

Mind-Mirror: EEG-Guided Image Evolution

Nima Bigdely Shamlo, Scott Makeig

¹Swartz Center for Computational Neuroscience, Institute for Neural
Computation, University of California, San Diego, USA
{nima, scott}@scn.ucsd.edu

Abstract. We propose a brain-computer interface (BCI) system for evolving images in real-time based on subject feedback derived from electroencephalography (EEG). The goal of this system is to produce a picture best resembling a subject's 'imagined' image. This system evolves images using Compositional Pattern Producing Networks (CPPNs) via the NeuroEvolution of Augmenting Topologies (NEAT) genetic algorithm. Fitness values for NEAT-based evolution are derived from a real-time EEG classifier as images are presented using rapid serial visual presentation (RSVP). Here, we report the design and performance, for a pilot training session, of a BCI system for real-time single-trial binary classification of viewed images based on participant-specific brain response signatures present in 128-channel EEG data. Selected training-session image clips created by the image evolution algorithm were presented in 2-s bursts at 8/s. The subject indicated by subsequent button press whether or not each burst included an image resembling two eyes. Approximately half the bursts included such an image. Independent component analysis (ICA) was used to extract a set of maximally independent EEG source time-courses and their 100 minimally-redundant low-dimensional informative features in the time and time-frequency amplitude domains from the (94%) bursts followed by correct manual responses. To estimate the likelihood that the post-image EEG contained EEG 'flickers' of target recognition, we applied two Fisher discriminant classifiers to the time and/or time-frequency features. The area under the receiver operating characteristic (ROC) curve by tenfold cross-validation was 0.96 using time-domain features, 0.97 using time-frequency domain features, and 0.98 using both domain features.

Key words: human-computer interface (HCI), brain-computer interface (BCI), evolutionary algorithms, genetic algorithms (GA), electroencephalography (EEG), independent component analysis (ICA), rapid serial visual presentation (RSVP), genetic art, evolutionary art.

1 Introduction

Recent advances in Evolutionary Algorithms have lead to the introduction of methods for evolving images based on human feedback [6]. These methods present a succession of images to a user who rates each image on how closely it resembles an imagined target picture. These user fitness ratings are exploited to evolve the sequence of pictures towards the imagined target. An example of such a system is Picbreeder (picbreeder.com), a collaborative web application for evolving pictures. Although these evolutionary methods can be quite successful in maximizing fitness,

they are limited by the speed at which humans can provide fitness-value feedback and the amount of feedback they can generate before becoming fatigued.

Several brain-computer interface (BCI) systems for broadening user feedback performance in image processing tasks have been developed in last few years (Bigdely-Shamlo et al. [2] and Sajda et al. [7], [8], [9]). These interfaces use a rapid serial visual presentation (RSVP) paradigm to present a stream of image ‘clips’ drawn from a larger search image at high presentation rates (about 12/s). Subjects look for clips containing a set of predefined target types (e.g., airplanes or helipads) among the clips presented in the stream while EEG signals are recorded and analyzed in near-real time to detect dynamic brain signatures of target recognition. These signatures are used to find the image clip in the sequence in which the subject spotted a target feature.

Basa et al. [5] proposed an EEG-based genetic art evolution algorithm based on evaluation of ‘interestingness’ of each image. They trained SVM classifiers on EEG and an array of other physiological measures recorded during exposure to emotion-inducing images. In test sessions, their classifiers were able to distinguish positive and negative emotions elicited by genetic art images, produced without subject online feedback, with 61% accuracy.

Here we propose a real-time EEG-based system for evolving images based on subject feedback produced by EEG classifier output during an RSVP stream of generated images. Bursts of images generated by an image evolution algorithm are presented to the subject. Images resembling the targeted or ‘imagined’ image elicit some level of target response which may be recognized by the EEG classification algorithm, generating feedback to the image evolution algorithm guiding the generation of the next RSVP images. This process is repeated continuously until the images in the RSVP stream wholly or mostly resemble the ‘imagined’ image with adequate detail and accuracy. At that point, the presented images would ‘mirror’ the original target image in the mind of the viewer.

2 Methods

2.1 Training Experiment

In previous work [2], we demonstrated online EEG classification to detect ‘flickers of recognition’ associated with spotting a target feature in a sequence of images displayed using rapid serial visual presentation (RSVP) at a rate of 12 images per second. Here, we report initial training session results of individualized EEG-based classification of images in an RSVP sequences that resemble an imagined target, here “two eyes facing the subject.” We propose that the derived classification model or one like it can be used in on-line test sessions to drive image evolution of imagined pictures.

In the training session, EEG and manual responses were recorded as the subject viewed the RSVP sequence. Fig. 1 shows a trial timeline. After fixating a cross (left) for 1 s, the participant viewed a 2-s RSVP image burst (presentation rate, 8 images per sec), and was then asked to indicate, by pressing one of two (yes/no) finger buttons, whether or not he/she had detected an image resembling the imagined (two-eyes) target among the burst images. Half of these bursts included a two-eyes image. Visual error/correct feedback was provided after each burst. The training session was comprised of 467 RSVP bursts organized into 52 9-burst bouts with a self-paced break after each bout. In all, each session thus included 253 (~3%) two-eyes and 7219 (~97%) not-eyes image presentations. No burst contained more than one target (two-eyes) image.

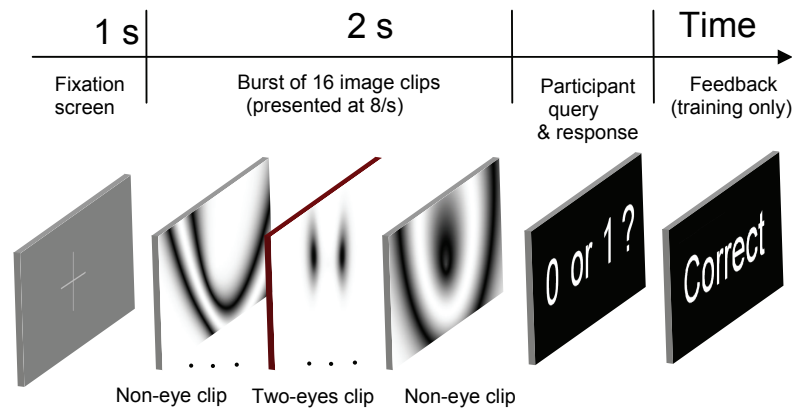


Fig. 1. Time-line of each RSVP burst. Participant response feedback ('Correct' or 'Incorrect') was delivered only during Training sessions (rightmost panel).

2.2 Stimuli

We used a modified version of the GeneticArt program written in Delphi by Mattias Fagerlund [14] to create images used in the training session. It generates gray-scale images produced by a Compositional Pattern Producing Network (CPPN) [10] evolved using the NeuroEvolution of Augmenting Topologies (NEAT) algorithm [1]. The CPPN used in the experiment is a network of Gaussian function nodes connected by a directed graph. The input to each node is the sum of values from input connections weighted by connection weights. The output of each node is sent via output connections to other nodes in the network. The function of the network is to generate a gray-scale image. Its final output is a matrix of light intensities at each pixel of the output image. To generate an image, the light intensity at each pixel is

calculated by providing values indicating the (x,y) location of the pixel in the image ($\epsilon \in [0,1]$) and its distance to the center of the image.

The topology of the CPPN network (its connections, connection weights, and node types) is evolved by the NEAT algorithm. Each new generation of images is created by applying genetic operators (mating, mutation...) to the individual images (or image pairs) in the previous generation based on their received fitness values. NEAT begins by evolving simple (low node- and connection-count) networks, gradually creating more a complex network that produce higher received fitness values. By starting the search from a low-dimensional space, NEAT increases the chance of finding a desirable solution by preventing the problems associated with searching in a high dimensional space (many nodes and connection weights). In our case, the network starts with simple shapes and can increase the amount of visual detail in later image generations.

We used two sets of two-eyes and not-eyes images generated by the network in previous behavioral pilot sessions to prepare the stimuli for the training session. To create non-eye images, we ran the GeneticArt program for many generations giving equal fitness values to all images in each generation. This produced thousands of ‘random’ images not shaped by a specific fitness function. We then went through these images and manually deleted any that had a resemblance to the target image. Two-eyes images were evolved by providing appropriate manual feedback to the GeneticArt program. After generating many of these images, we picked a variety of two-eyes images with different shapes for use as targets in the training session. We hypothesized that sudden changes in image size and contrast might induce EEG responses to visual surprise [11] that were not related to subject recognition of target resemblance. To minimize variations in feature size and contrast, we normalized the center of mass location and angular momentum magnitude of the images.

Two-second RSVP bursts, each containing 16 (547x343 pixel) images, were arranged using the normalized two-eyes and non-eye images. In half of these bursts, a two-eyes image was displayed at a randomly selected latency between 250 ms and 1750 ms. The 250-ms edge buffers prevented most interference during classification from EEG responses to burst onsets and offsets.

2.3 Experimental Setup

RSVP bursts were presented at a rate of 8 images per second and synchronized with EEG and behavioral data collection using our custom *DataRiver* interactive stimulus presentation and data recording suite (A. Vankov). Events corresponding to image display, image type (two-eyes or not-eyes), and subject post-burst manual responses were recorded synchronously with EEG into a single file. EEG was recorded using a BIOSEMI Active View 2 system with 128 electrodes mounted in a whole-head elastic electrode cap with a custom near-uniform montage across the scalp and bony parts of the upper face. Computer data acquisition was performed via USB using a customized

acquisition driver at a sampling rate of 512 Hz. In this pilot experiment, we recorded one hour of data from one (28-year old male) subject.

2.4 Data Preprocessing

Preprocessing of training session data was performed in the Matlab environment (Mathworks, Inc.) using EEGLAB [12] functions and data structures. The scalp channel data were first high-passed using an FIR filter above 2 Hz. Two channels containing invalid EEG signals were then removed. The rest of the EEG data were then re-referenced from the Biosemi active reference to average reference. EEG data epochs from 1.5 s before burst onset until 2 s after (2-s) burst offset were extracted from the continuous EEG and concatenated. Data epochs were then decomposed using extended-infomax ICA [4][17][18][19] to obtain 126 maximally independent components (ICs), spatial filters separating EEG sensor data into maximally temporally independent processes, most appearing to predominantly represent the contribution to the scalp data of one brain EEG or non-brain artifact source, respectively. Bursts followed by a correct subject response (438 of 467, ~94%) were used for classifier training and cross-validation.

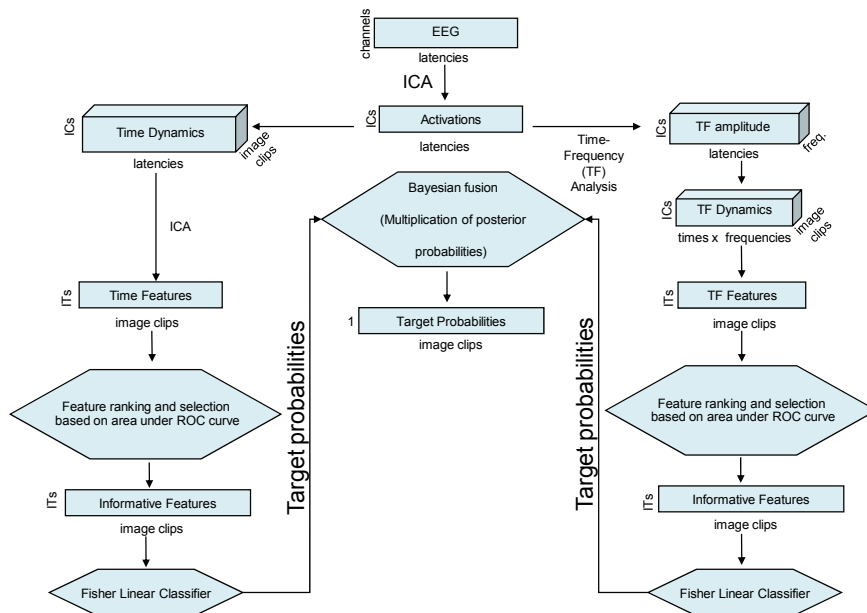


Fig. 2. Flowchart of the classifier system. Both (left) temporal activity patterns (phase-locking following stimulus onsets) and (right) event-related changes in spectral amplitude are used by the classifier (center). [From N. Bigdely Shamlo et al., *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2008].

2.5 Image Classification

We used methods described previously [2] for classifying image responses as target (two-eyes) or non-target (not-eyes). Fig. 2 summarizes the classification process. We exploited information from both time-domain and frequency-domain features of EEG signals following image presentations. Time-course templates were selected by a method that involved performing, for every IC, a second-level deconvolutive ICA decomposition on the time-domain IC activation following image presentations. To do this, for each IC we constructed matrices (of size 410 lags by 6570 images) by concatenating columns consisting of lagged IC activation time series in the 800 ms following each image presentation. Applying principal component analysis (PCA), we then reduced the rows of this matrix to 50. A second-level ICA was then performed on the dimension-reduced matrix to find 50 independent time course templates (IT_{TCS}) for each IC. These features form a natural spatiotemporal decomposition of the EEG signal with low interdependency, making them natural candidate features for an EEG classifier. Time-frequency amplitude ITs (IT_{AS}) were found by decomposing the single-trial spectrograms of the IC activations following image onsets, vectorizing these, then concatenating them into a large (frequencies \times latencies, images) matrix that was reduced to its first 50 principal dimensions by PCA before second-level ICA decomposition.

As expected, the weights for a number of time-domain and time-frequency domain dimensions contained robust information for target/non-target image classification. We used the 100 most informative ITs (over all ICs) for each domain (as determined by computing individual area under the ROC curves for the target/non-target discrimination) as input to a Fisher discriminant linear classifier. For time-frequency domain features, we used the independent template weights for each image as input. For time-domain features, we used columns of the transpose of the second-level ICA mixing matrix as matched filters (Fig. 3) since in earlier testing this had given slightly better results than using the decomposed IC activations themselves.

3 Results

3.1 Classification Performance

After preprocessing the data and extracting classification features as described above, we performed ten-fold cross-validation. Using only the most informative 100 time-domain features in each fold, an area under the ROC curve of 0.96 was achieved for this subject and session. Ten-fold cross-validation results for frequency-domain features alone (the top 100 in each validation fold) produced an area under the ROC curve of 0.97. Maximum performance was achieved when posterior probabilities of

being a target obtained from the time-domain and frequency-domain classifiers in each validation fold were combined using naive Bayesian fusion (i.e., when the two probabilities were multiplied). In this case, the area under the ROC curve was 0.98.

3.2 EEG ‘flickers’ of target recognition

Both time-domain and frequency-domain contain information regarding EEG signatures of target detection. To investigate the information content of sources used in the classifier, we calculated the ten-fold cross validation area under ROC for individual ICs. All independent templates (ITs) for each IC were used in the manner described in Fig. 2.

Equivalent dipole sources for ICs were then localized using Spherical Four shell (BESA) model and components with residual variance less than 0.2 were selected. Fig. 3 shows these dipoles in the standard MNI head space. Both size and color of each dipole reflect the amount of information relevant to target detection, measured by the area under ROC of the single-IC classifier. Most likely Brodmann areas associated with dipole locations, identified using the Talairach Client [13], are displayed on the most informative IC dipole spheres.

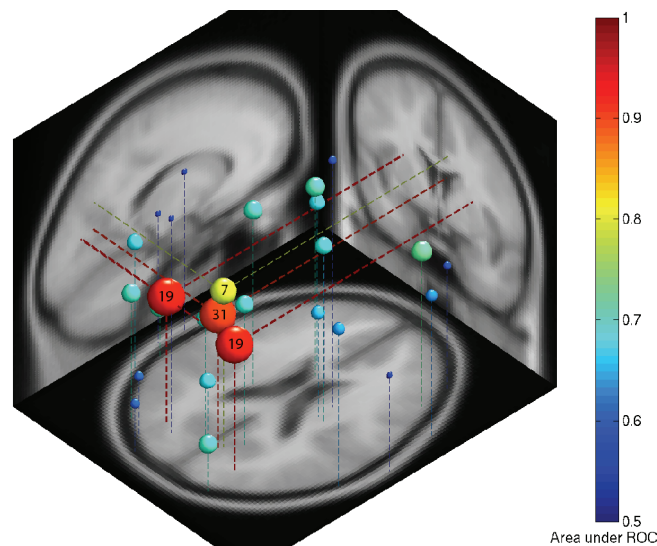


Fig. 3. Equivalent dipoles for subject ICs with low (< 0.2) dipole model residual variance. The color and size of each dipole sphere shows the relative amount of information about target classification in the IC activation. Numbers on the most informative (red-yellow) dipole markers indicate the Brodmann area most probably associated with the dipole location (BA19, BA31, etc.).

The most informative IC (Fig. 3) required a dual-symmetric dipole model with source areas located in or near BA 19, a secondary visual cortical area associated with visual processing, specifically, higher order visual processing involving object recognition and movement sensing [16]. Another informative IC was located in or near the precuneus (BA 31). Activity in this area is associated with attentional processes, including shifts of attention between object features [15]. A third most target-information bearing IC was located in or near right BA 7 (superior parietal lobule), an area associated with visual-spatial processing and use of spatial imagery. Note, that the most informative areas involved responses peaking 200 ms and more after target image onsets in higher regions of the brain visual system hierarchy, areas known to be involved in visual target recognition. A number of medial and frontal area sources (light blue balls in Fig. 3) also contributed information via later potentials that maximally projected to the frontal scalp.

4 Discussion

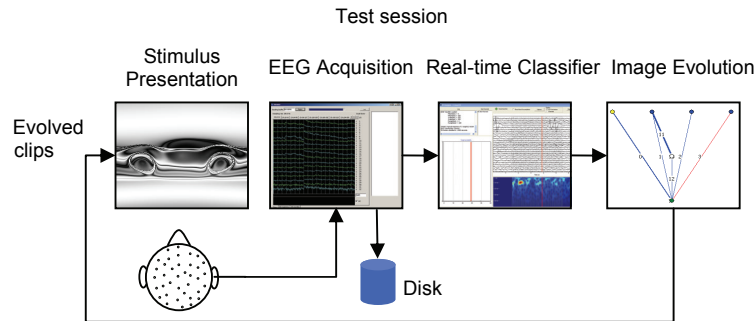


Fig. 4. Schematic of the proposed Mind-Mirror system in intended operation. Here, the subject is imagining an image of a sports car. EEG target recognition signals fed to the image evolution algorithm from the real-time EEG classifier continually shape the evolution of the image sequence to resemble more and more closely the subject’s imagined car image. In the test session, the subject observes RSVP image bursts produced by the image evolution algorithm. If subject detects an image with some resemblance to imagined picture, it will be reflected in a high target probability value output of the EEG classification algorithm. These values are then sent to the image evolution program to evolve a new generation of images, which are then displayed in the next RSVP burst. This establishes a feedback loop, shown in Fig. 3 (right), which favors image features resembling the imagined target picture. After the subject views many bursts, the system may consistently generate a ‘mirror image’ of the image the subject has been mentally imagining. By eliminating the need for subject manual feedback in the Mind-Mirror system, its speed and user convenience might be increased.

4.1 The Proposed Mind-Mirror System

The pilot training session we report here is the first step in creating the proposed Mind-Mirror system. Further system development involves testing the use of

classifier output online for image evolution (Fig. 4). Based on our initial RSVP studies [2], we hypothesize that the subject-specific classifier trained using a certain image target can likely be applied to online evolution of other image targets. Thus, although the Mind-Mirror classifier reported here was trained on a specific ('two-eyes') target image, we predict that during test sessions it will correctly detect an EEG 'target recognition' activity pattern for this subject in sessions using different target images. For example, the (manually guided) evolution algorithm can produce images resembling a car, house, or butterfly.

Despite the near-perfect image classification obtained by our method in this session, several factors could limit the success of the derived EEG model in guiding online image evolution. The EEG features used by the classifier might be attenuated, altered, or even absent in response to early sequence images that only vaguely resemble the image target, or in response to later target-like images presented more frequently than one target per burst. Such feature changes could degrade image classification and evolution performance in early and/or later stages of image evolution. To avoid these pitfalls, the training session images might better resemble early- or later-stage image evolution sequences. Solutions for several practical issues for online classification related to dealing with session-to-session shifts in the electrode montage, novel within-session artifacts, and channel drop-outs have been described in our earlier RSVP paper [2].

Based on the strong recent acceleration of technical, business, and popular interest, appears that applications of EEG-based human-computer interaction are just beginning. Very likely, the best uses for EEG in brain-computer interfaces should be those that respect the natural brain functions associated with the features used for interaction. The proposed Mind Mirror exemplifies this by making use of EEG features that constitute the complex EEG 'flicker of recognition' that immediately follows recognition of a sought visual target. The Mind Mirror system, if sufficiently robust, might be used as a game, to create art, or using a somewhat different image evolution algorithm, to guide production of a face image resembling a suspect as recalled by a witness in a criminal investigation.

References

1. Stanley, K. O.: Evolving Neural Networks through Augmenting Topologies. *Evolutionary Computation*, vol. 10(2), pp. 99-. (2002)
2. Bigdely-Shamlo, N., Vankov, A., Ramirez, R., Makeig, S.: Brain Activity-Based Image Classification From Rapid Serial visual Presentation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 16, no 4 (2008)
3. Lee, T., Girolami, M., Sejnowski, T. J.: Independent Component Analysis Using an Extended Infomax Algorithm for Mixed Subgaussian and Supergaussian Sources. *Neural Computation*, vol. 11, no. 2, pp. 417--441 (Feb. 1999)
4. Bell, A., Sejnowski, T.: An Information Maximization Approach to Blind Separation and Blind Deconvolution. *Neural Computation*, vol. 7, pp. 1129--1159 (July 1995)
5. Basa, T., Go, C., Yoo, K., Lee, W.: Using Physiological Signals to Evolve Art., *LNCS*, vol. 3907, pp. 633--641 (2006)

6. Secretan, J., Beato, N., D'Ambrosio, D. B., Rodriguez, A., Campbell, A., Stanley, K. O.: Picbreeder: Evolving Pictures Collaboratively Online. Proc. Computer Human Interaction Conf. (CHI), New York, NY, ACM, 10 pages (2008)
7. Sajda, P., Gerson, A., Parra, L.: High-throughput image search via single-trial event detection in a rapid serial visual presentation task. In: Proc. 1st Inter. IEEE EMBS Conf. on Neural Engineering, Capri Island, Italy (2003)
8. Gerson A., Parra L., Sajda P.: Cortically-coupled computer vision for rapid image search. IEEE Transactions on Neural Systems and Rehabilitation Engineering, vol. 14 (2), pp. 174--179 (2006)
9. Parra, L.C., Christoforou, C., Gerson, A. D., Dyrholm, M., Luo, A., Wagner, M., Philiastides, M. G., Sajda, P.: Spatio-temporal linear decoding of brain state: Application to performance augmentation in high-throughput tasks. IEEE Signal Processing Magazine, vol. 25, no. 1, pp. 95--115 (Jan. 2008)
10. Stanley, K. O.: Exploiting Regularity Without Development. Proc. AAAI Fall Symposium on Developmental Systems, Menlo Park, CA, AAAI Press, 8 pages (2006)
11. Einhuser, W., Mundhenk, T. N., Baldi, P., Koch, C., Itti, L.: A bottom-up model of spatial attention predicts human error patterns in rapid scene recognition. J. of Vision, vol.7, no.10, pp.1--13 (2007)
12. Delorme, A., Makeig, S.: EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. J. of Neuroscience Methods, vol. 134 (1), p. 9, <http://sccn.ucsd.edu/eeglab>
13. Talairach Client, <http://www.talairach.org/client.html>
14. DelphiNEAT GeneticArt program, <http://www.mattiasfagerlund.com/DelphiNEAT/>
15. Cavanna, A.E., Trimble, M.R.: The precuneus: a review of its functional anatomy and behavioural correlates. Brain 129 (2006) 564-583
16. Worden, M.S., Foxe, J.J., Wang, N., Simpson, G.V.: Anticipatory biasing of visuospatial attention indexed by retinotopically specific alpha-band electroencephalography increases over occipital cortex. J Neurosci 20 (2000) RC63
17. Makeig, S., Bell, A.J., Jung, T-P, Sejnowski, T.J.: Independent component analysis of electroencephalographic data In: D. Touretzky, M. Mozer and M. Hasselmo (Eds). Advances in Neural Information Processing Systems 8:145-151 (1996).
18. Makeig, S., Jung, T.P., Bell, A.J., Ghahremani, D., Sejnowski, T.J.: Blind separation of auditory event-related brain responses into independent components. Proc Natl Acad Sci U S A 94 (1997) 10979-10984
19. Makeig, S., Enghoff, S., Jung, T.P., Sejnowski, T.J.: A natural basis for efficient brain-actuated control. IEEE Trans Rehabil Eng 8 (2000) 208-211