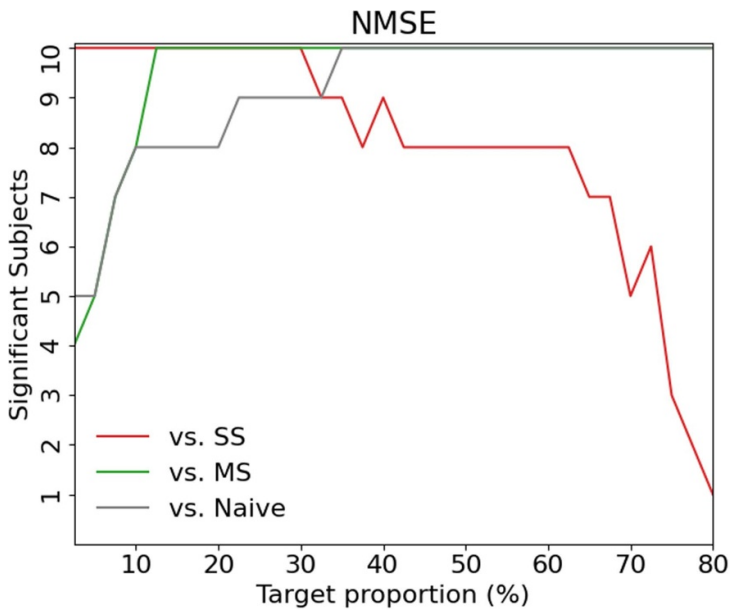

Figure 5. Number of subjects for which Meta's NMSE is statistically better than that of MS, SS and Naive for a *source proportion* of 100% and increasing *target proportion*.

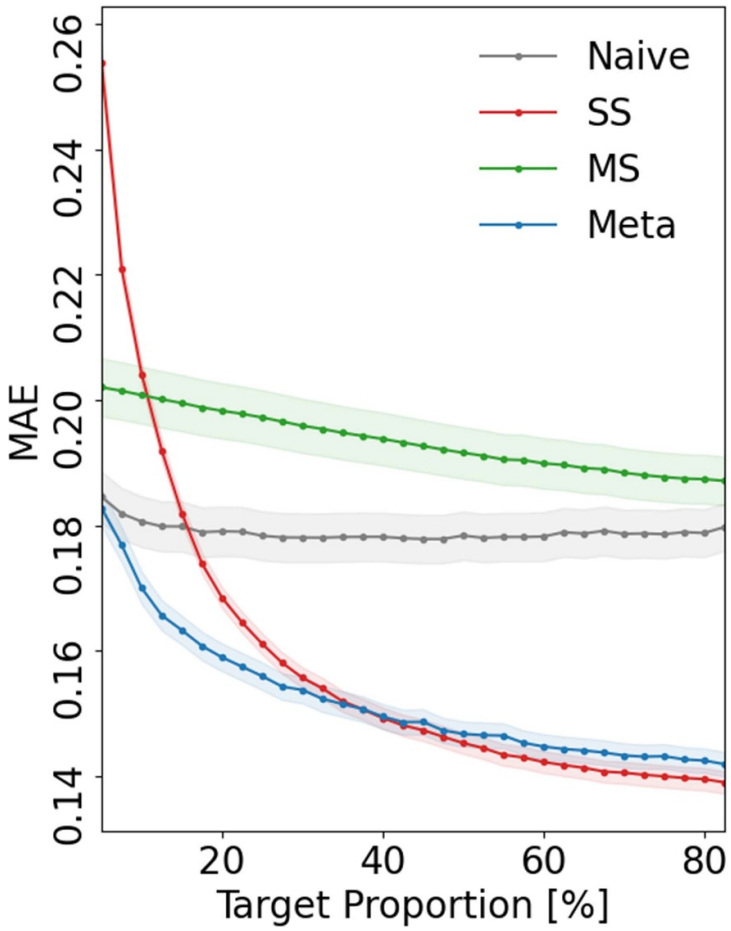

Figure 6. MAE for the four different approaches with a *source proportion* of 100% and increasing *target proportion*.

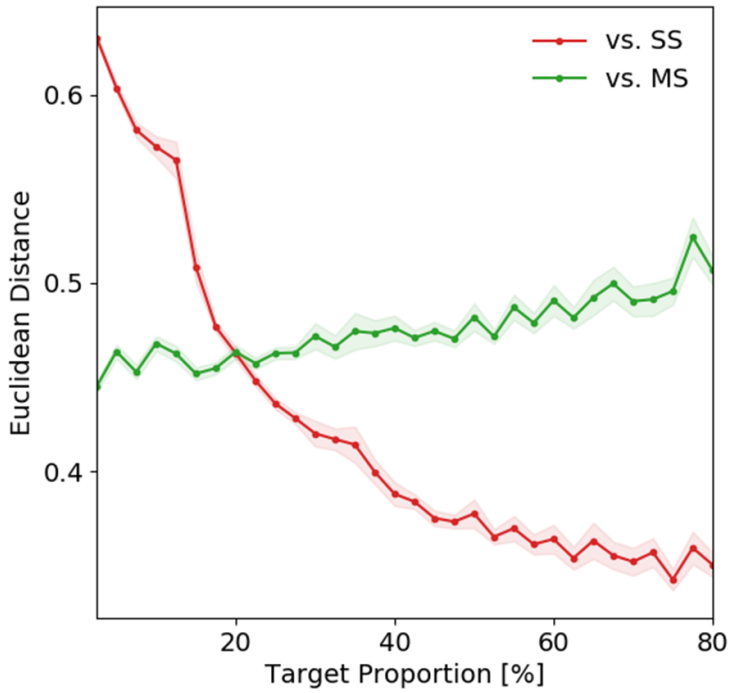

Figure 7. Euclidean distance between the vectors of regression coefficients produced by Meta and those produced by SS and MS as a function of the *target proportion*.

distance between the vectors of regression coefficients found by Meta and those found by SS and MS for varying *target proportions*. Results are shown in figure 7.

While we see that the distance between the Meta and SS models decreases as the target proportion increases, the distance by no means drops down to zero. This suggests that there are many model configurations that roughly produce the same quality of predictions, probably as a result of redundancies in the input feature set. We also see, as expected, that the MS and Meta models start different and remain different (in fact even diverging slightly) as the target proportion increases.

To look at this with a finer level of granularity, in figure 8 we report the average (across participants and seeds) of the sum (across time steps) of the coefficients associated to each area of the brain (feature group) defined in section 3.4 as a function of the *target proportion*. The figure also reports the average (across participants and seeds) of the bias term. For each confidence-prediction method, the sign of the values being plotted indicates whether a positive voltage/ERP in the corresponding areas of the brain increases or decreases the confidence estimation (and vice versa for negative voltages).

At this level of granularity it appears that MS is very weakly influenced by the data from the target participant, irrespective of its proportion. Also, MS's average coefficients (by brain area) appear to be quite different from those used in SS and Meta. This is the main reason for the dissimilarity between Meta and MS highlighted in figure 7.

Also Meta's brain-region average coefficients and bias term appear to be weakly influenced by the data from the target participant. However, adjustments take place in the horizontal and vertical EOG coefficients as the target proportion varies.

Finally we see that SS's coefficients typically start far from the corresponding ones in MS and Meta. However, the EEG coefficients of SS tend to become quite similar to those of Meta as the target proportion is increased. Also, SS's EOG coefficients and bias term tend to move in the direction of the corresponding terms in Meta. However, in this case they never get all the way there. This is the main cause of the residual differences between SS and Meta observed in figure 7.

4.4. Source analysis

The performance of machine learning algorithms depend crucially on the amount of data available for training. As explained above, in a transfer learning setting, there are two types of training data: source data (in our case, data from different participants) and target data (data from the participant that will use the BCI after training). In the previous sections we have analysed in depth the performance and behaviour of different algorithms as we varied the *target proportion*. In this section we want to briefly look at the effects on performance of varying the proportion of source data. So, for this analysis we varied the *source proportion* from 5% to 100% in steps of 5% while keeping the *target proportion* constant at 25% (the latter value was chosen as it represents a typical offline training session for an online BCI).

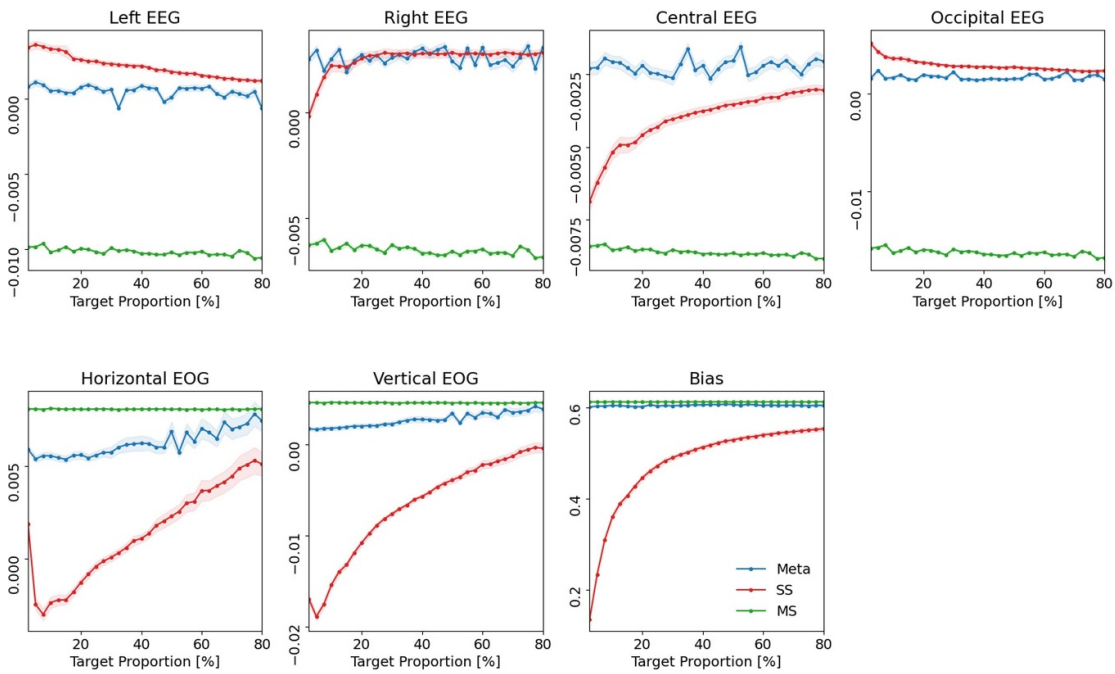

Figure 8. Normalized weights summed up for each feature group representing the importance of each feature group for *source proportion* of 100% and increasing *target proportion*. The EEG groups are displayed in the top row while vertical and horizontal EOG as well as the bias are displayed in the bottom row.

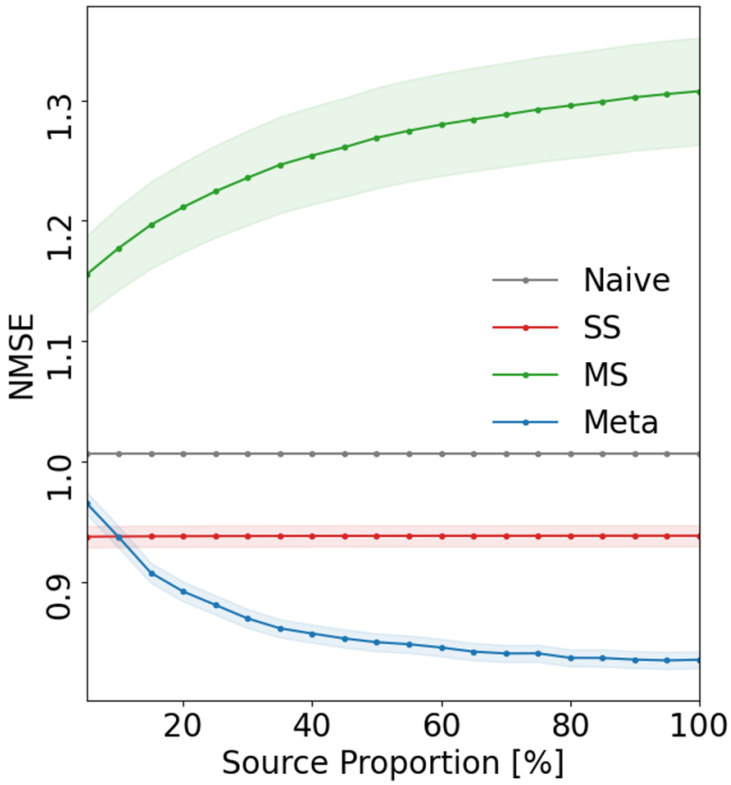

Figure 9. NMSE for the four different confidence-prediction approaches as a function of the *source proportion* (for 25% *target proportion*).

Figure 9 shows how the NMSE performance metric varied for the four different confidence-prediction approaches considered in this study as a function of the *source proportion*. Because Naive and SS do not use any source data their NMSE is constant. Unsurprisingly, we see that MS's performance is worse than for other approaches and gets worse as more and more source data 'dilute' that target data available to

the algorithm. On the contrary, the performance of Meta improves significantly as the *source proportion* is increased, tending to plateau after the *source proportion* reaches 70% or 80% indicating that the algorithm has no problem correctly handling any amount of source data.

5. Discussion

Due to the variability and noise of EEG/EOG data, usually BCIs are trained with data from a specific participant (i.e. using a single-subject approach). Attempts to learn from multiple participants normally lead to BCIs that generalise somehow across participants but at the cost of much reduced individual performance. For instance, this is what recently happened in [36], in the area of confidence prediction.

When considering the main performance criterion, NMSE, the present work shows that for the task of confidence prediction even better performance can be obtained by training a BCI using a specialisation of meta-learning by biased regularisation [59]. The meta-learning method was adapted from [59] and has two phases: *meta-training* where the goal is to iteratively train a BCI with data from previous participants to learn domain-invariant features, and *fine-tuning*, which uses data from a new participant to quickly fine-tune the BCI. Our specialisation of it is the addition of a *validation fine-tuning* step that uses equation (1) but optimizes the hyperparameter λ with unseen data. In effect this algorithm combines the strengths of the mainstream *single-subject* approach with the generality of *multi-subject*, without compounding their weaknesses, so that it outperforms them both in most conditions. Meta also performs better that a state-of-the-art algorithm called, DANN, under all criteria considered.

It is also important to note that Meta does not only perform significantly better than most algorithms, it also has unique positive properties setting it further apart and making it harder to directly compare it to other algorithms. For instance, the training (with source data) and the fine-tuning (with target data) can be done sequentially, i.e. the initial training can be done in advance, leaving only the, relatively fast, fine-tuning step to be performed during the deployment, as only the Meta coefficients need to be saved. Conversely, MS and DANN, require both the source and target data to train the model. This also means that Meta does not have to be fully retrained and can be updated with each new data point as opposed to, for example, TCA. Since Meta only utilizes one hyperparameter it does not need to be specifically tailored to a problem as compared to methods based on deep neural networks such as adversarial discriminative domain adaptation. Finally, the low computational demand allows it to be used easily as opposed to the more demanding algorithms such as TCA or deep neural networks.

Additionally, we found that, with meta-learning, much less subject specific training data is required to achieve a given level of performance than for the traditional single-subject approach. The superiority of meta-learning algorithm in situations where limited data from a new participant are available for training/tuning is a real-plus in BCI applications, as new users need to train the BCI for a much shorter period.

Also, we found that meta-learning delivered reliable performance improvements across participants having the best performance for most participants in most conditions (see figure 5).

Another issue we explored is how sensitive to the amount of source data (data from other participants) our meta-learning algorithm was. We explored this in section 4.4 and found that the meta-learning algorithm asymptotically improves as more and more source data are available. Unsurprisingly, we found that the *multi-subject* transfer-learning method is too simple to make good used of large volumes of source data.

While the performance of the machine learning algorithms was optimised during training using NMSE, meta-learning provided a welcome superiority with respect to the two additional metrics, FDR and correlation, with respect to its most serious contender, *single-subject*, for low data regimes (below a *target proportion* of 27% for FDR and 60% for correlation). Superior FDR indicates that the confidence predictions produced by meta-learning are better modulated by the correctness of the decision, i.e. it is a superior (self) evaluation of task performance. Superior correlation indicates that meta-learning predictions better capture the variations in confidence relative to the participant's average confidence.

A question that naturally comes to mind is: where does the superiority of the regression models produced by meta-learning comes from? In other words, what is special about such models that is not available in either the *single-subject* and *multi-subject* regression models? Because each model is a 121-dimensional vector and it is participant-dependent, it is not trivial to answer such questions. The analyses reported in section 4.3 are our best first attempt to understand differences and similarities between the optimal regression models found by different approaches. In this work we found that *meta-learning* and *multi-subject* tend to produce very different models, particularly in relation to the EEG-feature coefficients. Because there is so much more source data than target data in *multi-subject* training sets, unsurprisingly, MS's models do not change much as more and more data from the target participant are available. In the *meta-learning* case, instead, we can see that fine-tuning can make use of the data, particularly for adjusting the EOG coefficients of the model. The

single-subject case is more extreme. Here we see a very big influence of the target data, with the regression models being quite different from the ones produced by *meta-learning* at low data regimes, but asymptotically getting nearer and nearer them as more and more target data are made available. They never fully converge (see figure 7), and yet their performance nearly converges (see red and blue lines in figure 3(left)). A possible interpretation of this effect is that effect in that there exist many local optimal solutions in the feature space and since Meta is also minimizing with respect to the meta-parameter θ it finds a different solution than *single-subject*. An alternative hypothesis is that there is in fact only one global optimum but that the performance surface around it is nearly flat and, had more (say twice as much) target data been available, *single-subject* would have eventually found the same solution as the *meta-learning* algorithm.

6. Conclusions

In this study we investigated whether a new powerful transfer-learning algorithm, *meta-learning by biased regularisation*, could improve the performance of BCIs for confidence prediction in a difficult target discrimination task based on video feeds with respect to other machine learning methods.

We adapted the algorithm to the confidence prediction problem from EEG and EOG data on a decision-by-decision basis. The method exploits previous participants' data to produce an algorithm that is then quickly tuned to new participants. We compared it with the traditional single-subject training almost universally adopted in BCIs, a transfer-learning adaptation of a zero-training method we used recently for a similar task, a state-of-the-art transfer learning algorithm—DANNs—and with a simple baseline algorithm. We found that meta-learning is significantly better than other approaches in most conditions, and much better in situations where limited training/tuning data from a new participant are available.

Meta-learning was also very robust in the presence of small training sets and across participants. Furthermore, compared with many other transfer learning techniques, meta learning requires less computation, does not have to save all participant data, can be partially trained beforehand, does not to be specifically tailored to one problem and can be updated fast with each new recorded data point. All of these properties are beneficial in BCI applications.

Additionally, here confidence estimation was performed on a continuous scale, as opposed to the few discrete levels used in previous work.

The study has also some limitations. Firstly, we used linear regression models. While we have not seen large performance differences between our current model and the artificial neural network that we have used in previous work [36], it is not unlikely that non-linear models could further improved performance. Secondly, we did not explore in detail how the models found by different methods differed at participant level. Thirdly, we did not explore all possible proportions of source and target data. Fourthly, we did not test the BCI in an online experiment. We hope to address all these limitations in future work. In future work we will also attempt to adapt the meta-learning algorithm to other forms of BCI and and compare it to deep neural network methods such as adversarial discriminative domain adaptation that can be specifically designed for confidence prediction.

Data availability statement

The data generated and/or analysed during the current study are not publicly available for legal/ethical reasons but are available from the corresponding author on reasonable request.

Acknowledgments

This article is an overview of UK MOD sponsored research and is released for informational purposes only. The contents of this article should not be interpreted as representing the views of the UK MOD, nor should it be assumed that they reflect any current or future UK MOD policy. The information contained in this article cannot supersede any statutory or contractual requirements or liabilities and is offered without prejudice or commitment. The Authors acknowledge support of the UK Defence Science and Technology Laboratory (Dstl) and Engineering and Physical Research Council (EPSRC) under Grant EP/P009204/1. This research is part of the collaborative project funded by US DoD, UK MOD and UK EPSRC (EP/P009204/1 and EP/P009069/1) under the Multidisciplinary University Research Initiative (MURI).

ORCID iDs

Christoph Tremmel https://orcid.org/0000-0002-0324-6626
Jacobo Fernandez-Vargas https://orcid.org/0000-0002-6941-1233
Caterina Cinel https://orcid.org/0000-0001-8668-4017

References

[1] Gold J I and Shadlen M N 2007 The neural basis of decision making *Annu. Rev. Neurosci.* **30** 535–74
[2] de Lange F P, Jensen O and Dehaene S 2010 Accumulation of evidence during sequential decision making: the importance of top-down factors *J. Neurosci.* **30** 731–8
[3] Cheadle S, Wyart V, Tsetsos K, Myers N, De Gardelle V, Herce Castañón S H and Summerfield C 2014 Adaptive gain control during human perceptual choice *Neuron* **81** 1429–41
[4] Polanía R, Krajbich I, Grueschow M and Ruff C C 2014 Neural oscillations and synchronization differentially

support evidence accumulation in perceptual and value-based decision making *Neuron* **82** 709–20
[5] Sanders J I, Hangya B and Kepecs A 2016 Signatures of a statistical computation in the human sense of confidence *Neuron* **90** 499–506
[6] Graziano M, Parra L C and Sigman M 2015 Neural correlates of perceived confidence in a partial report paradigm *J. Cogn. Neurosci.* **27** 1090–103
[7] Kepecs A, Uchida N, Zariwala H A and Mainen Z F 2008 Neural correlates, computation and behavioural impact of decision confidence *Nature* **455** 227–31
[8] Grimaldi P, Lau H and Basso M A 2015 There are things that we know that we know and there are things that we do not know we do not know: confidence in decision-making *Neurosci. Biobehav. Rev.* **55** 88–97
[9] Adler W T, Ma W J and Gershman S J 2018 Comparing Bayesian and non-Bayesian accounts of human confidence reports *PLoS Comput. Biol.* **14** e1006572
[10] Ratcliff R and Starns J J 2013 Modeling confidence judgments, response times and multiple choices in decision making: recognition memory and motion discrimination *Psychol. Rev.* **120** 697
[11] Baranski J V and Petrusic W M 1994 The calibration and resolution of confidence in perceptual judgments *Percept. Psychophys.* **55** 412–28
[12] Zander T O, Krol L R, Birbaumer N P and Gramann K 2016 Neuroadaptive technology enables implicit cursor control based on medial prefrontal cortex activity *Proc. Natl Acad. Sci.* **113** 14898–903
[13] Heekeren H R, Marrett S, Bandettini P A and Ungerleider L G 2004 A general mechanism for perceptual decision-making in the human brain *Nature* **431** 859–62
[14] Christopoulos G I, Tobler P N, Bossaerts P, Dolan R J and Schultz W 2009 Neural correlates of value, risk and risk aversion contributing to decision making under risk *J. Neurosci.* **29** 12574–83
[15] Philiastides M G and Sajda P 2006 Temporal characterization of the neural correlates of perceptual decision making in the human brain *Cereb. Cortex* **16** 509–18
[16] Boldt A and Yeung N 2015 Shared neural markers of decision confidence and error detection *J. Neurosci.* **35** 3478–84
[17] Gherman S and Philiastides M G 2015 Neural representations of confidence emerge from the process of decision formation during perceptual choices *Neuroimage* **106** 134–43
[18] Gherman S and Philiastides M G 2018 Human VMPFC encodes early signatures of confidence in perceptual decisions *Elife* **7** e38293
[19] Herding J, Ludwig S, von Lautz A, Spitzer B and Blankenburg F 2019 Centro-parietal EEG potentials index subjective evidence and confidence during perceptual decision making *Neuroimage* **201** 116011
[20] Wolpaw J R, Birbaumer N, McFarland D J, Pfurtscheller G and Vaughan T M 2002 Brain–computer interfaces for communication and control *Clin. Neurophysiol.* **113** 767–91
[21] Müller-Putz G R, Schwarz A, Pereira J and Ofner P 2016 From classic motor imagery to complex movement intention decoding: the noninvasive Graz-BCI approach *Prog. Brain Res.* **228** 39–70
[22] Birbaumer N 2006 Breaking the silence: brain-computer interfaces (BCI) for communication and motor control *Psychophysiology* **43** 517–32
[23] Vaid S, Singh P and Kaur C 2015 EEG signal analysis for BCI interface: a review *2015 5th Int. Conf. on Advanced Computing and Communication Technologies* (IEEE) pp 143–7
[24] Wang X W, Nie D and Lu B L 2014 Emotional state classification from EEG data using machine learning approach *Neurocomputing* **129** 94–106
[25] Berka C, Levendowski D J, Lumicao M N, Yau A, Davis G, Zivkovic V T, Olmstead R E, Tremoulet P D and Craven P L 2007 EEG correlates of task engagement and mental workload in vigilance, learning and memory tasks *Aviat. Space Environ. Med.* **78** B231–44
[26] Tremmel C, Herff C, Sato T, Rechowicz K, Yamani Y and Krusienski D J 2019 Estimating cognitive workload in an interactive virtual reality environment using EEG *Front. Hum. Neurosci.* **13** 401
[27] Bhattacharyya S, Valeriani D, Cinel C, Citi L and Poli R 2021 Anytime collaborative brain-computer interfaces for enhancing perceptual group decision-making *Sci. Rep.* **11** 17008
[28] Cinel C, Valeriani D and Poli R 2019 Neurotechnologies for human cognitive augmentation: current state of the art and future prospects *Front. Hum. Neurosci.* **13** 13
[29] Valeriani D, Poli R and Schwenker F 2019 Cyborg groups enhance face recognition in crowded environments *PLoS One* **14** e0212935
[30] Valeriani D, Cinel C and Poli R 2017 Group augmentation in realistic visual-search decisions via a hybrid brain-computer interface *Sci. Rep.* **7** 1–12
[31] Valeriani D, Poli R and Cinel C 2017 Enhancement of group perception via a collaborative brain-computer interface *IEEE Trans. Biomed. Eng.* **64** 1238–48
[32] Valeriani D, Poli R and Cinel C 2015 A collaborative brain-computer interface for improving group detection of visual targets in complex natural environments *2015 7th Int. IEEE/EMBS Conf. on Neural Engineering (NER)* (IEEE) pp 25–28
[33] Poli R, Valeriani D, Cinel C and Chacron M J 2014 Collaborative brain-computer interface for aiding decision-making *PLoS One* **9** e102693
[34] Krumpe T, Gerjets P, Rosenstiel W and Spüler M 2019 Decision confidence: EEG correlates of confidence in different phases of an old/new recognition task *Brain Comput. Interfaces* **6** 162–77
[35] Kubanek J, Hill J, Snyder L H and Schalk G 2015 Cortical alpha activity predicts the confidence in an impending action *Front. Neurosci.* **9** 243
[36] Fernandez-Vargas J, Tremmel C, Valeriani D, Bhattacharyya S, Cinel C, Citi L and Poli R 2021 Subject- and task-independent neural correlates and prediction of decision confidence in perceptual decision making *J. Neural Eng.* **18** 046055
[37] Ganin Y, Ustinova E, Ajakan H, Germain P, Larochelle H, Laviolette F, Marchand M and Lempitsky V 2016 Domain-adversarial training of neural networks *J. Mach. Learn. Res.* **17** 2096–30
[38] Krusienski D J, Grosse-Wentrup M, Galán F, Coyle D, Miller K J, Forney E and Anderson C W 2011 Critical issues in state-of-the-art brain-computer interface signal processing *J. Neural Eng.* **8** 025002
[39] Rodrigues P L C, Jutten C and Congedo M 2019 Riemannian procrustes analysis: transfer learning for brain-computer interfaces *IEEE Trans. Biomed. Eng.* **66** 2390–401
[40] Pan S J and Yang Q 2010 A survey on transfer learning *IEEE Trans. Knowl. Data Eng.* **22** 1345–59
[41] Waytowich N R, Lawhern V J, Bohannon A W, Ball K R and Lance B J 2016 Spectral transfer learning using information geometry for a user-independent brain-computer interface *Front. Neurosci.* **10** 430
[42] Rodrigues P L C, Jutten C and Congedo M 2019 Riemannian procrustes analysis: transfer learning for brain-computer interfaces *IEEE Trans. Biomed. Eng.* **66** 2390–401
[43] Ming Y, Ding W, Pelusi D, Wu D, Wang Y K, Prasad M and Lin C T 2019 Subject adaptation network for EEG data analysis *Appl. Soft Comput.* **84** 105689
[44] Jayaram V, Alamgir M, Altun Y, Scholkopf B and Grosse-Wentrup M 2016 Transfer learning in brain-computer interfaces *IEEE Comput. Intell. Mag.* **11** 20–31
[45] Dai M, Zheng D, Liu S and Zhang P 2018 Transfer kernel common spatial patterns for motor imagery brain-computer interface classification *Comput. Math. Methods Med.* **2018** 1–9

[46] Waytowich N R, Faller J, Garcia J O, Vettel J M and Sajda P 2016 Unsupervised adaptive transfer learning for steady-state visual evoked potential brain-computer interfaces *2016 IEEE Int. Conf. on Systems, Man and Cybernetics (SMC)* (IEEE) pp 004135–40
[47] Chai X, Wang Q, Zhao Y, Li Y, Liu D, Liu X and Bai O 2017 A fast, efficient domain adaptation technique for cross-domain electroencephalography(EEG)-based emotion recognition *Sensors* **17** 1014
[48] Lin Y P and Jung T P 2017 Improving EEG-based emotion classification using conditional transfer learning *Front. Hum. Neurosci.* **11** 334
[49] Lin Y P, Jao P K and Yang Y H 2017 Improving cross-day EEG-based emotion classification using robust principal component analysis *Front. Comput. Neurosci.* **11** 64
[50] Li X, Song D, Zhang P, Zhang Y, Hou Y and Hu B 2018 Exploring EEG features in cross-subject emotion recognition *Front. Neurosci.* **12** 162
[51] Song T, Zheng W, Song P and Cui Z 2020 EEG emotion recognition using dynamical graph convolutional neural networks *IEEE Trans. Affective Comput.* **11** 532–41
[52] Hospedales T, Antoniou A, Micaelli P and Storkey A 2020 Meta-learning in neural networks: a survey (arXiv:2004.05439)
[53] Thrun S and Pratt L 1998 Learning to learn: introduction and overview *Learning to Learn* (Berlin: Springer) pp 3–17
[54] An S, Kim S, Chikontwe P and Park S H 2020 Few-shot relation learning with attention for EEG-based motor imagery classification (arXiv:2003.01300)
[55] Duan T, Chauhan M, Shaikh M A, Chu J and Srihari S 2020 Ultra efficient transfer learning with meta update for cross subject EEG classification (arXiv:2003.06113)
[56] Wu D, Lawhern V J, Gordon S, Lance B J and Lin C T 2016 Spectral meta-learner for regression (SMLR) model aggregation: Towards calibrationless brain-computer interface (BCI) *2016 IEEE Int. Conf. on Systems, Man and Cybernetics (SMC)* (IEEE) pp 000743–9
[57] Denevi G, Ciliberto C, Stamos D and Pontil M 2018 Learning to learn around a common mean *Advances in Neural Information Processing Systems* vol 31 (NIPS Proceedings)
[58] Kuzborskij I and Orabona F 2017 Fast rates by transferring from auxiliary hypotheses *Mach. Learn.* **106** 171–95
[59] Denevi G, Ciliberto C, Grazzi R and Pontil M 2019 Learning-to-learn stochastic gradient descent with biased regularization *Proc. 36th Int. Conf. on Machine Learning* pp 1566–75 (arXiv:1903.10399)
[60] Denevi G, Ciliberto C, Stamos D and Pontil M 2018 Incremental learning-to-learn with statistical guarantees (arXiv:1803.08089)
[61] Alquier P *et al* 2017 Regret bounds for lifelong learning *Artificial Intelligence and Statistics* (PMLR) pp 261–9 (arXiv:1610.08628)
[62] Boldt A, Schiffer A M, Waszak F and Yeung N 2019 Confidence predictions affect performance confidence and neural preparation in perceptual decision making *Sci. Rep.* **9** 1–17
[63] Zheng W L, Gao K, Li G, Liu W, Liu C, Liu J Q, Wang G and Lu B L 2020 Vigilance estimation using a wearable EOG device in real driving environment *IEEE Trans. Intell. Transp. Syst.* **21** 170–84
[64] Zhu X, Zheng W L, Lu B L, Chen X, Chen S and Wang C 2014 EOG-based drowsiness detection using convolutional neural networks *2014 Int. Joint Conf. on Neural Networks (IJCNN)* (IEEE) pp 128–34
[65] Huo X Q, Zheng W L and Lu B L 2016 Driving fatigue detection with fusion of EEG and forehead EOG *2016 Int. Joint Conf. on Neural Networks (IJCNN)* (IEEE) pp 897–904
[66] Zheng W L and Lu B L 2017 A multimodal approach to estimating vigilance using EEG and forehead EOG *J. Neural Eng.* **14** 026017
[67] Croft R J and Barry R J 2000 Removal of ocular artifact from the EEG: a review *Clin. Neurophysiol.* **30** 5–19
[68] Nolan H, Whelan R and Reilly R B 2010 FASTER: fully automated statistical thresholding for EEG artifact rejection *J. Neurosci. Methods* **192** 152–62
[69] Pan S J, Tsang I W, Kwok J T and Yang Q 2011 Domain adaptation via transfer component analysis *IEEE Trans. Neural Netw.* **22** 199–210
[70] Tan C, Sun F and Zhang W 2018 Deep transfer learning for EEG-based brain computer interface *2018 IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)* (IEEE) pp 916–20
[71] Özdenizci O, Wang Y, Koike-Akino T and Erdoğmuş D 2019 Transfer learning in brain-computer interfaces with adversarial variational autoencoders *2019 9th Int. IEEE/EMBS Conf. on Neural Engineering (NER)* (IEEE) pp 207–10 (arXiv:1812.06857)
[72] Wu D, Xu Y and Lu B L 2020 Transfer learning for EEG-based brain-computer interfaces: a review of progress made since 2016 *IEEE Trans. Cogn. Dev. Syst.* **14** 4–19
[73] Li J, Qiu S, Du C, Wang Y and He H 2020 Domain adaptation for EEG emotion recognition based on latent representation similarity *IEEE Trans. Cogn. Dev. Syst.* **12** 344–53