

# End-to-end learnable EEG channel selection with deep neural networks

Thomas Strypsteen

Alexander Bertrand

**Abstract**—Many electroencephalography (EEG) applications rely on channel selection methods to remove the least informative channels, e.g., to reduce the amount of electrodes to be mounted, to decrease the computational load, or to reduce overfitting effects and improve performance. Wrapper-based channel selection methods aim to match the channel selection step to the target model, yet they require to re-train the model multiple times on different candidate channel subsets, which often leads to an unacceptably high computational cost, especially when said model is a (deep) neural network. To alleviate this, we propose a framework to embed the EEG channel selection in the neural network itself to jointly learn the network weights and optimal channels in an end-to-end manner by traditional backpropagation algorithms. We deal with the discrete nature of this new optimization problem by employing continuous relaxations of the discrete channel selection parameters based on the Gumbel-softmax trick. We also propose a regularization method that discourages selecting channels more than once. This generic approach is evaluated on two different EEG tasks: motor imagery brain-computer interfaces and auditory attention decoding. The results demonstrate that our framework is generally applicable, while being competitive with state-of-the-art EEG channel selection methods, tailored to these tasks.

**Index Terms**—Channel selection, deep neural networks, EEG, wireless EEG sensor network

## I. INTRODUCTION

Electroencephalography (EEG) is a widely used neuro-monitoring technique that measures the brain’s electrical activity in a noninvasive way. Its applications are numerous, including detection of epileptic seizures [1], monitoring sleeping patterns [2], studying brain disorders after injuries [3], providing communication means for motor-impaired patients through brain-computer interfaces (BCI’s) [4] and many more. However, acquiring these EEG signals typically involves wearing bulky, heavy EEG caps containing a large amount of electrodes with conductive gel, resulting in an uncomfortable user experience and restricting the monitoring to hospital or lab settings. These limitations of classical EEG have led to a growing desire for ambulatory EEG, allowing for continuous neuromonitoring in daily life [5].

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 802895). The authors also acknowledge the financial support of the FWO (Research Foundation Flanders) for project G.0A49.18N, and the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” programme.

T. Strypsteen and A. Bertrand are with KU Leuven, Department of Electrical Engineering (ESAT), STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics and with Leuven.AI - KU Leuven institute for AI, Kasteelpark Arenberg 10, B-3001 Leuven, Belgium (e-mail: thomas.strypsteen@kuleuven.be, alexander.bertrand@kuleuven.be).

This shift to mobile applications means the EEG cap is replaced by a number of lightweight, concealable mini-EEG devices, possibly organized in a wireless EEG sensor network (WESN) [6], [7], [8]. Since recording and transmitting all possible channels would incur enormous energy costs, selecting an optimal subset of channels to perform the given task constitutes a crucial step in this wireless setting. Even in more traditional EEG settings, reducing the number of channels offers numerous advantages: it reduces the setup time in clinical settings, helps prevent overfitting effects, decreases the computational load and improves interpretability of the model by removing uninformative channels.

The last few years, deep learning models have emerged as a popular EEG analysis tool [9]. For several applications, it has been shown that replacing the classical signal processing approaches with deep neural networks (DNNs) can substantially improve performance [1,4,10]. In this paper, we focus on the EEG channel selection problem in DNNs. A major problem that arises when performing channel selection - which can be viewed as a grouped feature selection - for neural networks is that many popular feature selection techniques are wrapper approaches. This means that a certain heuristic search is performed on the space of possible feature subsets, training a model on each of these candidate subsets and selecting the one where the model’s performance is optimal. However, training a neural network is computationally a lot more demanding than traditional machine learning algorithms, rendering these procedures far too time-consuming for practical usage.

In this work, we propose a procedure to learn the optimal channel subset in an end-to-end manner, training the network weights while simultaneously learning to select the optimal channels. To this end, we will extend the neural network with a layer of selection neurons [11], that each learn to select one of its inputs rather than learn weights to linearly combine them. To learn the discrete parameters of these selection neurons through standard backpropagation, we will employ the Gumbel-softmax trick [12] to make continuous approximations of the discrete parameters involved. While this approach has been successfully applied to high-dimensional feature spaces [11,13], we demonstrate that it often leads to the selection of duplicate channels in the case of EEG channel selection. This is because the selection neurons act independently and are not ‘aware’ of each others’ selection.

In high-dimensional data sets as in [13], the probability of such a collision is negligible. However, EEG channel selection typically involves selecting from a relatively small pool of input channels. This means that the probability of different selection neurons selecting the same input channel is no longer negligible and should be addressed. To this end, we introduce a novel regularization function that couples the different selection neurons and encourages them to select distinct channels.

To demonstrate the general applicability of this method, we study its performance on two different EEG tasks: motor imagery and auditory attention decoding (AAD). We demonstrate that the proposed end-to-end learnable EEG channel selector, despite its generic nature, achieves competitive results (better or at least equally good) compared to state-of-the-art channel selection methods. Furthermore, the latter are often tailored or constrained to specific tasks or input feature types, while the Gumbel-softmax channel selector is widely applicable. It can be used for both regression and classification tasks, and can be placed behind any type of input layer, be it raw EEG time series or pre-computed per-channel features. We provide a Pytorch [14] implementation for the interested reader who wants to use our method <sup>1</sup>.

The remainder of this manuscript is organized as follows. Section II presents related work in the field of channel selection and discrete optimization in neural networks. Section III describes the proposed channel selection layer in detail. In section IV, we discuss the tasks used to validate this method along with the baseline algorithms we compare the proposed method with, before presenting our experimental results in section V. We conclude this paper with a discussion of these results in section VI.

## II. RELATED WORK

### A. Channel selection

The problem we aim to solve here is selecting the optimal subset  $K$  channels to solve a given EEG task. This is inherently a grouped feature selection problem, with each channel containing multiple features to be selected together. In general, this can be solved with a filter, wrapper or embedded approach.

In filter approaches, the relevance of each feature in predicting the correct class label is determined using some model-independent criterion, such as, e.g., mutual information (MI). In the channel selection case however, this is complicated by the fact that each channel contains multiple features and thus requires multi-dimensional entropy estimators to determine the relevance of a single channel. One solution that has been successfully applied to EEG channel selection is applying Independent Component Analysis (ICA) to the features of each channel to transform them to new,

independent features. After this transformation, their joint entropy can be estimated as the sum of their marginal entropies [15]. While this approach solves the problem of multi-dimensional entropy estimation, it still requires the computation of handcrafted features for the specific task at hand and is therefore not directly applicable in cases where raw EEG data is used directly as an input for a DNN<sup>2</sup>.

An important drawback of filter approaches is that the channel selection is performed on a surrogate metric that is not matched to the target model that will act as the eventual classifier. To alleviate this drawback, wrapper approaches aim to find the optimal feature subset by performing a heuristic search through the space of possible subsets and evaluating the target model's performance on multiple candidate sets. For instance, Qiu et al. propose a sequential forward floating selection approach (SFFS) in combination with Common Spatial Pattern (CSP) feature extraction and Support Vector Machine (SVM) classification to select the optimal channel subset for a motor imagery BCI problem [16]. In each iteration of the algorithm, the channel that would improve the model's crossvalidation accuracy the most is added to the selected subset, requiring the CSP bank and SVMs to be trained multiple times. While wrapper approaches generally lead to better selections, they are expensive from a computational point of view, due to the numerous re-training of the model on a large number of feature subsets. Applying such methods when the model to be trained is a DNN would be far too time-consuming to be used in practice and will not be further discussed here.

Finally, embedded approaches jointly train the model and select the optimal features, typically by adding a regularization term to the training objective. A well-known example is the LASSO [17], which induces sparsity in a model's weights by penalizing their  $L_1$ -norm. LASSO can be used to perform feature selection in DNNs by driving all the weights associated with uninformative features to zero [18] and can also be extended to select groups of features together [19]. Feature selection can be even more explicitly modeled by using continuous relaxations of  $L_0$ -regularization, learning each input neuron to either be 'on' or 'off' [20]. The downside of these sparsity-inducing methods is that the amount of features selected depends on the weight of the regularization in the objective function. This means we cannot supply the model with the number of features to be selected *a priori*, which makes them unfit for the given subset selection problem. If a target number of channels is to be selected, such methods have to be retrained multiple times on a trial-and-error basis. Similar to wrapper methods, this makes such methods often too time-consuming, in particular for training of DNNs.

<sup>2</sup>Note that one of the main advantages of DNNs is the fact that the network can jointly learn a classifier and a proper feature embedding in an end-to-end fashion, where only raw data is provided as an input of the network.

<sup>1</sup><https://github.com/Strypsteen/Gumbel-Channel-Selection>

### B. Discrete optimization in neural networks

Directly solving the channel selection problem is a discrete optimization problem, while learning in neural networks is based on backpropagation, a procedure that requires a differentiable loss function and by extension, continuous parameters to be learned. However, it is possible to integrate discrete parameters in this framework by using categorical reparametrization with the Gumbel-softmax trick [12]. This process encodes the discrete parameters as a discrete distribution to be learned. This discrete distribution is approximated by some continuous relaxation, for instance the concrete distribution [21]. The parameters of this distribution are then learned through standard backpropagation by employing the reparametrization trick [22]. A general overview of this class of methods can be found in [23]. This framework has been used by Abid et al. to build a concrete selector layer that learns to select  $K$  features from its input [11]. By stacking this layer on top of an autoencoder, it is possible to learn the subset of features that allow for optimal reconstruction of the complete feature set. In contrast to the embedded methods described above, this model explicitly models the amount of features to be selected and will serve as the basis for our channel selection method, as described in the next section.

### III. PROPOSED METHOD

#### A. Channel selection layer

Let  $\mathcal{D} = \{(X^{(1)}, y^{(1)}), (X^{(2)}, y^{(2)}), \dots, (X^{(M)}, y^{(M)})\}$  be a dataset of  $M$  EEG samples  $X^{(i)}$  with class labels  $y^{(i)}$ . Each  $X \in \mathbb{R}^{N \times F}$  contains  $N$  channels and  $F$  features per channel. These features could be anything ranging from the raw time samples to, e.g., power features in certain frequency bands. Let  $S$  indicate a subset of  $K$  channels and  $X_S \in \mathbb{R}^{K \times F}$  the reduced EEG samples containing only the rows of  $X$  corresponding to the channels in  $S$ . Also assume we have a neural network model  $f_\theta(X_S)$ , where  $\theta$  contains all the learnable parameters of the model. Our goal is then to learn the optimal  $S^*$  and  $\theta^*$  such that

$$S^*, \theta^* = \arg \min_{S, \theta} \mathcal{L}(f_\theta(X_S), y) \quad (1)$$

with  $\mathcal{L}(p, y)$  any loss function between the predicted label  $p$  and the ground truth  $y$ .

To accomplish this, we extend the neural network model with a channel selection layer. We propose the use of a so-called concrete selector layer [11], in which  $K$  selection neurons are stacked on top of each other, one for each channel to be selected (see Fig. 1). Each of these selection neurons takes all channels as input and produces a single output channel. Each selection neuron is parametrized by a learnable vector  $\alpha_k \in \mathbb{R}_{>0}^N$ . When being fed a sample  $X$ , each selection

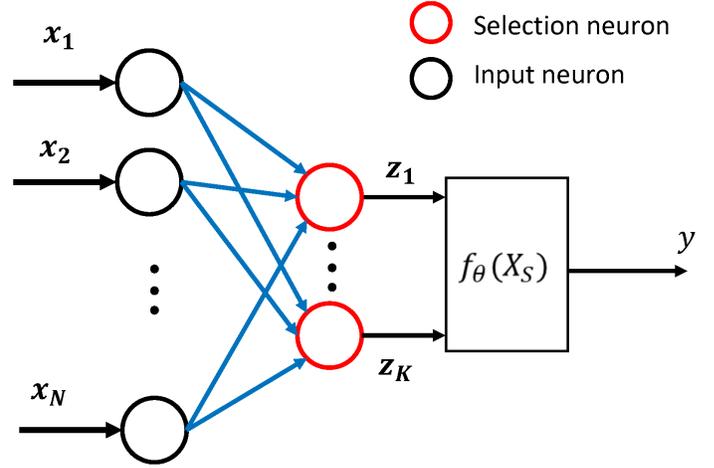


Fig. 1. Illustration of the channel selection layer:  $x_i$  indicates the feature derived from channel  $i$ . During training, the output of each selection neuron is given by  $z_k = w_k^T X$ , with  $w_k \sim \text{Concrete}(\alpha_k, \beta)$ .

neuron samples a weight vector  $w_k \in \mathbb{R}^N$  from the concrete distribution [21]:

$$w_{ik} = \frac{\exp((\log \alpha_{ik} + G_{ik})/\beta)}{\sum_{j=1}^N \exp((\log \alpha_{jk} + G_{jk})/\beta)} \quad (2)$$

with  $G_{ik}$  independent and identically distributed (i.i.d) samples from the Gumbel distribution [24] and  $\beta \in (0, +\infty)$  the temperature of the concrete distribution. Thus, in a manner similar to Dropout [25], a different weight vector is sampled for each observation during training. Each neurons then computes its output channel as  $z_k = w_k^T X$ .

Equation (2) can be viewed as a softmax operation, which produces weight vectors whose elements sum to one as continuous relaxations of one-hot vectors. The temperature  $\beta$  controls the extent of this relaxation. It can be shown that, as  $\beta$  approaches 0, the distribution will become more discrete, the sampled weights will converge to one-hot vectors and the neuron will go from linearly combining to *selecting* a certain input channel [21]. The probability  $p_{nk}$  of neuron  $k$  selecting a certain channel  $n$  is then given by

$$p_{nk} = \frac{\alpha_{nk}}{\sum_{j=1}^N \alpha_{jk}}. \quad (3)$$

During training, the temperature is decreased by an exponentially decreasing curve as in [11], i.e.  $\beta(t) = \beta_s(\beta_e/\beta_s)^{t/T}$  with  $\beta(t)$  the temperature at epoch  $t$ ,  $\beta_s$  and  $\beta_e$  the start and end temperatures and  $T$  the number of epochs. This allows the network to explore various combinations of channels in the beginning of the training while forcing it to a selection operation by the end of training. At the same time, as the probability of sampling from a certain channel increases, its gradient will start to dominate the gradient of the batches, causing a positive feedback effect that drives the probability vectors  $p_k$  to (approximately) one-hot vectors as well. This means that the uncertainty of each selection neuron progressively decreases

until it almost only selects one specific channel. At test time, the stochastic nature of the network is dropped entirely and the continuous softmax is replaced by a discrete argmax. This means the weights of the neurons are replaced with fixed one-hot vectors.

$$w_{ik} = \begin{cases} 1, & \text{if } i = \arg \max_j \alpha_{jk} \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

### B. Duplicate channel selection

The downside of this construction is that, since each neuron samples its weights independently, it is possible for multiple selection neurons to select the same channel, introducing redundancy in the network’s input. Originally, the Gumbel-softmax selection layer was proposed for dimensionality reduction in dense auto-encoder architectures [11] or high-dimensional input features [13], where the probability of having duplicate selections is negligible. However, in the case of EEG channel selection, the probability of having a ‘collision’ between two selection neurons is high, as we will show in Section IV. We will refer to this problem from this point on as the duplicate channel selection problem. A straightforward (yet naive) fix for this problem would be to replace all duplicate channels with different ones, and to retrain the network weights for the newly added channels. The replacement could be done arbitrarily or one can let the network select the missing channels in a second training step from a reduced candidate pool (and repeating this process until no duplicate channels are selected anymore). However, such ad hoc fixes require additional training steps (at least one), which can be time-consuming in a context of DNNs. Furthermore, the occurrence of duplicate channels usually implies that the training process is unable to escape from a local minimum, possibly leading to a suboptimal initial channel selection, which is then carried over to all subsequent iterations. To avoid such local optima and to avoid expensive extra training phases, duplicate channels should ideally be avoided *during* training instead of using post hoc (and ad hoc) fixes. To this end, we propose a regularization function that encourages the selection neurons to learn distinct channel selections within a single training phase.

To avoid this issue, we instead introduce a regularization function that encourages the selection neurons to learn distinct selections during training. Consider the selection matrix  $P$ , constructed by normalizing the parameter vector  $\alpha_k$  of each selection neuron and putting them in the columns of  $P$ , such that the entry in the  $n$ -th row and  $k$ -th column is defined as  $p_{nk} = \frac{\alpha_{nk}}{\sum_{j=1}^N \alpha_{jk}}$ . Thus, column  $k$  of  $P$  represents the probability distribution over the input channels that neuron  $k$  will select as the temperature goes to 0. By the end of the training procedure, the columns of this matrix approximate one-hot vectors, with the position of the 1 indicating which channel is selected for this selection neuron. It can be observed that choosing  $K$  unique channels corresponds to the *rows* of  $P$  not containing more than a single 1-entry. During training

however, the entries of the selection matrix are still continuous probabilities, so we encourage the selection neurons to pick distinct channels by penalizing the sum of the selection matrix’s rows:

$$\mathcal{L}(P) = \lambda \sum_{n=1}^N \text{Relu}(\sum_{k=1}^K p_{nk} - \tau) \quad (5)$$

with  $\lambda$  the weight of the regularization loss,  $\text{Relu}$  the rectified linear unit operation  $f(x) = \max(0, x)$  and  $\tau$  a threshold parameter. This regularization function only applies a penalization when the sum of a channel’s probabilities across the selection neurons exceeds a threshold  $\tau$ . Like the temperature  $\beta$  of the concrete distribution, this threshold is decayed exponentially during training, becoming more stringent as the distribution becomes more discrete and the selection matrix becomes more one-hot. By the end of the training, the threshold approaches 1 and duplicate channel selection is explicitly penalized. The weight  $\lambda$  on the other hand controls the strength of the regularization: a higher  $\lambda$  prevents more duplicate channels, but raising it too high might result in the network ignoring the original loss (1), thereby pushing the selection process towards distinct channels, regardless of how informative they are. Important to note is that the regularization function is constructed in such a way that it does not interfere in the training at all when no duplicate channels start to arise, avoiding the introduction of unnecessary bias in the network.

## IV. MATERIALS AND METHODS

To demonstrate the generic nature of this method to any EEG channel selection task, we validate our method on two different EEG-BCI paradigms: motor imagery and auditory attention decoding.

### A. Motor Imagery

The first task we discuss is motor imagery, a popular paradigm in the field of BCI. The goal is to decode EEG signals in motorsensory brain areas associated with imagined body movement. For our experiments, we make use of the High Gamma Dataset [26]. This dataset contains 128-channel EEG recordings from 14 subjects, with about 1000 trials of executed movement per subject following a visual cue. These movements belong to one of four classes: left hand, right hand, feet and rest. Similar to [26], we only employ the 44 channels covering the motor cortex in our experiments. We employ the same preprocessing procedure as [26], that is, resampling the data to 250 Hz, highpass-filtering the data at 4 Hz, standardizing the data for each electrode and epoching the data in 4.5 second segments, taking the 0.5 seconds before each visual cue and the 4 seconds after. We then classify the filtered time series with the parallel multiscale filter bank convolutional neural network (MSFBCNN) proposed in [27]. We report mean test accuracy across the subjects on a held-out test set of about 180 trials per subject.

For motor imagery, we compare the performance of the proposed Gumbel-softmax method with the mutual information (MI) based channel selection approach described in [15]. In short, it can be described as follows: the joint MI each (per-channel) block of features and the class label is computed and the channel with the maximal MI is added to the set of selected channels. Then, the joint MI of the currently selected set with each remaining channel is computed and the channel that maximizes this value is added to the selected set. This process is repeated until the desired number of channels are selected. Since we first need to craft informative features to perform this procedure, we compute the spectral power of 9 4-Hz wide frequency bands between 4 and 40 Hz for each channel, which were previously employed in filter bank CSPs for motor imagery [28]. Note that the feature extraction is only used here to inform the MI channel selection procedure, whereas the classification by the MSFBCNN is performed directly on the filtered time series.

### B. Auditory match-mismatch

A second, entirely different task we discuss here is speech decoding, more specifically, the auditory match-mismatch paradigm [29]. Given two candidate speech envelopes and an EEG recording, the goal is to classify which of the envelopes actually elicited the EEG response and which one is an ‘imposter’. We employ the dataset described in [30], containing speech stimuli and the corresponding 64-channel EEG recordings for 48 normal hearing subjects. For classification, we follow the approach of [31], using a dilated convolutional neural network (DCNN) model. Similar to the previous network, this model uses the raw EEG traces as inputs and therefore does not require a prior feature construction step. As before, we report mean test accuracy across the subjects on a held-out test set of 10% of the full recordings.

Unfortunately, using MI as a channel selection benchmark for this application is not straightforward, as MI requires a set of discriminative features to work on. For motor imagery, it is generally known that the power in certain spectral bands can indeed serve as a discriminative feature. In the field of AAD on the other hand, there is no clear set of EEG/speech features that perform the same role. Instead, we employ a greedy channel selection procedure based on the least-squares utility metric as described in [7], which is currently the state-of-the-art channel selection method in EEG-based speech decoding paradigms [32]. In this setting, a linear decoder is trained to reconstruct the matching speech stimulus from the EEG signal, which constitutes a least-squares (LS) regression problem. The utility metric of a channel is then defined as the increase in LS-cost when this channel is dropped from the regression and the regression parameters are re-optimized, which can be calculated very efficiently as shown in [33]. We can use this to select  $K$  channels by iteratively removing the channels with the lowest utility metric until there are  $K$  channels left.

### C. Training procedure

For both tasks, we build the network for Gumbel-softmax channel selection by inserting the channel selection layer between the input layer and the baseline networks as described above. Using the data of all subjects simultaneously, we jointly train the selection layer and the network weights using the Adam optimizer (with a learning rate of 0.001) [34]. During training, the temperature  $\beta$  is decayed from 10 to 0.1 and the regularization threshold  $\tau$  from 3 to 1.1. For the regularization weight  $\lambda$  we chose a value of 0.1, which worked well for our applications, where the supervised loss was typically between 1 and 0.1. For applications where the supervised loss is higher or lower than this,  $\lambda$  should be scaled accordingly. We track the convergence of the channel selection by analyzing the normalized entropy of the distribution of each selection neuron, computed as:

$$H(\alpha_{\mathbf{k}}) = -\frac{1}{\log N} \sum_{j=1}^N \alpha_{jk} \log(\alpha_{jk}). \quad (6)$$

When the mean entropy of all selection neurons drops below a predefined threshold (we took 0.05 in our experiments) or the maximum amount of epochs is reached (150 for the motor imagery and 50 for the auditory match-mismatch task), we consider the selection process to have converged and the training finished. At that point, the parameters of the channel selection layer are frozen and the layer operates in its deterministic selection mode for evaluation (see section III). In the motor imagery case, we also proceed to fine-tune the network weights with the data of each subject separately to construct subject-dependent decoders, with the selected channels remaining the same for all subjects. We performed 10 runs of each training session and report results across all runs through box plots.

## V. EXPERIMENTAL RESULTS

### A. Duplicate channel selection

We first study the effect of our proposed regularization to avoid duplicate channel selection on the motor imagery problem. Fig. 2 shows how many unique channels are selected by the selection neurons for the normal Gumbel-softmax method and the regularized version, while the corresponding test accuracy is shown in Fig. 3. It can be seen that as early as selecting 10 out of 44 channels, duplicate channels occur regularly, while adding the regularization term postpones this behaviour. Putting the two figures side by side, the onset of the performance gap between the base algorithm and the regularized version does indeed coincide with the occurrence of these duplicate channels, implying that, without the regularization, the algorithm appears to get stuck in a local minimum and converges to a suboptimal performance.

### B. Channel selection task

1) *Motor imagery task:* Fig. 4 compares the test accuracies reached by the regularized Gumbel-softmax and the MI algorithm. The Gumbel-softmax method generally

performs better than the MI algorithm, while having two main advantages over it. Firstly, the MI channel selection method is computationally expensive, performing multiple iterations of ICA’s of growing size to perform the channel selection and then still needing the network to be trained on the selected channel subset. Secondly, the MI algorithm requires the computation of handcrafted features since computing accurate entropies of the raw time series would pose a far too highdimensional problem. The Gumbel-softmax method on the other hand is able to perform the channel selection on the raw time series and jointly learns the network weights alongside the optimal selection. Also interesting to note is that at only 10 channels, the network’s performance differs only about 3% from the full 44-channel performance, showing that deep neural network models can still be powerful tools in the low-channel settings of mobile EEG.

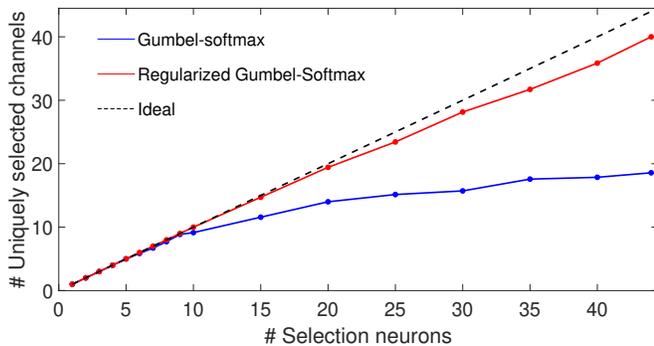


Fig. 2. Comparison of the amount of duplicate channels occurring in the Gumbel-softmax and regularized Gumbel-softmax method in function of the amount of channels to be selected. The displayed boxplots are computed over 10 runs.

2) *Auditory match-mismatch task*: The results of the auditory match-mismatch task are shown in figure 5. Here, we can make the same observations as in the motor imagery case. Firstly, the Gumbel-softmax in most cases performs better or at least as well as the utility metric despite the former being a more generic method. Secondly, high accuracies compared to the full channel baseline can already be achieved with a low amount of channels, with 10 channels already achieving an accuracy close to that of the full-channel baseline.

## VI. CONCLUSION AND FUTURE OUTLOOK

We have proposed the use of selection neurons based on concrete distributions to solve the EEG channel selection problem for neural network models. This method embeds the channel selection as part of the training of the model, dealing with its discrete nature by employing categorical reparametrization with Gumbel-softmax. We have demonstrated that directly applying this method to the task of EEG channel selection results in redundant selections containing duplicate channels. We addressed this issue by introducing a novel regularization function that encourages

the selection neurons to choose distinct channels, which was shown to increase the performance of the algorithm.

The performance of this method has been evaluated on two different EEG tasks, motor imagery and auditory match-mismatch. On both these tasks, the experimental results indicate that the Gumbel-softmax generally performed better or at least as well as ad hoc benchmarks tailored for these tasks: mutual information and greedy channel selection with the utility metric. Important to note here is that both these benchmark algorithms are not easily usable on the other task. MI requires the computation of task-specific, per-channel features to compute the mutual information with the class label with, features that are not readily available in the case of auditory match-mismatch. The utility metric requires the problem at hand to be formulated as a least-squares regression problem, which is not possible in the case of motor imagery. The Gumbel-softmax method on the other hand is a very general method and can be readily applied to any EEG regression or classification task, whether the inputs are pre-computed features or raw time series.

A second advantage of the Gumbel-softmax approach is that it is an embedded method that jointly learns the optimal channel selection and the weights of the neural network classifier/regressor model. In contrast, traditional filter methods require a separate channel selection step and a subsequent training of the model on the selected subset, and therefore the channel selection strategy is not necessarily matched to the target model that will eventually be used for classification. On the other hand, wrapper methods require the model to be trained multiple times, a computationally infeasible demand when dealing with neural networks.

Finally, the Gumbel-softmax method has a great plug-and-play value: applying it to an existing model simply requires putting a channel selection layer in front of it, given that the input layer of the neural network can be scaled to accommodate the lower input dimensionality of the required number of channels. Additionally, our use of a regularization function for distinct selections can be extended with additional constraints on the channels to be selected. For example, one possibility is selecting the channels that not only optimize performance, but also minimize the inter-electrode distance as much as possible, as is required in the design of miniaturized EEG sensor networks [7,32].

## ACKNOWLEDGEMENTS

We would like to thank Professor Tom Francart, Bernd Accou and Mohammad Jalilpour Monesi for contributing the dataset for the auditory match-mismatch task.

## REFERENCES

- [1] A. H. Ansari, P. J. Cherian, A. Caicedo, G. Naulaers, M. De Vos, and S. Van Huffel, “Neonatal seizure detection using deep convolutional neural networks,” *International journal of neural systems*, vol. 29, no. 04, p. 1850011, 2019.

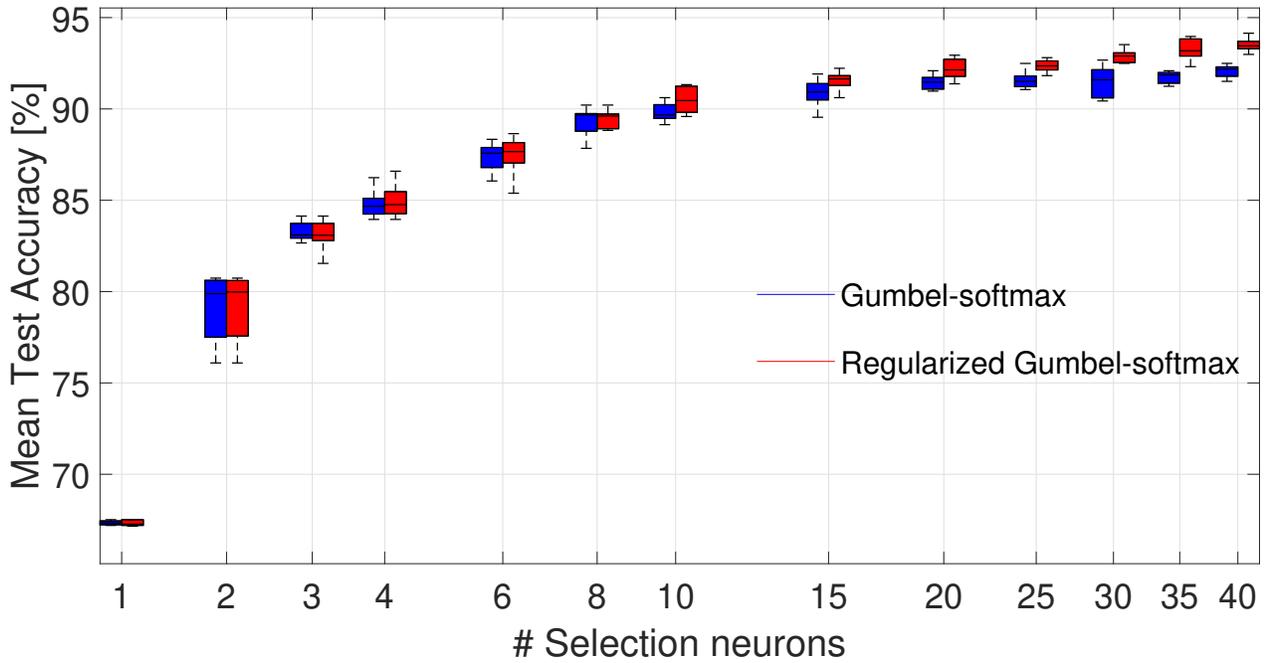


Fig. 3. Comparison of the Gumbel-softmax channel selection strategies on the motor imagery dataset. Mean test accuracies across the subjects are plotted as a function of the number of selection neurons. Note that duplicate channels can occur in the Gumbel-softmax methods, so the x-axis represents the *maximum* number of channels the algorithm can select. The displayed boxplots are computed over 10 runs.

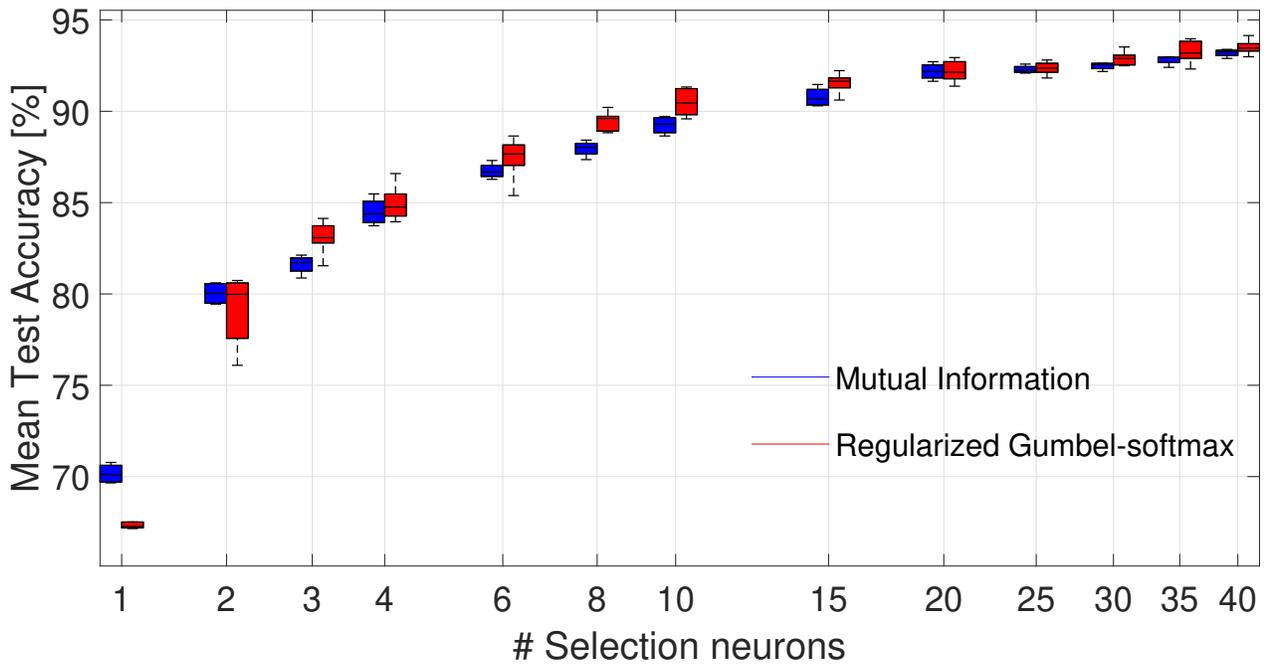


Fig. 4. Comparison of regularized Gumbel-softmax with the MI algorithm on the motor imagery dataset. Mean test accuracies across the subjects are plotted as a function of the number of selection neurons. The displayed boxplots are computed over 10 runs.

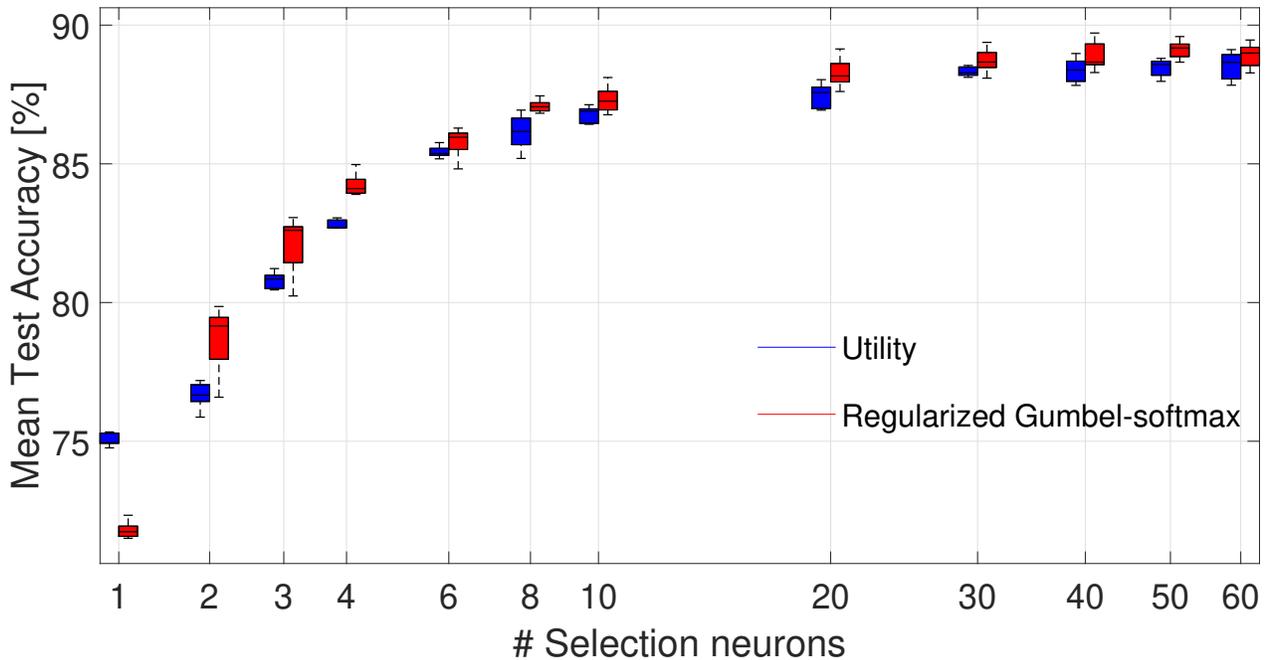


Fig. 5. Comparison of the channel selection strategies on the auditory match-mismatch dataset. Mean test accuracies across the subjects are plotted as a function the number of selection neurons. The displayed boxplots are computed over 10 runs.

- [2] O. De Wel, M. Lavanga, A. C. Dorado, K. Jansen, A. Dereymaeker, G. Naulaers, and S. Van Huffel, "Complexity analysis of neonatal EEG using multiscale entropy: applications in brain maturation and sleep stage classification," *Entropy*, vol. 19, no. 10, p. 516, 2017.
- [3] J. T. Giacino, J. J. Fins, S. Laureys, and N. D. Schiff, "Disorders of consciousness after acquired brain injury: the state of the science," *Nature Reviews Neurology*, vol. 10, no. 2, p. 99, 2014.
- [4] V. J. Lawhern, A. J. Solon, N. R. Waytowich, S. M. Gordon, C. P. Hung, and B. J. Lance, "EEGNet: a compact convolutional neural network for EEG-based brain-computer interfaces," *Journal of neural engineering*, vol. 15, no. 5, p. 056013, 2018.
- [5] A. J. Casson, D. C. Yates, S. J. Smith, J. S. Duncan, and E. Rodriguez-Villegas, "Wearable electroencephalography," *IEEE engineering in medicine and biology magazine*, vol. 29, no. 3, pp. 44–56, 2010.
- [6] A. Bertrand, "Distributed signal processing for wireless EEG sensor networks," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 23, no. 6, pp. 923–935, 2015.
- [7] A. M. Narayanan and A. Bertrand, "Analysis of miniaturization effects and channel selection strategies for EEG sensor networks with application to auditory attention detection," *IEEE Transactions on Biomedical Engineering*, vol. 67, no. 1, pp. 234–244, 2019.
- [8] T. Tang, L. Yan, J. H. Park, H. Wu, L. Zhang, H. Y. B. Lee, and J. Yoo, "34.6 EEG dust: A BCC-based wireless concurrent recording/transmitting concentric electrode," in *2020 IEEE International Solid-State Circuits Conference-(ISSCC)*. IEEE, 2020, pp. 516–518.
- [9] Y. Roy, H. Banville, I. Albuquerque, A. Gramfort, T. H. Falk, and J. Faubert, "Deep learning-based electroencephalography analysis: a systematic review," *Journal of neural engineering*, vol. 16, no. 5, p. 051001, 2019.
- [10] S. Vandecappelle, L. Deckers, N. Das, A. Ansari, A. Bertrand, and T. Francart, "EEG-based detection of the attended speaker and the locus of auditory attention with cnns," in *DeepLearn 2019, Date: 2019/07/22-2019/07/26, Location: Warsaw, PL*, 2019.
- [11] A. Abid, M. F. Balin, and J. Zou, "Concrete autoencoders for differentiable feature selection and reconstruction," *arXiv preprint arXiv:1901.09346*, 2019.
- [12] E. Jang, S. Gu, and B. Poole, "Categorical reparameterization with gumbel-softmax," *arXiv preprint arXiv:1611.01144*, 2016.
- [13] D. Singh and M. Yamada, "Fsnet: Feature selection network on high-dimensional biological data," *arXiv preprint arXiv:2001.08322*, 2020.
- [14] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in neural information processing systems*, 2019, pp. 8026–8037.
- [15] T. Lan, D. Erdogmus, A. Adami, M. Pavel, and S. Mathan, "Salient EEG channel selection in brain computer interfaces by mutual information maximization," in *2005 IEEE Engineering in Medicine and Biology 27th Annual Conference*. IEEE, 2006, pp. 7064–7067.
- [16] Z. Qiu, J. Jin, H.-K. Lam, Y. Zhang, X. Wang, and A. Cichocki, "Improved sffs method for channel selection in motor imagery based bci," *Neurocomputing*, vol. 207, pp. 519–527, 2016.
- [17] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, 1996.
- [18] S. Scardapane, D. Comminiello, A. Hussain, and A. Uncini, "Group sparse regularization for deep neural networks," *Neurocomputing*, vol. 241, pp. 81–89, 2017.
- [19] L. Zhao, Q. Hu, and W. Wang, "Heterogeneous feature selection with multi-modal deep neural networks and sparse group lasso," *IEEE Transactions on Multimedia*, vol. 17, no. 11, pp. 1936–1948, 2015.
- [20] C. Louizos, M. Welling, and D. P. Kingma, "Learning sparse neural networks through  $l_0$  regularization," *arXiv preprint arXiv:1712.01312*, 2017.
- [21] C. J. Maddison, A. Mnih, and Y. W. Teh, "The concrete distribution: A continuous relaxation of discrete random variables," *arXiv preprint arXiv:1611.00712*, 2016.
- [22] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [23] M. B. Paulus, D. Choi, D. Tarlow, A. Krause, and C. J. Maddison, "Gradient estimation with stochastic softmax tricks," *arXiv preprint arXiv:2006.08063*, 2020.
- [24] E. J. Gumbel, *Statistical theory of extreme values and some practical applications: a series of lectures*. US Government Printing Office, 1948, vol. 33.
- [25] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from over-

- fitting,” *The journal of machine learning research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [26] R. T. Schirrmeister, J. T. Springenberg, L. D. J. Fiederer, M. Glasstetter, K. Eggensperger, M. Tangermann, F. Hutter, W. Burgard, and T. Ball, “Deep learning with convolutional neural networks for EEG decoding and visualization,” *Human brain mapping*, vol. 38, no. 11, pp. 5391–5420, 2017.
- [27] H. Wu, F. Li, Y. Li, B. Fu, G. Shi, M. Dong, and Y. Niu, “A parallel multiscale filter bank convolutional neural networks for motor imagery EEG classification,” *Frontiers in Neuroscience*, vol. 13, p. 1275, 2019.
- [28] K. K. Ang, Z. Y. Chin, H. Zhang, and C. Guan, “Filter bank common spatial pattern (fbcs) in brain-computer interface,” in *2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*. IEEE, 2008, pp. 2390–2397.
- [29] D. D. Wong, G. M. Di Liberto, and A. de Cheveigné, “Accurate modeling of brain responses to speech,” *bioRxiv*, p. 509307, 2018.
- [30] M. J. Monesi, B. Accou, J. Montoya-Martinez, T. Francart, and H. Van Hamme, “An LSTM based architecture to relate speech stimulus to EEG,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 941–945.
- [31] B. Accou, M. J. Monesi, J. Montoya-Martinez, H. Van Hamme, and T. Francart, “Modeling the relationship between acoustic stimulus and EEG with a dilated convolutional neural network,” in *Proceedings of the 28th European Signal Processing Conference, EUSIPCO 2020*. IEEE, 2020, pp. 1175–1179.
- [32] A. M. Narayanan, P. Patrinos, and A. Bertrand, “Optimal versus approximate channel selection methods for EEG decoding with application to topology-constrained neuro-sensor networks,” *IEEE transactions on neural systems and rehabilitation engineering: a publication of the IEEE Engineering in Medicine and Biology Society*.
- [33] A. Bertrand, “Utility metrics for assessment and subset selection of input variables for linear estimation [tips & tricks],” *IEEE Signal Processing Magazine*, vol. 35, no. 6, pp. 93–99, 2018.
- [34] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.