

Improving classification performance of BCIs by using stationary Common Spatial Patterns and unsupervised bias adaptation

Wojciech Wojcikiewicz^{†‡◊}, Carmen Vidaurre[†], and Motoaki Kawanabe^{††*}

[†]Technical University of Berlin, Franklinstr. 28 / 29, 10587 Berlin, Germany

[‡]Fraunhofer Institute FIRST, Kekuléstr. 7, 12489 Berlin, Germany

[◊]Bernstein Center for Computational Neuroscience, Philippstr. 13, 10115 Berlin, Germany

wojwoj@mail.tu-berlin, carmen.vidaurre@tu-berlin.de, nabe@first.fraunhofer.de

Abstract. Non-stationarities in EEG signals coming from electrode artefacts, muscular activity or changes of task involvement can negatively affect the classification accuracy of Brain-Computer Interface (BCI) systems. In this paper we investigate three methods to alleviate this: (1) Regularization of Common Spatial Patterns (CSP) towards stationary subspaces in order to reduce the influence of artefacts. (2) Unsupervised adaptation of the classifier bias with the goal to account for systematic shifts of the features occurring for example in the transition from calibration to feedback session or with increasing fatigue of the subject. (3) Decomposition of the CSP projection matrix into a whitening and a rotation part and adaptation of the whitening matrix in order to reduce the influence of non-task related changes. We study all three approaches on a data set of 80 subjects and show that stationary features with bias adaptation significantly outperforms the other combinations.

Keywords: Brain-Computer Interface, Common Spatial Patterns, stationary features, adaptive classification

1 Introduction

Brain-Computer Interface (BCI) systems [4] aim to translate the intent of a subject measured from brain signals e.g. EEG into control commands for a computer application or a neuroprosthesis. A popular paradigm for BCI communication is motor imagery i.e. subjects perform the imagination of movements with their feet or hands, the imagined movements are detected by the system and translated into computer commands. The detection step often involves the extraction

* We thank Klaus-Robert Müller for valuable discussions. This work was supported by the German Research Foundation (GRK 1589/1), the European Union under the project TOBI (FP7-ICT-224631) and the Federal Ministry of Economics and Technology of Germany under the project THESEUS (01MQ07018). This publication only reflects the authors' views. Funding agencies are not liable for any use that may be made of the information contained herein.

of relevant features by a spatial filtering method called Common Spatial Patterns (CSP) [1]. The CSP filter is computed during the calibration session, but can be adversely affected by noise and non-stationarities e.g. coming from changes in impedance, muscular activity or eye movements. Moreover, even optimal filters may result in bad classification accuracies as the distribution of the features can change over time due to differences between sessions e.g. no feedback vs. feedback or changes of task involvement. Therefore this paper aims at both, we propose a regularized version of CSP to alleviate the impact of bad trials and noisy electrodes and we combine it with methods which adapt to changes in the feedback session.

Recently, several approaches were proposed to reduce the impact of non-stationarities in BCI applications. For example [10] uses techniques for co-adaptive learning of user and machine, [9] investigates different methods for unsupervised adaptation of the classifier and [8] uses adaptive spatial filtering. Other approaches use extra measurement like EOG or EMG to remove artefacts [2] or apply covariate shift adaptation to account for shifts of the features [5].

This paper is organized as follows. In Section 2 we present stationary CSP and the adaptation methods. Section 3 describes the experimental setup. After that we present and analyse the results in Section 4. Finally we concludes with a short summary and future research ideas.

2 Methods

Stationary CSP

Stationary Common Spatial Patterns (sCSP) is inspired by invariantCSP [2, 6]. However, in contrast to the latter it is completely data driven without using neurophysiological prior knowledge or extra measurements. The objective function of sCSP trades-off discriminativity and stationarity as it maximizes the variance of one class while minimizing the variance of the other class, but at the same time it penalizes non-stationary directions:

$$\max_{\mathbf{w}} \frac{\mathbf{w}^\top \boldsymbol{\Sigma}_+ \mathbf{w}}{\mathbf{w}^\top \{\boldsymbol{\Sigma}_+ + \boldsymbol{\Sigma}_- + \lambda \overline{\boldsymbol{\Delta}}\} \mathbf{w}}, \quad (1)$$

$$\max_{\mathbf{w}} \frac{\mathbf{w}^\top \boldsymbol{\Sigma}_- \mathbf{w}}{\mathbf{w}^\top \{\boldsymbol{\Sigma}_+ + \boldsymbol{\Sigma}_- + \lambda \overline{\boldsymbol{\Delta}}\} \mathbf{w}}. \quad (2)$$

Note that \mathbf{w} is the sCSP filter, $\boldsymbol{\Sigma}_+$ and $\boldsymbol{\Sigma}_-$ are average covariance matrices of the two classes, $\overline{\boldsymbol{\Delta}}$ is the regularization term and λ is a trade-off parameter.

In order to estimate the regularization term $\overline{\boldsymbol{\Delta}}$, we first compute the trial-wise covariance matrices $\boldsymbol{\Sigma}_+^{(k)}$ and $\boldsymbol{\Sigma}_-^{(k)}$ for both classes. After that we compute the deviations of the trial-wise covariance matrices from the class-average covariance matrix i.e. $\boldsymbol{\Delta}_i^{(k)} := \mathcal{P}(\boldsymbol{\Sigma}_i^{(k)} - \boldsymbol{\Sigma}_i)$ with i being the class indicator. Note that \mathcal{P} is an operator to force symmetric matrices be positive definite. This assures that the penalty term is always positive. In our study we simply flip the sign of

negative eigenvalues. The regularization term $\overline{\Delta}$ is finally defined as the average of the difference matrices $\Delta_i^{(k)}$ over all trials k and the two classes i .

Unsupervised Bias Adaptation

In this paper we use the bias adaptation method PMean introduced in [9]. This method adapts the bias b of an Linear Discriminant Analysis (LDA) classifier $\mathbf{w}^T \mathbf{x} + b(t)$ by updating the global mean $\boldsymbol{\mu}$ of the features:

$$b(t) = \mathbf{w}^T \boldsymbol{\mu}(t-1), \quad (3)$$

$$\boldsymbol{\mu}(t) = (1 - \eta) \cdot \boldsymbol{\mu}(t-1) + \eta \cdot \mathbf{x}(t), \quad (4)$$

where t is the trial number, $\eta \in [0, 1]$ is a constant, \mathbf{w} are the weights of the LDA classifier and \mathbf{x} is the feature vector. Note that the adaptation is only performed in the feedback session and only affects the bias and not the weights.

Whitening Adaptation

The main idea of whitening adaptation is to decompose the CSP projection into a whitening \mathbf{P} and a rotation \mathbf{B} part. The rotation part is assumed to be fixed, but the whitening matrix is updated during feedback. This idea was introduced in [8]. There the covariance matrix was updated using blocks of data. In this work, we update the covariance using the same formula as for PMean, i.e. equation (4), because this approach only needs to store one sample rather than a block of data and allows a continuous adaptation:

$$\mathbf{P}(t) = \mathbf{V} \mathbf{D}^{-1/2} \mathbf{V}^T, \quad (5)$$

$$\boldsymbol{\Sigma}(t) = (1 - \eta) \cdot \boldsymbol{\Sigma}(t-1) + \eta \cdot \boldsymbol{\Sigma}^t, \quad (6)$$

where t is the trial number, \mathbf{V} and \mathbf{D} are the eigenvectors and eigenvalues of $\boldsymbol{\Sigma}(t-1)$, $\eta \in [0, 1]$ is a constant and $\boldsymbol{\Sigma}^t$ is the covariance matrix of trial t . As before the adaptation is only performed in the feedback session and is unsupervised.

3 Experimental Setup

The experiments are based on data from a joint study [3] with University Tübingen. We asked 80 volunteers to perform motor imagery tasks with the left and right hand or with the feet. For each user we select the best binary task-combination and estimate parameters like frequency band or time interval of interest in a calibration session (150 trials) without providing feedback. After that we perform a test session consisting of 300 trials and provide 1D visual feedback i.e. a moving arrow on a screen. All subjects in this study are BCI novices. We use recordings of 68 preselected electrodes densely covering the motor cortex, three CSP directions per class, power features, an LDA classifier and error rate to measure performance. We select the λ parameter for sCSP from the set of candidates $\{0,$

$\{0.1, 0.025, 0.05, 0.075, 0.1, 0.25, 0.5, 0.75, 1, 2.5, 5, 10\}$ by 5-fold cross-validation on the calibration data. As the authors in [9] we use $\eta = 0.05$ for PMean. For the whitening adaptation η is chosen to be one order of magnitude smaller i.e. 0.005 as the method is more sensitive to changes.

4 Results

At first we compare the performances of CSP and sCSP with and without using adaptation. Fig. 1 shows six scatter plots, each displaying the error rates of two methods which we want to compare. We see that in general sCSP performs much better than CSP, although for some subjects it can lead to slightly worse results. We will analyse the reasons for the performance increase further below. We also see that using bias adaptation improves classification results as it reduces shifts in the features which can occur between calibration and test session¹. However, except for a few users PMean is superior to whitening adaptation (see right bottom panel). We also investigated the combination of sCSP with both adaptation methods, but it did not bring any significant improvements over sCSP and PMean. We also observed unequal false positive and false negative rates when using CSP or sCSP. This bias comes from a shift in the distribution of features between calibration and test phase and can be resolved by bias adaptation.

In order to evaluate the significance of our results we use the Wilcoxon signed-rank test. Table 1 shows the p-values (one-sided) of different comparisons using either all subjects or dividing them into three groups based on their error rates: 0% - 15%, 15% - 30% and above 30%. Note that the null hypothesis of the test states that the median of the distribution of error rate differences is zero. Together with the scatter plots from Fig. 1 we can not only tell whether there is a significant difference between two methods or not, but also which one performs better. Most significant improvements are obtained for subjects in the above 30% error rate group. For these participants both using stationary features and adapting to changes is advantageous, whereas users who perform well do not benefit from sCSP. This is intuitive as subjects with low error rates usually have a clean signal i.e. it is not affected by artefacts and the signal to noise ratio is higher than in subjects who lack BCI efficiency. Therefore using sCSP does not bring any advantage. However, since these subjects can also suffer from shifts in the features (as introduced by changing from calibration of feedback conditions), adaptation methods may improve the classification results.

In the following we would like to analyse the reasons for the performance gain on a specific participant. This subject performs left vs. right motor imagery and has an error rate of 28.67% when using CSP. However, the error rate decreases to 8% when applying sCSP (with $\lambda = 0.25$), it is 6% for both sCSP_{pmean} and sCSP_{whitening} and it is 10.33% for CSP_{pmean} and 9.67% for CSP_{whitening}. So why does CSP perform so poorly compared to the other methods ?

In order to answer this question we will first identify the dimensions which are responsible for the bad performance of CSP. It seems that two out of the

¹ Note that visual feedback was only provided in the test session.

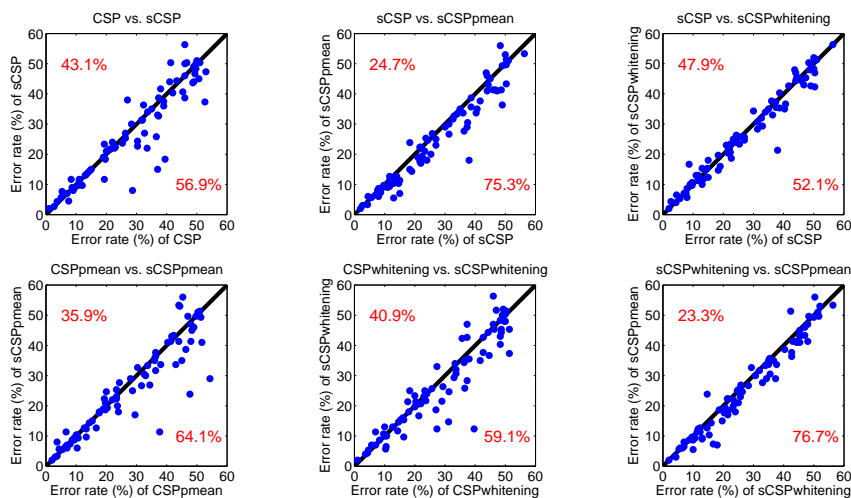


Fig. 1. Six scatter plots comparing the error rates of different methods. Each subject is represented by a blue dot. The percentage of points lying above or below the black line is shown in red and indicates which of the compared methods performs better. We see that both sCSP and bias adaptation, especially PMean, decrease error rates.

six CSP filters are corrupted by noise and non-stationarities and degrade the overall performance since when classifying the data without dimensions 2 and 3, the error rate of CSP decreases to 10%. If on the other hand we include one of the noisy dimensions, the error rate increases to more than 20%. Fig. 2 shows the filters of dimension 2 and 3 for CSP and sCSP using the same scale. We see that the CSP filters heavily weight electrodes at different locations including areas which are primarily not responsible for performing motor imagery like frontal and temporal areas. Since the filters are obviously corrupted, the error rates of the corresponding features are almost at chance level. In contrast to that the corresponding sCSP filters are much less affected by artefacts in the data which leads to a more reasonable weighting of the channels and smaller error rates.

A remaining question is what effect does the reduction of the weights have on the features. In Fig. 3 we plot the calibration and test features of the noisy dimensions 2 and the most discriminant² dimension 4 for CSP and sCSP. As can be seen from the left panel there is a significant shift in dimension 2 between calibration and test features. Since the influence of bad trials and noisy electrodes is reduced when using sCSP, this shift is much weaker in the right panel. So we conclude that for this subject non-stationarities in the data, most probably eye movements and some noisy electrodes, have corrupted two CSP filters and thus lead to shifts in the features. These shifts can be either removed by applying adaptation methods like PMean or by computing stationary features with sCSP.

² The error rate of this dimension is 16% for CSP and 17% for sCSP.

Compared Methods	Groups based on error rate (in %)			
	0 – 15	15 – 30	> 30	all
sCSP better than CSP	0.4924	0.4510	0.0118	0.0301
sCSPpmean better than sCSP	0.0033	0.0076	0.0006	0
sCSPwhitening better than sCSP	0.4925	0.0098	0.0103	0.0022
sCSPpmean better than CSPpmean	0.2429	0.2035	0.0192	0.0117
sCSPwhitening better than CSPwhitening	0.2146	0.4691	0.0089	0.0138
sCSPpmean better than sCSPwhitening	0.0605	0.0011	0.0130	0
sCSPpmean better than CSP	0.0026	0.0014	0.0022	0

Table 1. Overview of p-values for different method combinations using Wilcoxon signed-rank test (one-sided). Bold values are significant when $\alpha = 0.05$. Grouping is performed based on the error rates of the worse method (x-axis in Fig. 1).

If we analyse Fig. 4 we see that using whitening adaptation also reduces the shift in dimension 2 and 3. This figure shows the mean absolute difference between calibration and test features for each dimension after a different number of updates of the covariance matrix. The reduction of the shift in dimension 2 and 3 is probably the main reason for the performance gain of CSPwhitening. However, whitening adaptation does not only reduce the shift between dimensions but can also increase it e.g. for dimension 6. We conjecture that this may be a reason why whitening adaptation is inferior to PMean.

We would further like to note that not in all cases sCSP and PMean have the same effect, namely the reduction of the shift in the features. For some subjects applying sCSP helps a lot, whereas CSPpmean or CSPwhitening does not perform any better than CSP. In these subjects the influence of noise and non-stationarities is often so strong that all CSP filters are more or less corrupted. Thus adaptation methods can not improve results as no dimension is discriminative. In our example on the other hand there is a discriminative dimension after application of CSP, but other noisy dimensions lead to a shift in the features and affect the performance negatively. In both cases sCSP reduces the influence of noisy and non-stationarity directions and can improve classification accuracy. However, since sCSP is computed on calibration data it is not able to capture changes which does not appear during this phase. Therefore combining sCSP with PMean provides best results as it reduces the adverse effects on non-stationarities in both the calibration and the test session.

5 Conclusion

We presented an extension of CSP which explicitly measures non-stationarities and regularizes the CSP directions towards stationary subspaces. We showed that the main reason why sCSP improves performance is that it reduces the influence of bad trials (artefacts) and noisy electrodes which can corrupt the CSP filters. Unlike other methods, such as invariantCSP, our method is completely data-

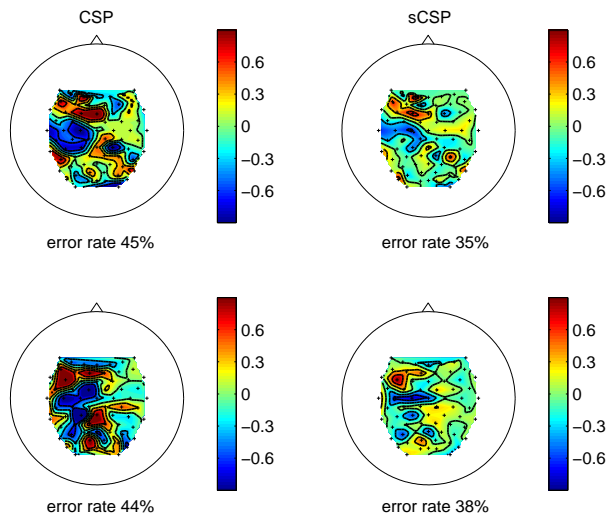


Fig. 2. Filters of noisy dimensions 2 (top) and 3 (bottom) for CSP and sCSP.

driven and does not need additional recordings or models of the expected change that occurs in the EEG. However, it cannot capture changes which occur only in the feedback session since sCSP is computed on calibration data. Therefore further methods of adaptation are required. We showed that combining sCSP with PMean significantly outperforms the baseline and is also superior to sCSP with whitening adaptation.

In future research we would like to investigate other data-driven regularization criteria and compare sCSP to invariantCSP with additional EOG and EMG recordings.

References

1. B. Blankertz, R. Tomioka, S. Lemm, M. Kawanabe, and K-R. Müller, “Optimizing spatial filters for robust eeg single-trial analysis,” in *IEEE Signal Proc. Magazine*, 2008, pp. 581–607.
2. B. Blankertz, M. Kawanabe R. Tomioka, F. U. Hohlefeld, V. Nikulin, and K-R. Müller, “Invariant common spatial patterns: Alleviating nonstationarities in brain-computer interfacing,” in *Ad. in NIPS 20*, 2008, p. 2008.
3. B. Blankertz, C. Sannelli, S. Halder, E.M. Hammer, A. Kübler, K-R. Müller, G. Curio, T. Dickhaus “Neurophysiological Predictor of SMR-Based BCI Performance,” in *NeuroImage*, 51(4):1303-1309, 2010.
4. G. Dornhege, J. del R. Millán, T. Hinterberger, D. McFarland, and K-R. Müller, Eds., *Toward Brain-Computer Interfacing*, MIT Press, Cambridge, MA, 2007.
5. Y. Li, H. Kambara, Y. Koike, and M. Sugiyama, “Application of covariate shift adaptation techniques in brain-computer interfaces,” *IEEE Trans. Biomed. Eng.*, vol. 57, no. 6, pp. 1318–24, 2010.




Fig. 3. Visualization of feature dimension 2 and 4 for CSP and sCSP. The shift between calibration and test features in dimension 2 is due to eye movements artefacts which adversely affect the CSP filter, but have less impact when using sCSP.




Fig. 4. Mean absolute differences between calibration and test features for different feature dimensions after a different number of update steps of the covariance matrix.

6. S. Mika, G. Rätsch, J. Weston, B. Schölkopf, A. Smola, and K.-R. Müller, “Invariant feature extraction and classification in kernel spaces,” in *Ad. in NIPS 12*, 2000, pp. 526–532.
7. H. Ramoser, J. Müller-Gerking, and G. Pfurtscheller, “Optimal spatial filtering of single trial EEG during imagined hand movement,” *IEEE Trans. Rehab. Eng.*, vol. 8, pp. 441–446, 1998.
8. R. Tomioka, J.N. Hill, B. Blankertz, and K. Aihara, “Adapting Spatial Filter Methods for Nonstationary BCIs,” *Proc. of Workshop on Information-Based Induction Sciences (IBIS)*, 65-70, 2006.
9. C. Vidaurre, M. Kawanabe, P. von Bünau, B. Blankertz, and K-R. Müller, “Toward an unsupervised adaptation of LDA for Brain-Computer Interfaces,” *IEEE Trans. Biomed. Eng.*, 2010, in press.
10. C. Vidaurre, C. Sannelli, K-R. Müller, and B. Blankertz, “Machine-learning based co-adaptive calibration,” *Neural Comput.*, 2010, in press.