Blind Separation of Spatio-Temporal Synfire Sources and Visualization of Neural Cliques

Hilit Unger , Yehoshua Y. Zeevi

Department of Electrical Engineering, Technion - Israel Institute of Technology, Haifa 32000, Israel

Abstract

A dominating paradigm in neuroscience attributes components of perception and behavior to synchronous spatio-temporal activities of subsets of neurons within neural networks – the so-called Synfire chains. Synfire chains cohere to generate *neural cliques* within the simultaneously active Synfires. The present study is concerned with blind separation of Synfire activities and identification of neural cliques. Assuming stationarity and, to a first approximation, linearity, we extend the Blind Source Separation (BSS) technique to the spatio-temporal domain, to deal with dynamic signals, and apply it on our analysis of neural networks. To demonstrate the concept of dynamic BSS, we first apply it in a relatively simple physical example of separation of dynamic reflections. In the latter case, the physics of the problem is well understood and the assumptions of linearity and stationarity are valid. In the case of neural networks, the problem is much more involved, and the assumptions of linearity and stationarity are not fully justified. Nevertheless, we demonstrate that the basic approach yields interesting insight into the function of complex neural activity. To better approximate the function of neural activity, we also investigate the effect of non-linearities on blind separation of neural activity.

Key words: neurons, neural network, Synfire chains, cliques, ICA, BSS, sparse ICA, SCA, dynamic sources, spatio-temporal, 3D ICA, Wavelet Packet Transform.
PACS:

1 Introduction

In recent years, there is concerted effort to decipher messages signaled simultaneously by the spatio-temporal firing patterns, typical of the dynamics

Email addresses: hilitg@tx.technion.ac.il (Hilit Unger), zeevi@ee.technion.ac.il (Yehoshua Y. Zeevi).

of massively connected neural networks. The 'music' of such neural activity sounds at first as a cacophonous neural cocktail-party. However, as the Synfire doctrine indicates, the global asynchronous activity masks the synchronous spatio-temporal activities of subsets of neurons that constitute functional Synfire chains. Our hypothesis is that Synfire chains cohere to generate more complex neural patterns of activity for which we coin the term 'neural cliques'. A clique is a more complex function of perception, concerning the identification of a specific concept. Understanding cliques may be a useful tool for understanding higher brain functions. The concept of cliques may also enhance further research concerning the capacity of ANN. It is therefore important to develop new techniques for the analysis of such activity, within the context of analysis of spatio-temporal firing patterns recorded by optical [1] or electrical [2] techniques from a massive number of neurons, as well as for the analysis of simulated large scale neural networks [3]. We examine the applicability of the Blind Source Separation (BSS) technique for this interesting and challenging problem.

Most of the research devoted to the problem of Blind Source Separation (BSS) has been concerned with either one-dimensional functions of time or static with images ([4],[5]). Yet, many physical systems generate linear mixtures of dynamic data sets, and it is therefore desirable to extend the BSS techniques to functions of both time and space. In biomedical applications such as those encountered, for example, in functional MRI, one is interested in the dynamic activity of specific loci of the brain. Another application concerns video sequences acquired through a semireflective medium and thereby masked by superimposed reflections, such as a movie recorded through a glass windshield of a moving vehicle, or the visual environment observed through the canopy of an airplane. The video sequence acquired in such cases can be represented as a three-dimensional (volumetric) data cube, in which spatial images are stacked along a third axis (Fig. 1).

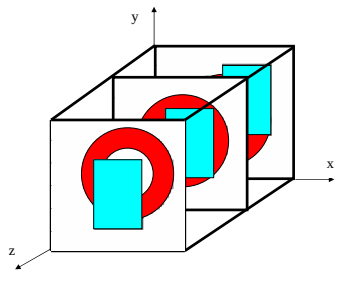

Fig. 1. Video sequence considered as a volumetric (cubic) date set. Shown is a data set comprised of three consecutive frames obtained from the sequence. Note the relative movement of the objects.

2 Independent Component Analysis (ICA)

In problems requiring BSS, an N-channel sensor signal $\{x_i\}_{i=1}^{N}$ is generated by M unknown scalar source signals $\{s_j\}_{j=1}^{M}$, linearly mixed by an unknown constant $N \times M$ mixing matrix, where each source or mixture is a function of the independent variables $\{\xi_k\}_{k=1}^{K}$.
In matrix notation, the N-dimensional vector of mixtures, $\mathbf{X}$, is equal to the product of the $N \times M$ mixing matrix by the M-dimensional sources vector, $\mathbf{S}$:

$$\mathbf{X}(\xi_1, \xi_2, \ldots, \xi_K) = \mathbf{A} \cdot \mathbf{S}(\xi_1, \xi_2, \ldots, \xi_K). \tag{1}$$

Under the assumption that the sources are statistically independent, the BSS method yields an estimate of $\tilde{A}$, the unknown mixing matrix, without prior knowledge of the sources and/or the mixing process. The sources are recovered (up to permutation and scale) by using an inverse of the estimated mixing matrix, provided it exists:

$$\tilde{\mathbf{S}}(\xi_1, \xi_2, \ldots, \xi_K) = \tilde{\mathbf{W}} \cdot \mathbf{X}(\xi_1, \xi_2, \ldots, \xi_K) = \tilde{\mathbf{A}}^{-1} \cdot \mathbf{X}(\xi_1, \xi_2, \ldots, \xi_K), \tag{2}$$

where $\tilde{\mathbf{W}}$ is the estimated 'unmixing' matrix.

2.1 Sparse Component Analysis (SCA)

It has been shown that when sources are sparse, they can be easily recovered from their linear mixtures using simple geometrical methods [6],[7],[8]. This is based on the observation that whenever sources are sparse, there is a high probability that most data points of the mixtures will result from the contribution of only one source. Using such geometrical methods, one can relax the condition of statistical independence. If we plot the N-dimensional scatter plot wherein each axis represents one of the mixtures, a co-linear cluster emerges for each subset of mixtures' data points that are contributed by one source only. Recall that the projection onto the space of sparse representation decouples the contributions of the sources to most of the mixtures' data points; this is the essence of the implementation of sparsity in the context of BSS. It can be shown that the coordinates of the vectors representing the centroids of these clusters correspond to the columns of the mixing matrix A (Fig. 2).

The simplest way to estimate the mixing matrix is to calculate the orientations of the clusters and select the optimal M angles from the histogram of angles. Another algorithm projects the data points onto a hemisphere, then uses clustering (such as Fuzzy C-means) in order to recover the orientations.

Another related maximum-likelihood-based approach is the well-known Infomax [9],[10].

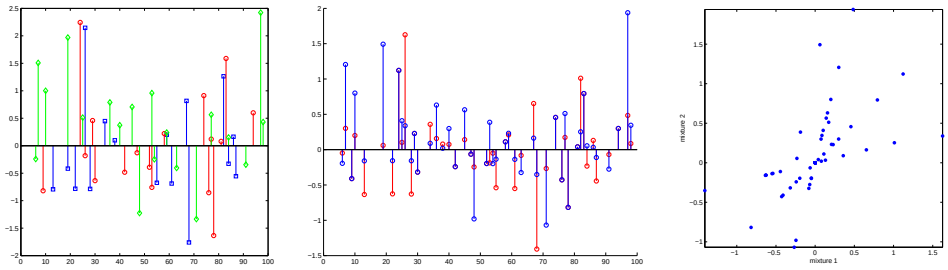

(a) Three sources (b) Two mixtures (c) scatter plot

Fig. 2. Three sparse sources mixed into two mixtures. The orientations of the three co-linear centroids of the sub-clusters in the scatter plot correspond to the columns of the mixing matrix.

2.2 Sparse Decompositions

2.2.1 Overcomplete Representations

Natural images and image sequences are not typically sparse. In order to exploit the methods previously described, we have to apply a transformation that yields a sparse representation of the signals. It has been shown that for a wide range of natural images, smoothed derivative operators yield a good, and even optimal, sparsification results [11]. However, an overcomplete representation obtained, for example, by the Wavelet Packet transform (WPT, proposed in [7]) matches better the specific structure of a given set of images and thereby yields better sparsification. The latter, in turn, facilitates and improves the estimation of the mixing matrix. The local nature of the wavelet-type transforms can highlight specific features of distributions (such as the distinct orientations in a scatter plot) based on a subset of data points of the transformed signal. Such highly structural distributions are not clearly present in the highly correlated original, non-transformed, signal (Fig. 3).

2.2.2 WP transform

The Wavelet Packet family consists of the triple-indexed family of functions:

$$\begin{aligned} &\varphi_{jnl}(\xi_i) = 2^{j/2}\varphi_n(2^j\xi_i - l) \ , j, l \in Z, \ n \in N \\ &\xi_1 \equiv x, \ \xi_2 \equiv y, \ \xi_3 \equiv t \\ &\boldsymbol{\varphi}_{jnl} = \prod_i \varphi_{jnl}(\xi_i) \end{aligned} \tag{3}$$

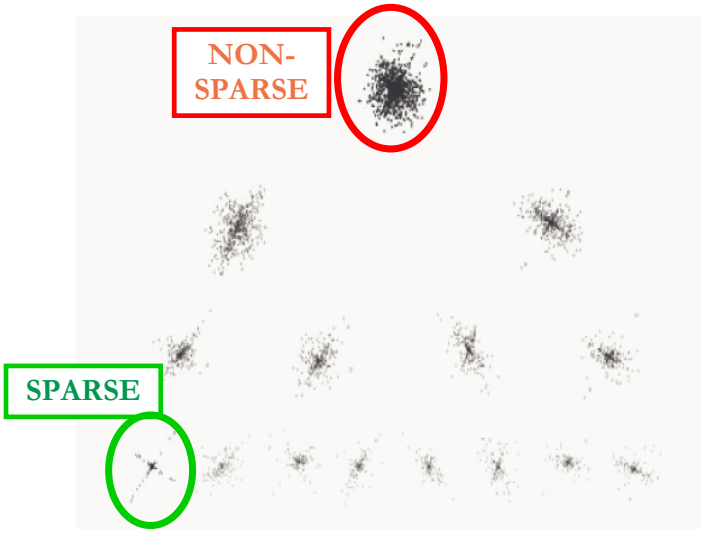

Fig. 3. Wavelet Packets (WP)-based multinode data structure of a 1D signal (adopted from [7]). The upper node is not sparse, while the leftmost lower node highlights the orientations of the two centroids very vividly.

According to the formalism of the Wavelet Packet transform, a signal is recursively decomposed into its approximation (L) and detail (H) components. In the case of 2D signals, using separable wavelets, the signal is decomposed into its approximation and vertical, horizontal and diagonal detail sub-images. For three-dimensional data cube, the signal is decomposed into 8 sub-volumes (Fig. 4).

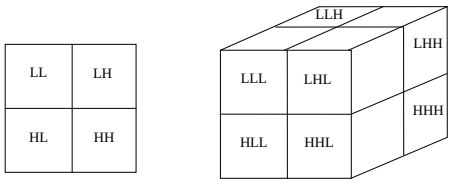

Fig. 4. WP decomposition. Left: 2D decomposition. Right: 3D decomposition.

We used a separable transformation, for the sake of simplicity, by transforming rows first, then columns and then time (depth) axis. Nonseparable wavelets offer important advantages in that they are inherently endowed with more degrees of freedom that can be exploited in their design. However, nonseparable wavelets are much more complex to deal with [12] and their application in the context of sparsification is therefore beyond the scope of this study.

2.2.3 Source Separation using the WPT

After the mixture signals are decomposed into WP tree nodes using the WPT [7], we search for the sparsest node (Fig. 3). A quality criterion that assigns high values for sparse nodes and lower values for less sparse nodes is computed

for every node. Common choices for such quality criteria are entropy or global distortion. The best node (or the top few nodes) is chosen and used as input data for the geometrical BSS algorithm.
Using the WPT has another advantage: because of downsampling in the process of the signal transformation, the number of data points in each node is significantly smaller than the number of data points in the mixture signals. This reduction in the number of data points speeds up the processing.

2.3 BSS of Dynamic Reflections

In order to test our method of dynamic (spatio-temporal) BSS, we explore the example of separation of dynamic reflections: a situation wherein a virtual (reflected) image is superimposed on a dynamic visual scene. To this end we extend the study concerned with separation of reflections from static images [11] to the case of 3D dynamic images, i.e. video. The problem of dynamic reflections is a good test case since the physics of it is well understood and the assumptions of both linearity and stationarity are valid.

In the context of separation of reflections, the BSS problem usually reduces to the case of M=2 sources. The observed mixture is then given by

$$x(\xi 1, \xi 2, t) = a_{11} s_1(\xi 1, \xi 2, t) + a_{12} s_2(\xi 1, \xi 2, t), \tag{4}$$

where x, s_1 and s_2 are dynamic images, usually acquired as video sequences. It is assumed here that the dynamics of the image and of the superimposed reflections are limited to planar translation of rigid bodies. The more difficult problem of non-planar motion and rotation as well as non-rigid distortions

are beyond the scope of this paper, and will be dealt with elsewhere. Likewise, the coefficients a_{11} and a_{12} are assumed to be constant, approximating spatial invariance and linear mixing [11].

Since the reflected light is polarized, by using a linear polarizer, the relative weights of the two mixed video sequences can be varied to yield N different mixtures of the form:

$$x_n(\xi 1, \xi 2, t) = a_{n1} s_1(\xi 1, \xi 2, t) + a_{n2} s_2(\xi 1, \xi 2, t) \; : \; n = 1, \ldots, N \, . \tag{5}$$

Thus, we can use two or more video sequences obtained with different polarizations and separate objects and reflections. Simulation results are shown in Fig. 5. The results (video movies) of an experiment with real data are available on the web at `http://visl.technion.ac.il/~hilitg/dynamic_scenes.htm`.

(a) (b)

Fig. 5. Simulation of blind separation of dynamic (moving) image from a superimposed reflection. (a) Shown are four frames from one mixture (top) and four frames extracted from the corresponding sequence of the recovered source (bottom). (b) Data cube of one mixture. The arrows trace the trajectories of movements of the image and reflection, relatively to a stationary background.

3 BSS of Neural Cliques

New optical [1] and electrical [2] imaging techniques for simultaneous recording of activity of populations of neurons in the brain tissue have been developed in recent years. Whereas traditional methods for detection of action potentials in neurons were limited to a small number of neurons, it is now possible to record massive neural activity with spatial resolution of a single cell and a temporal resolution of a single action potential. It is therefore important to develop new techniques for the analysis of such activity.

The study of large populations of neurons enables the identification and analysis of neural phenomena such as Synfire chains [13]: waves of synchronous neural activity that propagate through biological neural network (Fig. 6). Since the messages of concepts signaled by Synfire chains are often associated with each other, examining such spatio-temporal patterns of firing neurons led us to the assumption that simultaneous activity of Synfire chains coheres into representation of components of more complex functions of behavior and perception in the form of 'neural cliques'. A clique is thus a synchronous spatio-temporal spiking activity across and within several different Synfire chains.

To highlight the meaning of the clique, consider the example of recognition of different digits. Either when we hear the word "one" or when we look at an image of the digit "1" (Fig. 7), our brain must identify the same concept signaled by neural patterns associated with this concept. The Clique is, in a way, higher order Synfire and neural activity within a neural network of a critical size. In sorting the dynamic patterns of neural activity, it is useful to consider a conceptual matrix of Synfire chains and cliques. Our goal in this paper is to illustrate these concepts by means of blind separation of spatio-temporal neural activities.

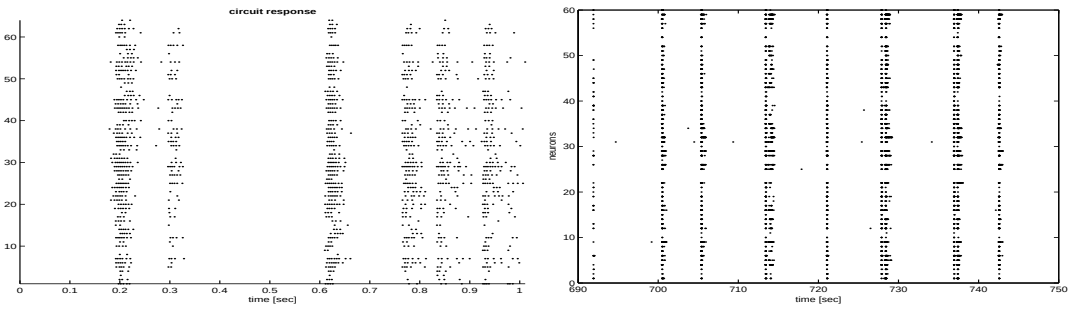

Fig. 6. A raster plot of Synfire chains in a biological neural network. Waves of synchronous activity and between them quiescent periods. Left, simulated data. Right, spontaneous activity of a biological neural network recorded from multi-electrode array (Adopted with permission from S. Marom and D. Eytan.)

A neural network maps, at each moment, sensory or other inputs impinging on it, into meaningful neural activity. We assume, to a first order approximation, that every spiking pattern is a linear combination of cliques, with the addition of clutter and some noise. To better understand the concept of cliques in the context of spatio-temporal neural network activity, recall the representation of spatio-temporal data as a cubical data set (Fig. 1). Here each frame corresponds to a slice along the time axis of duration $\triangle t$. A clique then corresponds to correlated pattern of activity of two or more such slices of duration $T > \triangle t$.

To provide some intuitive insight into the analysis of neural cliques by means of BSS technique, we generate data using CSIM circuit-tool; a simulator for neural networks [14]. The network connectivity is randomized and then paths for each clique are randomly chosen and selectively strengthened. Each clique is stimulated by at least two different input neurons (Fig. 7).

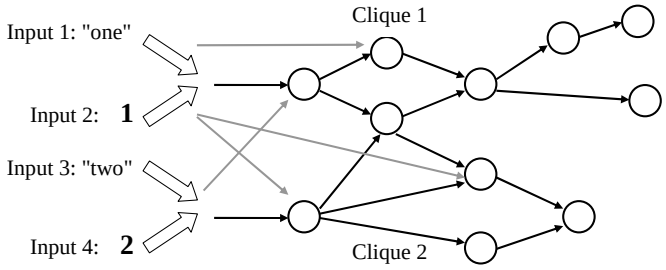

Fig. 7. Simulation of a two-cliques scenario, each being characterized by a different spatio-temporal sequence. Black arrows indicate synaptic chains associated with Synfire activities related to the clique, while gray (thin) arrows are samples of unrelated network connections.

The discrete output spiking activity is converted into continuous analog signal which, in turn, is quantized for further computation. Since every input neuron also stimulates other neurons except those participating in the clique, the network also exhibits background activity that is not related to cliques (gray arrows in Fig. 7).

Using our spatio-temporal BSS approach, we then project the data deduced from slices onto a scatter plot wherein each axis represents activity in one mixture slice. Each point then represents the activity of a neuron at a specific time in the slice.
Investigating the mutual activity of two slices, one often observes that two slices that are selected within the duration of co-activation do not necessarily exhibit coincidence of spike activity. In fact, the spatio-temporal activity may be almost exclusively restricted to only one slice. Under these circumstances, the distribution of activity projected onto the scatter plot will form either a vertical or a horizontal cluster (Fig. 8). The diagonal cluster corresponds to the coincident activity.

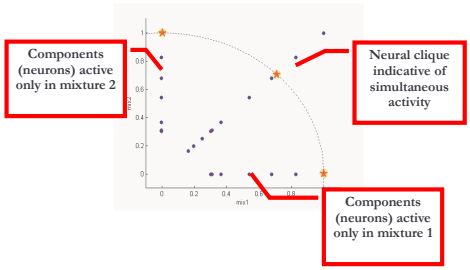

Fig. 8. A scatter plot of two slices of data.

In order to see how cliques appear on the scatter plot, we simulated a network and tested its response to two inputs that activate the same clique. We then projected the data points onto a 2D scatter plot. To better highlight the clique we also show a 3D visualization of the same scatter plot, in which the height of each point (z-axis) represents the number of data points at each coordinate. The spiking patterns and results are shown in Fig. 9. The 3D scatter plot shows that the main diagonal has relatively many data points which constitute a clique. We highlight those data points by a piecewise-linear trajectory for clarity.

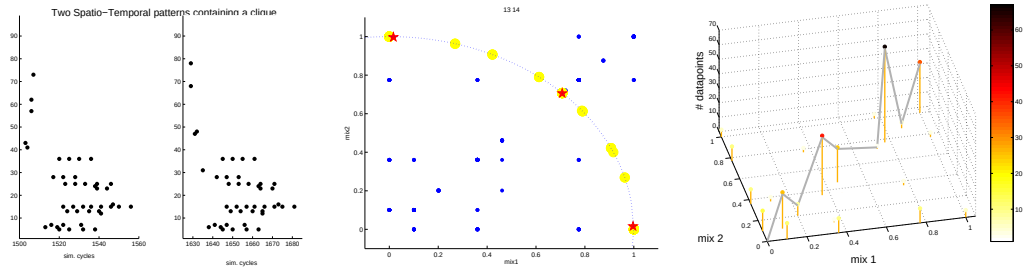

(a) (b) (c)

Fig. 9. (a) Comparison of two spiking patterns. (b) A 2D scatter plot. (c) A 3D scatter plot of the same clique. Note that the diagonal elements are much higher than the other scatterplot data points.

To further illustrate the emergence of a clique from the co-activation of 3 different activity slices, we depict in Fig. 10 a 3D scatter plot wherein the color and size of the dot in the plot depends on the number of data points

along each coordinate. The main diagonal (red dots) shows data points of the emerging clique, while the other points show points not related to the clique activity.

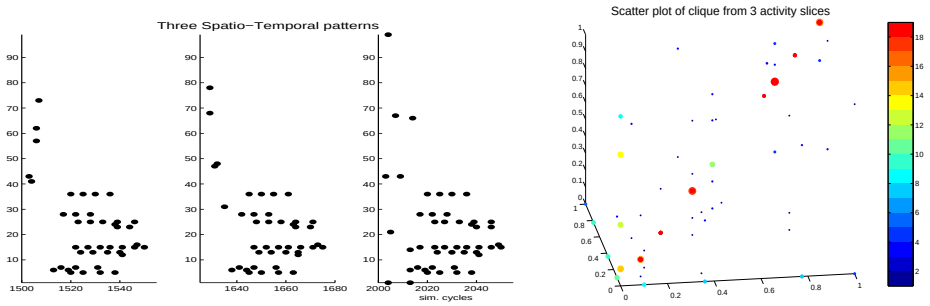

Fig. 10. Left: Comparison of three spatio-temporal spiking patterns. Right: The emergence of a clique illustrated in a 3D scatter plot.

It should be stressed, however, that unlike the example of blind separation of 'sources' out of video data, here we face a non-linear phenomenon that limits the power of ICA-type techniques. The projection of nonlinearly-mixed data onto the scatter plots generates distributions that are not yet well understood. Nevertheless, the formalism and approach of projecting the data onto a scatter plot is powerful in gaining some insight into the structure of non-linearly interacting sources or cliques.

4 Non-linearity of the neural network

Both biological and artificial neural networks process their inputs non-linearly [15]. We therefore examine the effect of a non-linear mixing process on a pattern of neural spikes. We consider a simple model for a neuron, with a shifted version of the parametric non-linear logistic function:
$f_\lambda(x) = \frac{1}{1+\exp(x/\lambda)}$, (Fig. 11).

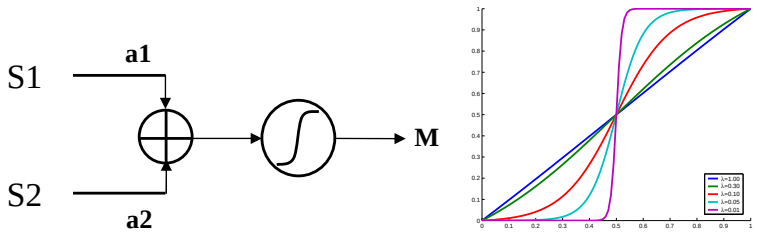

Fig. 11. (a) A model of a non-linear mixing of spatio-temporal spike patterns S1 and S2. (b) The non-linear logistic function is depicted for various values of the parameter λ.

In order to see the effects of non-linearity, we consider the mixing of two spatio-temporal spike patterns, characterized by having some mutual spikes within the two patterns (Fig. 12). We now test the effect of nonlinear processing of the

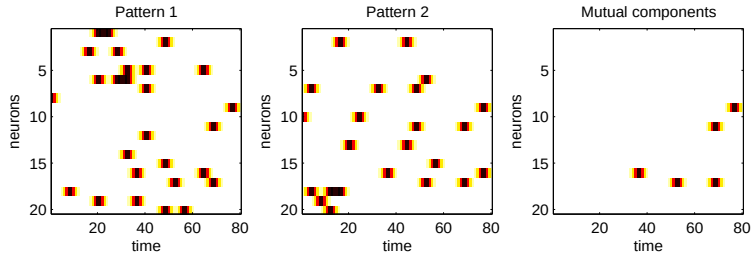

Fig. 12. Two spatio-temporal spike patterns and their mutual activity pattern. Patterns are shown as a 2D image for clarity, where the x-axis is time and y-axis is the neuron number.

mixtures using different weights of the linear mixing and for different values of the parameter λ that changes the nonlinearity from 'soft' to 'hard'. Results such as those depicted in Fig. 13 indicate that the mutual components are most salient when the mixing coefficients, i.e. the synaptic weights, are roughly of the same strength (rightmost column). The hard non-linearity (bottom row) then preserves only the highest values of the signal. Hard limiter (i.e. non-linearity) quantizes the signal to an approximately 1-bit signal.

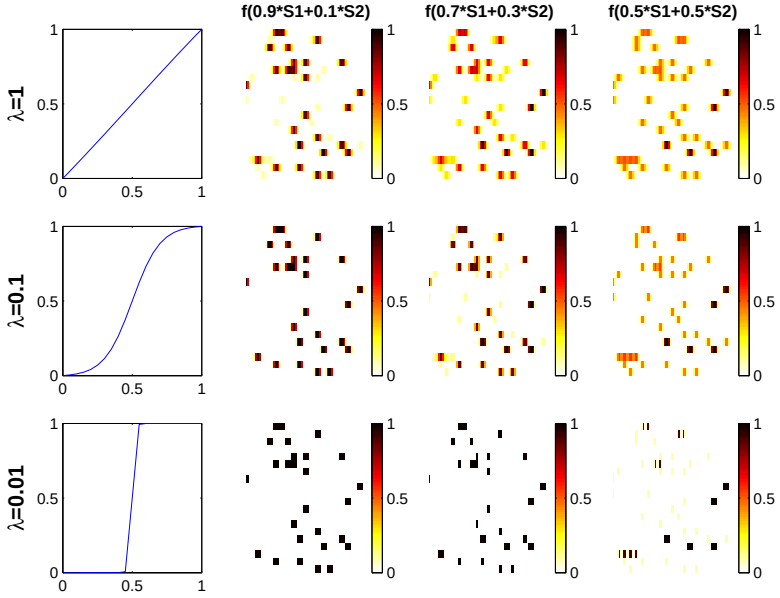

Fig. 13. Post-nonlinear mixture of two spike patterns for different mixing coefficients and non-linearity constants.

It is interesting to compare the correlation coefficient between the mutual components of the patterns (right pattern in Fig. 12) and the different outputs of mixing process:

correlation	f(0.9*S1+ 0.1*S2)	f(0.7*S1+ 0.3*S2)	f(0.5*S1+ 0.5*S2)
$\lambda = 1$	0.4477	0.5225	0.5692
$\lambda = 0.1$	0.4280	0.5136	0.6567
$\lambda = 0.01$	0.3947	0.4949	0.7671

Table 1
A comparison of correlation coefficients between the mutual components of neural patterns and the output of a non-linear neuron model.

When the mixing process is linear (Fig. 13 top row), both patterns can be identified in the output pattern, according to the linear mixing weights. Hard non-linearity emphasizes components of activity with higher amplitudes.
As the mixing coefficients become symmetric (right column of Fig. 13), the mutual components become more salient and the correlation to the mutual components become larger. The highest correlation is achieved when the interaction is highly non-linear and for symmetric mixing, i.e. equal coefficients ($a_1 = a_2 = 0.5$). In this case, the mutual components are dominant while the non-mutual components almost disappear.

5 Discussion

Identifying coherent sources of neural activity in neural networks is a challenging and demanding task. Our understanding of neural activity still lacks a clear definition of what type of activities constitute coherent sources. Here we coin the term 'neural cliques', where sources are defined by the conceptual meaning of the activity, i.e. each source is associated with a concept in our brain. The cliques can be composed by several Synfire chains and can be represented by a matrix associating clique activity with Synfire activity.

We extend the SCA approach [8] to three-dimensional sources, and show good separation results for the simple case of removal of reflections from a video sequence. The extension to 3D problems is rather simple, and it broadens the kind of problems dealt with until now to dynamic sources, which are usually the sources generated by real-world systems. However, unlike the physics of separation of superimposed reflections, which can, to a good approximation, be considered linear and stationary, the neural cliques separation is necessarily non-linear, and most likely non-stationary. Nevertheless, as we have demonstrated here, the novel approach of using BSS techniques in isolation of the fingerprints of coherent neural activity, can be instrumental in highlighting the functions of biological neural networks. It may be also instrumental in studies attempting to reverse engineer the structures of linear skeletons of such networks using spatio-temporal spiking activity. We show how practical is it to use scatter-plots in order to visualize the clique and find coherent mutual

activity in neural networks. We conclude by examining how various degrees of post non-linearity (PNL) processing affects the output of a simple neuron. We further show that hard non-linearity and symmetric neural connections emphasize mutual components in neuronal patterns.

Acknowledgement: Research supported in part by the Ollendorff Minerva center, by the HASSIP Research Network Program HPRN-CT-2002-00285, sponsored by the European Commission and by the Fund for Promotion of Research at the Technion.

References

[1] D. Smetters, A. Majewska, R. Yuste, Detecting action potentials in neuronal populations with calcium imaging, Methods 18 (2) (1999) 215–221.

[2] G. Shahaf, S. Marom, Learning in networks of cortical neurons, J. of Neuroscience 21 (22) (2001) 8782–8788.

[3] Y. Ikegaya, G. Aaron, R. Cossart, D. Aronov, I. Lampl, D. Ferster, R. Yuste, Synfire chains and cortical songs: Temporal modules of cortical activity, Science 304 (2004) 559–564.

[4] Special issue on Independent Components Analysis, Vol. 4.

[5] G. D.Brown, S. Yamada, T. J. Sejnowski, Independent component analysis at the neural cocktail party, Trends Neurosci 24 (1) (2001) 54–63.

[6] M. Zibulevsky, B. A. Pearlmutter, Blind source separation by sparse decomposition in a signal dictionary, Neural Comp. 13 (4) (2001) 863–882.

[7] P. Kisilev, M. Zibulevsky, Y. Y. Zeevi, A multiscale framework for blind separation of linearly mixed signals, J. Mach. Learn. Res. 4 (2003) 1339–1363.

[8] A. Bronstein, M. Bronstein, M. Zibulevsky, Y. Zeevi, "Unmixing" Tissues: Sparse Component Analysis in Multi-Contrast MRI, submitted to ICIP 2005.

[9] A. J. Bell, T. J. Sejnowski, An information-maximization approach to blind separation and blind deconvolution, Neural Computation 7 (6) (1995) 1129–1159.

[10] J. Cardoso, Infomax and maximum likelihood for blind source separation, IEEE Signal Processing Letters 4 (4) (1997) 112–114.

[11] A. Bronstein, M. Bronstein, M. Zibulevsky, Y. Y. Zeevi, Separation of reflections via sparse ICA, in: ICIP03, 2003, pp. 313–316.

[12] D. Stanhill, Y. Y. Zeevi, Two-dimensional orthogonal wavelets with vanishing moments, IEEE Transactions on Signal Processing 44 (9) (1996) 2579–2590.

[13] M. Abeles, Corticonics, neural circuits of the cerebral cortex, Cambridge University Press, 1991.

[14] Http://www.lsm.tugraz.at/csim/index.html.

[15] S. Haykin, NEURAL NETWORKS A Comprehensive Foundation, Prentice Hall, 1994.