Temporally Regularized Filters for Common Spatial Patterns by Preserving Locally Linear Structure of EEG Trials

Minmin Cheng[1], Haixian Wang[1,⋆], Zuhong Lu[1], and Deji Lu[2]

[1] Key Lab. of Child Development and Learning Science of Ministry of Education, Research Center for Learning Science, Southeast University, Nanjing, Jiangsu 210096, China
[2] Medical Electronics Lab. , Southeast University, Nanjing, Jiangsu 210096, China
{mmcheng,hxwang,zhlu}@seu.edu.cn,deji.lu@gmail.com

Abstract. Common spatial patterns (CSP) is a commonly used method of feature extraction for motor imagery–based brain computer interfaces (BCI). However, its performance is limited when subjects have small training samples or signals are very noisy. In this paper, we propose a new regularized CSP: temporally regularized common spatial patterns (TRCSP), which is an extension of the conventional CSP by preserving locally linear structure. The proposed method and CSP are tested on data sets from BCI competitions. Experimental results show that the TRCSP achieves higher average accuracy for most of the subjects and some of them are up to 10%. Furthermore, the results also show that the TRCSP is particularly effective in the small–sample data sets.

Keywords: brain–computer interfaces (BCI), common spatial patterns (CSP), locally linear structure, regularization.

1 Introduction

Brain computer interfaces (BCI) have emerged as a promising way of non-muscular communication with external world for severely paralyzed persons [1]. Electroencephalogram (EEG)–based BCI transfers intents of an individual, reflected in distinguishable EEG signals directly, into control commands of an assistive device. The successful decoding of the mental tasks heavily relies on a robust classification of the EEG signals. Among the plenty of decoding methods [2], common spatial patterns (CSP) is a widely used feature extraction method that can learn spatial filters maximizing the discriminability of two classes. Its effectiveness has been demonstrated by the BCI competitions [3], [4].

Despite its popularity and efficiency, CSP is also known to be highly sensitive to noise and outliers [5]. Mathematically, CSP is formulated as the simultaneous diagonalization of two covariance matrices. There is an inherent drawback for the estimation of covariance matrices in using the conventional strategy. Specifically,

⋆ Corresponding author.

C.K. Loo et al. (Eds.): ICONIP 2014, Part III, LNCS 8836, pp. 167–174, 2014.

CSP does not take the temporal structure information of EEG time courses into account in the estimation of covariance matrices. In other words, CSP is a time-independent global method, and the temporal information is completely ignored.

In this paper, we propose a temporally regularized CSP (TRCSP), which incorporates the temporal structure information into the CSP learning process under the umbrella of regularization [6]. The temporal structure of EEG trials is characterized by using local linear embedding (LLE) [7]. Considering the advantage of LLE in successful discovery of manifold structure in machine learning, we aim to capture the locally linear structure of EEG trials with the LLE–based regularization. It is expected that such a prior information would help finding discriminative spatial filters, even with noisy EEG signals or small number of training samples, since the locally linear structure explicitly considers the temporal manifold behind the generation of EEG signals.

The framework of this paper is arranged as follows. Section 2 describes the conventional CSP algorithm and the proposed TRCSP algorithm. Section 3 gives details about the EEG data sets used for evaluation. Then the comparison results of the two methods are presented in Section 4. And finally Section 5 concludes the paper.

2 Methods

2.1 Common Spatial Patterns

Common spatial patterns(CSP) uses a linear transform to project multi–channel EEG data points into a low–dimensional spatial subspace with a projection matrix, of which each row consists of weights for channels. This transformation is to maximize the variance of band–pass filtered EEG signals of one class while minimizing the variance of EEG signals of the other class. Let $\boldsymbol{X}^i = \left\{\boldsymbol{x}_l^i \in R^d | l = 1, 2, \cdots, s\right\} (i = 1, 2, \cdots, n_x)$ be the EEG trials of one class, and $\boldsymbol{Y}^j = \left\{\boldsymbol{y}_l^j \in R^d | l = 1, 2, \cdots, s\right\} (j = 1, 2, \cdots, n_y)$ another class, where d denotes the number of channels, s is the number of samples within a trial, and n_x and n_y are the numbers of trials corresponding to the two classes. The trial segments are assumed to be already band–pass filtered, centered and scaled. The spatial covariance matrices of the two classes are calculated as

$$\overline{C_x} = \frac{1}{n_x} \sum_{i=1}^{n_x} \frac{\boldsymbol{X}^i {\boldsymbol{X}^i}^T}{tr\left(\boldsymbol{X}^i {\boldsymbol{X}^i}^T\right)} \qquad \overline{C_y} = \frac{1}{n_y} \sum_{j=1}^{n_y} \frac{\boldsymbol{Y}^j {\boldsymbol{Y}^j}^T}{tr\left(\boldsymbol{Y}^j {\boldsymbol{Y}^j}^T\right)} \tag{1}$$

where T represents the transpose operator, tr is the trace operator that sums up the diagonal entries of a matrix. The CSP approach aims to find a spatial filter $\boldsymbol{\omega} \in R^d$ to extract discriminative features. Mathematically, the spatial filter of CSP is formulated by maximizing (or minimizing) the criterion[8], [9]

$$J(\boldsymbol{\omega}) = \frac{\boldsymbol{\omega}^T \overline{C_x} \boldsymbol{\omega}}{\boldsymbol{\omega}^T \overline{C_y} \boldsymbol{\omega}} \tag{2}$$

The spatial filter is solved by the generalized eigenvalue equation

$$\overline{C_x}\boldsymbol{\omega} = \lambda\overline{C_y}\boldsymbol{\omega} \tag{3}$$

The few eigenvectors associated with eigenvalues from two ends of the eigenvalue spectrum are employed as spatial filters. The variances (possibly after a log–transformation) of the spatially filtered EEG data points are used as features for the purpose of classification.

2.2 Temporally Regularized Common Spatial Patterns

In this subsection we formulate the proposed TRCSP algorithm, which seeks to include temporal structure information into the learning process of the CSP. The EEG samples within a trial are actually a time course of signals. The temporally close samples usually correlated when recording a task–cued brain activity. It is beneficial to make use of the intrinsically temporal correlation to provide supplementary information and then regularize the computation of spatial filters. In other words, we try to keep the intrinsically temporal structure of EEG trials during the CSP filtering.

The temporal structure of EEG trials is captured by using LLE, which is well developed in the field of machine learning and has shown effective in manifold modeling. The basic idea is that we utilize LLE to consider temporally local relationship of EEG samples within the time course of EEG epochs. The relationship is expressed in terms of locally linear representation. Mathematically, LLE models each sample as a linear combination of its k nearest neighbors, and try to preserve this locally linear relationship in a transferred low–dimensional space. Different from the conventional LLE, in which the k nearest neighbors are identified with respect to Euclidean distance, we choose the k nearest neighbor EEG samples in terms of time points since we are interested in the temporal structure information of EEG time course. The reconstruction error is then measured by the cost function

$$\varepsilon(\boldsymbol{S}) = \sum_{l=1}^{s} \|\boldsymbol{x}_l - \sum_{m=1}^{s} S_{lm}\boldsymbol{x}_m\|^2 \tag{4}$$

where $\boldsymbol{S}$ is a matrix with real entries denoting representational weights. The weights S_{lm} summarize the contribution of the m th sample to the reconstruction of the l th sample in terms of linear representation. To compute the weights S_{lm}, we minimize the cost function subject to two constraints: (a) Each sample $\boldsymbol{x}_l$ is reconstructed only from its k nearest neighbors, resulting in $S_{lm} = 0$ if $\boldsymbol{x}_m$ does not belong to this set; (b) The row entries of the weight matrix sum to one, i.e., $\sum_{m=1}^{s} S_{lm} = 1$ for the purpose of transitional invariance. The matrix of weights $\boldsymbol{S}$ reflects the temporal structure information. Once $\boldsymbol{S}$ is obtained, LLE seeks a low–dimensional filtered space that preserves the temporal structure information of EEG trials as faithfully as possible. Let $\boldsymbol{z}_l\,(1 \le l \le s)$ be the filtered signal of $\boldsymbol{x}_l\,(1 \le l \le s)$ via the linear transformation $\boldsymbol{z}_l = \boldsymbol{\omega}^T\boldsymbol{x}_l$. One wishes to minimize the cost function

$$\Phi(\boldsymbol{Z}) = \sum_{l=1}^{s} \|\boldsymbol{z}_l - \sum_{m=1}^{s} S_{lm}\boldsymbol{z}_m\|^2 \tag{5}$$

where $\boldsymbol{Z} = [\boldsymbol{z}_1, \boldsymbol{z}_2, \cdots, \boldsymbol{z}_s]$. Note that the weights matrix $\boldsymbol{S}$ is fixed here and the transformation matrix $\boldsymbol{Z}$ is to be optimized. By substituting $\boldsymbol{z}_l = \boldsymbol{\omega}^T \boldsymbol{x}_l$ into (5), it follows that

$$\Phi(\boldsymbol{\omega}) = \sum_{l=1}^{s} \|\boldsymbol{\omega}^T \boldsymbol{x}_l - \sum_{m=1}^{s} S_{lm} \boldsymbol{\omega}^T \boldsymbol{x}_m\|^2 \tag{6}$$

With some matrix operations, (6) can be rewritten as

$$\Phi(\boldsymbol{\omega}) = \boldsymbol{\omega}^T \boldsymbol{X} \boldsymbol{L} \boldsymbol{X}^T \boldsymbol{\omega} \tag{7}$$

where $\boldsymbol{X} = [\boldsymbol{x}_1, \boldsymbol{x}_2, \cdots, \boldsymbol{x}_s]$, $\boldsymbol{L} = \left(\boldsymbol{I}_s - \boldsymbol{S}^T\right)(\boldsymbol{I}_s - \boldsymbol{S})$, and $\boldsymbol{I}_s$ is an $s \times s$ identity matrix.

We now incorporate $\Phi(\boldsymbol{\omega})$ into the objective function of the classical CSP in order to penalize solutions such that the temporal structure information is preserved. Formally, the objective function of our TRCSP is given by

$$J(\boldsymbol{\omega}) = \frac{\boldsymbol{\omega}^T \overline{C_x} \boldsymbol{\omega}}{\boldsymbol{\omega}^T \overline{C_y} \boldsymbol{\omega} + \alpha \left(\boldsymbol{\omega}^T \boldsymbol{X} \boldsymbol{L} \boldsymbol{X}^T \boldsymbol{\omega}\right)} = \frac{\boldsymbol{\omega}^T \overline{C_x} \boldsymbol{\omega}}{\boldsymbol{\omega}^T \left(\overline{C_y} + \alpha \boldsymbol{X} \boldsymbol{L} \boldsymbol{X}^T\right) \boldsymbol{\omega}} \tag{8}$$

Maximizing $J(\boldsymbol{\omega})$, would leads to the minimization of $\Phi(\boldsymbol{\omega})$, thus modifying spatial filters so as to satisfy the prior information. The parameter α is a user-defined positive constant which adjust the influence of the regularization term $\Phi(\boldsymbol{\omega})$. The higher the value of α is, the more favor the regularization term is given. The corresponding eigenvalue equation of (8) boils down to

$$\overline{C_x} \boldsymbol{\omega} = \lambda \left(\overline{C_y} + \alpha \boldsymbol{X} \boldsymbol{L} \boldsymbol{X}^T\right) \boldsymbol{\omega} \tag{9}$$

Thus, the filters $\boldsymbol{\omega}$ maximizing $J(\boldsymbol{\omega})$ are the leading eigenvectors corresponding to the largest eigenvalues. In the other hand, we need to accordingly maximize the dual objective function

$$J(\boldsymbol{\omega}) = \frac{\boldsymbol{\omega}^T \overline{C_y} \boldsymbol{\omega}}{\boldsymbol{\omega}^T \overline{C_x} \boldsymbol{\omega} + \alpha \left(\boldsymbol{\omega}^T \boldsymbol{X} \boldsymbol{L} \boldsymbol{X}^T \boldsymbol{\omega}\right)} = \frac{\boldsymbol{\omega}^T \overline{C_y} \boldsymbol{\omega}}{\boldsymbol{\omega}^T \left(\overline{C_x} + \alpha \boldsymbol{X} \boldsymbol{L} \boldsymbol{X}^T\right) \boldsymbol{\omega}} \tag{10}$$

Eventually, the spatial filters used are the leading eigenvectors corresponding to the eigenvalue problems of (8) and (10).

It is noted that in the above formulation of TRCSP, $\boldsymbol{X}$ denotes a general EEG trial. In implementation, we exploit the temporally local information of all the training trials. Specifically, we sum up all the locally linear structure expression as the final regularization term. Besides, TRCSP has two parameters: k which defines the number of the nearest neighbor samples, and α which defines the level of regularization. In the following experiments, the two parameters are specified with ten–fold cross validation on the training data. And we adopt linear discriminant analysis (LDA) as the classifier.

3 Materials for Evaluation

Three EEG data sets from public BCI competitions, recorded from totally 17 subjects, are used to assess the proposed TRCSP, Its performance is compared to the classic CSP algorithm.

3.1 EEG Data Sets

Data set IVa of BCI competition III is of two–class motor imagery (MI) paradigm by recording 5 subjects. Imagination of right hand and foot movements was performed after a visual cue per trial. The EEG measurements were recorded using 118 electrodes and sampled with 100 Hz. For each subject, there are totally 280 trials for two classes, 140 per class. Among them, 168, 224, 84, 56 and 28 training trials are respectively for subject 1 through 5.

Data set IIIa of BCI competition III contains EEG signals from 3 subjects, who performed 4 classes cued motor imagery, i.e., left hand, right hand, foot, and tongue MI. The EEG measurements were recorded using 60 sensors by a 64–channel EEG amplifier from Neuroscan. The EEG was sampled with 250 Hz and filtered between 1 and 50 Hz with Notchfilter on. In our study, only EEG data corresponding to right and left hands MI are used. In both of the training and testing sets, 45 trials per class are used for subject B1, and 30 trials per class for subject B2 and B3.

Data set IIa of BCI competition IV was constructed by recording 9 subjects, who carried out left hand, right hand, both feet and tongue MI tasks. 22 EEG channels were recorded. Signals were sampled with 250 Hz and bandpass filtered between 0.5 and 100 Hz with Notchfilter on. Only EEG signals of left and right hands MI are used for the present study. Each subject participated a training and a testing session, both sessions containing 72 trials for each class.

Table 1. Classification performances of CSP and TRCSP. The best percentage accuracy is displayed for each subject in the two Data sets of BCI competition III.

	BCI competition III								Overall	
	Data set IVa					Data set IIIa				
Subject	A1	A2	A3	A4	A5	B1	B2	B3	Mean	std
CSP	66.07	91.07	53.6	71.88	52.78	96.67	61.67	96.67	73.8	17.3
TRCSP	68.75	100	62.2	82.14	85.71	96.67	68.33	96.67	82.56	13.8

Table 2. Classification performances of CSP and TRCSP. The best percentage accuracy is displayed for each subject in Data set IIa of BCI competition IV.

	Data set IIa, BCI competition IV									Overall	
Subject	C1	C2	C3	C4	C5	C6	C7	C8	C9	Mean	std
CSP	86.11	57.64	96.5	70.1	60.42	70.14	82.64	93.0	93.75	78.92	13.9
TRCSP	87.5	63.89	97.9	70.1	65.97	68.75	81.94	95.83	92.36	80.47	12.7

3.2 Preprocessing

The EEG signals are band–pass filtered with cutoff frequencies 8 Hz and 30 Hz by using a fifth order Butterworth filter as recommended in [10]. Following the winner of BCI competition IV and [11], we use the time interval from 0.5 s to 2.5 s after the visual cue that indicates the start of imaginary as samples on all of the three data sets.

4 Results and Discussion

We use CSP and TRCSP to extract features on the data sets. Compared with CSP, there are two parameters in TRCSP which need to be configured. The parameters are selected by using ten–fold cross–validation method on the training sets. For each subject, the spatial filters are learnt on the training set available. As suggested in [7], three pairs of spatial filters for feature extraction are calculated in CSP and TRCSP. Then the log–variances of the spatially filtered EEG signals are used as input features for LDA. The results of classification accuracies and mean accuracies, as well as the corresponding standard deviations, are reported in Tables 1 and 2.

All the classification accuracies performed by TRCSP are larger than 60%. On average, TRCSP achieves better classification accuracies (mean: 81.45 ± 13.3) than CSP (mean: 76.51 ± 15.8). Whereas, it seems that TRCSP does not give high increase in classification accuracy for subjects who already have good performances (except for A2). With a closer look, results show that, for some subjects, using TRCSP leads to dramatic increase in performance as high as 10%, even higher than 30% for the subject whose performance is close to random by CSP (A5). It is interesting that the classification accuracy for A2 is always kept in 100% when using TRCSP in a wild range of parameters. Especially for the data set IVa of BCI competition III, the performance of TRCSP is much better than CSP. For the subjects A3, A4 and A5, TRCSP significantly enlarges the classification accuracies compared with CSP. It is probably because of the very small training set for these three subjects. It implies that adding a prior information, here a locally linear preserving penalty can help to find spatial filters despite the limited amount of training data, as agreed with [12].

Surprisingly, TRCSP leads to poorer performance than CSP on a few subjects, focusing on data set IIa of BCI competition IV. This might be due to the instability of EEG signal itself and the playing condition of subjects. Besides,

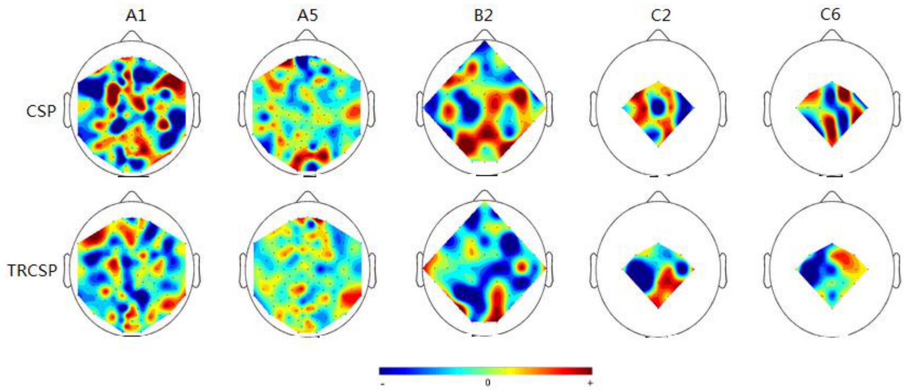

Fig. 1. Mappings of spatial filters obtained with CSP and TRCSP, for some subjects: A1, A5 (118 electrodes), B2 (60 electrodes), and C2, C9 (22 electrodes).

there is a very important point we can not ignore. That is, the best parameter may not be exactly found by the cross–validation strategy. It means that the real classification capacity TRCSP could possibly achieve better performances.

Some mappings of spatial filters obtained with both of CSP and TRCSP are presented in Fig. 1. The deeper color represents the greater weights. That is, they are more important for classification. In general, these pictures show that the CSP filters with large weights distribute in the whole brain, roughly and irregularly. Relatively, the TRCSP filters are generally smoother and more in line with the physiological characteristics. As expected from cerebral physiological theory, the weights are stronger over the motor cortex area. This suggests that the TPCSP algorithm lead to filters with more neurophysiological reality.

5 Conclusion

In this paper, we propose a new approach, called TRCSP, for optimizing spatial filers by incorporating temporal structure information to the conventional CSP. We add a locally linear regularization term to the CSP objective function. The experimental results confirm that TRCSP has the ability to obtain improved accuracies. In the future, much work is still needed to tune the appropriate parameters of TRCSP.

Acknowledgments. This work was supported in part by the National Basic Research Program of China under Grant 2015CB351704, the National Natural Science Foundation of China under Grant 61375118, the Natural Science Foundation of Jiangsu Province under Grant BK2011595, and the Program for New Century Excellent Talents in University of China under Grant NCET-12-0115. The authors thank Dr. F. Lotte for providing the code of filter mapping.

References

1. Wolpaw, J.R., Birbaumer, N., McFarland, D.J., Pfurtscheller, G., Vaughan, T.M.: Brain-computer interfaces for communication and control. Clinical Neurophysiology 113(6), 767–791 (2002)
2. Hwang, H.J., Kim, S., Choi, S., Im, C.H.: EEG-based brain-computer interfaces: A thorough literature survey. International Journal of Human-Computer Interaction 29(12), 814–826 (2013)
3. Blankertz, B., Muller, K.R., Curio, R., et al.: The BCI competition2003: Progress and perspectives in detection and discrimination of EEG single trials. IEEE Transactions on Bio-medical Engineering 51(6), 1044–1051 (2004)
4. Blankertz, B., Muller, K.R., Krusienski, D.J., et al.: The BCI competition III: Validating alternative approaches to actual BCI problems. IEEE Transactions on Neural Systems and Rehabilitation Engineering 14(2), 153–159 (2006)
5. Grosse-Wentrup, M., Liefhold, C., Gramann, K., Buss, M.: Beamforming in non-invasive brain computer interfaces. IEEE Transactions on Biomedical Engineering 56(4), 1209–1219 (2009)

6. Blankertz, B., Kawanabe, M., Tomioka, R., Hohlefeld, F., Nikulin, V., Muller, K.R.: Invariant common spatial patterns: Alleviating nonstationarities in brain-computer interfacing. Advances in Neural Information Processing Systems 20, 113–120 (2008)
7. Roweis, S.T., Saul, L.K.: Nonlinear dimensionality reduction by locally linear embedding. Science 290, 2323–2326 (2000)
8. Blankertz, B., Tomioka, R., Lemm, S., Kawanabe, M., Muller, K.R.: Optimizing spatial filters for robust EEG single-trial analysis. IEEE Signal Processing Magazine 25(1), 41–56 (2008)
9. Parra, L.C., Spence, C.D., Gerson, A.D., Sajda, P.: Recipes for linear analysis of EEG. NeuroImage 28(2), 326–341 (2005)
10. Ramoser, H., Muller-Gerking, J., Pfurtscheller, G.: Optimal spatial filtering of single trial EEG during imagined hand movement. IEEE Transactions on Rehabilitation Engineering 8(4), 441–446 (2000)
11. Lotte, F., Guan, C.: Regularizing common spatial patterns to improve BCI designs: Unified theory and new algorithms. IEEE Transactions on Biomedical Engineering 58(2), 355–362 (2011)
12. Lotte, F., Guan, C.: Spatially regularized common spatial patterns for EEG classification. In: Proceedings of the 20th International Conference on Pattern Recognition (ICPR), Istanbul, Turkey (2010)