Preprint version. Link to original publication: International Conference on Computational Advances in Bio and Medical Sciences (ICCABS), IEEE, 2012 E-ISBN : 978-1-4673-1319-3 Print ISBN: 978-1-4673-1320-9 http://dx.doi.org/10.1109/ICCABS.2012.6182646

Signal extraction from movies of honeybee brain activity by convex analysis.

Martin Strauch^{*†‡}, Julia Rein[†], C. Giovanni Galizia[†] *Bioinformatics and Information Mining, [†]Neurobiology, University of Konstanz, 78457 Konstanz, Germany [‡]Martin.Strauch@uni-konstanz.de

Abstract—Calcium-imaging enables us to record movies of brain activity from the antennal lobe, a region of the honeybee brain responsible for processing odor information. Here, we present a matrix factorisation framework to automatically detect the neural units in this region and to accurately estimate their signals. Based on a non-negative mixture model, the algorithmic approach is to construct a convex cone that contains the data. The generating vectors of the cone are the purest, least mixed timeseries from the movie and serve as basis vectors for the matrix factorisation. We show that vectors selected in this way correspond to the biological signals and evaluate the method on both artificial and biological data.

Keywords-imaging, matrix factorisation, non-negativity

I. INTRODUCTION

Olfaction, the sense of smell, is a formidable task for a sensory processing system that needs to map the highly multidimensional world of chemical properties to a neural code representing odor identity [12]. The olfactory code is based on neural units called glomeruli. In our model organism, the honeybee *Apis mellifera*, a particular odor elicits a characteristic activation pattern in the 160 glomeruli of a brain region called the antennal lobe (AL) [8].

These activation patterns can be observed by *in vivo* optical imaging using calcium-sensitive fluorescent dyes (see e.g. [6] [5] [13]). Here, we employ a ratiometric dye (Fura2-dextran) where the signal is the ratio of light emitted after excitation at wavelengths 340nm and 380nm.

The datasets are movies of brain activity over time (Figure 1a) recorded with a CCD camera through a confocal microscope. The experimental protocol (Figure 1b) is a concatenation of measurements during periods of odor presentation and periods of idleness. Glomeruli respond differentially to a given odor and they also have individual background activity in the idle state. Glomerulus-specific response properties are the basis for detecting glomeruli by observing correlations between pixel-timeseries from the movie [4].

Biological questions concern e.g. changes in the odor response after the odor has been learned through association with a reward [13], or reverberations of odor response patterns during the idle state as a form of short-term memory for odors [6]. Such effects can be subtle and thus signal processing needs to accurately detect glomerulus positions and to estimate their true signals from noisy measurements.



Figure 1. **a)** Frontal view onto a model of the honeybee AL (modified from [2]) and schematic for a calcium-imaging movie covering a part of the view (varies depending on focal depth). The circled glomeruli (labelled as 17 and 33 after [7]) exhibit differential responses to the odor stimulation during the interval marked with a black bar. **b)** Experimental protocol underlying the movie: periods without any stimulation (idle-state activity) alternate with periods of odor stimulation.

Traditionally, signal processing in this field has been performed in a semi-automatic fashion, using laboratoryspecific scripts to apply spatial filters and to visualise correlations between pixel-timeseries. Signals are estimated by averaging over pixel-timeseries within a radius around the manually selected glomerulus centroid [4] [5]. This approach may fail to detect all glomeruli, it is not optimal with respect to estimating the true glomerulus signals, and it does not allow for automated processing, e.g. in real-time systems.

Here, we propose an algorithm for automated signal extraction from imaging movies. We introduce a non-negative mixture model, such that glomerular signals occur either as pure, unmixed signals (plus noise), or as additive mixtures of source signals in regions of contact between glomeruli.

Adapting a concept known from the remote sensing literature [10] [9], we propose a convex cone fitting approach where the generating vectors of a convex cone containing the data are sought to serve as basis vectors in a matrix factorisation framework (Section II). They identify the purest, least mixed timeseries in the movie. On both artificial and real biological data we show that these vectors lead us to the glomerular signals and how they can be used for movie denoising with the matrix factorisation (Section III).

II. A CONVEX CONE FITTING ALGORITHM

A. Matrix factorisation framework for imaging movies

We consider a movie matrix A with m timepoints and n pixels. In this work, $n = 160 \times 120$, $m \approx 4000$, where m could be larger depending on the experimental protocol.

Our goal is to provide an interpretable rank-k approximation to A. For signal estimation and denoising, we would like to factorise A into a matrix T whose k columns are timeseries, and a matrix S whose k rows are images, where the images should be sparse and the timeseries should correspond to the glomerular signals.

$$A^{m \times n}: \quad A_k = T^{m \times k} S^{k \times n} = \sum_{r=1}^k T_{Ir} S_{rJ} \qquad (1)$$

We denote as A_{ij} the element at the intersection of the *i*th row and the *j*th column. A_{Ij} is the *j*th column of A and corresponds to a pixel, or more precisely, a pixel-timeseries vector of length *m*. For ease of notation, we also refer to the rows of S as $\mathbf{s}_{(r)}$ and the columns of T as $\mathbf{t}^{(r)}$.

B. Preprocessing with PCA

We zscore-normalised the n pixel-timeseries, i.e. we subtracted the mean of each pixel-timeseries and divided by the standard deviation. As a first approach to (1) and as a general preprocessing we then performed Principal Component Analysis (PCA) [11] on the movie A.

PCA provides optimal dimensionality reduction to a rank-k matrix with respect to common error measures. For glomerulus detection, it is relevant that PCA optimally preserves the covariation (between timeseries), i.e. $||P_k^T P_k - A^T A||_{F_r}$ is minimal for a given k, assuming that P_k contains the top-k principal components ([11], chap. 3.2).

Cast into the matrix factorisation framework (1), the k principal component images are in matrix S and the corresponding loadings per timepoint in matrix T (or vice versa). In Figure 2, we give examples for the images in S.

PCA is a beneficial preprocessing for imaging movies. Subsequent algorithms can be carried out efficiently on a small matrix, and glomerular signals are concentrated in the top principal components (see Figure 2). Transient signals will inevitably contribute to the variance, and the principal component is the variance-maximising projection [11]. I.e., variance is accumulated in the top components and by discarding components with lower eigenvalues (and no signal) we can improve the signal-to-noise ratio.

If only the top-k principal components are sought, PCA can be computed with a computational complexity of $\mathcal{O}(mnk)$ [18], which is in the same order of complexity as the convex cone fitting (Algorithm 1). I.e, the main aspect for this work is noise reduction by PCA. Note that, if necessary, fast approximate PCA by pixel importance sampling [17] can substantially reduce n and thus also the computational load for our method (PCA + Algorithm 1).



Figure 2. Principal component images (rows of S) of the movie from Figure 1 ranked by eigenvalue in decreasing order. Components 1-2 contain background, 3 - ca. 30: mixtures of background and glomeruli (black and white circles). Components with lower eigenvalue contain less structure. The color scale ranges from white (negative) to black (positive values).

C. Convex cone fitting

The principal components are useful in a technical sense, but they do not separate individual glomeruli (Figure 2), which motivates an alternative approach to (1).

Recall that we measure light. Neither can light intensity be negative, nor is it reasonable in an anatomical sense to speak of negative glomeruli in S. We thus assume that nonnegative combinations (coefficients in S^{0+}) of the basispixel-timeseries in T can explain the movie A up to the residual noise in N.

$$A = TS^{0+} + N \tag{2}$$

A pixel-timeseries contains the additive mixture of one or more glomerular signals plus noise. Mixtures can occur in case of light scatter (from neighbour glomeruli) at the fringes, but usually not in the middle of a glomerulus.

Regarding (2), we leverage on a concept from convex analysis: The columns of T generate a convex cone pointed at the origin, and this cone contains parts of A. We call a set of vectors V a convex cone, if $\alpha_1 \mathbf{v}_1 + \alpha_2 \mathbf{v}_2 \in V$ for non-negative α_1, α_2 , and any $\mathbf{v}_1, \mathbf{v}_2 \in V$.

The extreme vectors of a convex cone (on the boundary of the hull) are by definition those that cannot be expressed as conic (non-negative) combinations of the other vectors in the cone. Given the set of extreme vectors of the cone, we can reconstruct all others by conic combination of the extreme vectors. Different generators exist, but the set of extreme vectors is the minimal generator for the cone. [14] [3]

Our goal is now to choose T such that the convex cone generated by it contains (almost) all columns of A, and we do so by selecting extreme vectors from A into T.

By definition, the extreme vectors cannot be expressed by conic combination, i.e. they are not mixtures with nonnegative coefficients, but pure signal sources. Given the model in (2) and the assumption that pure, unmixed pixeltimeseries exist (neglecting N), these pure timeseries are the extreme vectors of A (cp. [1] [10]). Our approach to finding these extreme vectors is a greedy column selection strategy that corresponds to fitting a convex cone to the data: Algorithm 1. We sequentially select r = 0, ..., c-1 pixel-timeseries $\mathbf{t}^{(r)}$ into matrix T, at each step selecting the column that is least explained by conic combinations of the previous basis vectors. S contains the corresponding non-negative spatial mappings.

The residual matrix is formed by subtracting the "conic contribution" $\mathbf{t}^{(r)}\mathbf{s}_{(r)}^{0+}$ of the selected column vector $\mathbf{t}^{(r)}$, where the row vector $\mathbf{s}_{(r)} = A^T \mathbf{t}^{(r)}$, and $\mathbf{s}_{(r)}^{0+}$ is derived from $\mathbf{s}_{(r)}$ by setting negative entries to zero. Subtracting the conic contribution removes what can be explained by $\mathbf{t}^{(r)}$ with non-negative coefficients.

The maximum column norm in the residual matrix identifies the column that is least explained by the cone, and it is the locally optimal choice for the next extreme vector, as an extreme vector cannot be expressed by conic combination.

A natural, yet computationally expensive, initialisation is the column with the maximum conic contribution (in terms of $\|\mathbf{t} \mathbf{s}^{0+}\|$). In practice, we estimated an extreme column by finding the largest distance from a random start column.

For illustration, we apply Algorithm 1 to the movie from Figure 1a. Generally, we set $c \ge \#$ visible glomeruli ≈ 30 (see model in Figure 1a). In Figure 3a, we show the top-10 rows of matrix S. We can obtain a compact summary of all c components in S by forming the induced clustering in Figure 3b: each pixel is assigned to the component for which it has the highest coefficient.

Glomeruli form clusters, while inbetween the signal is discretised to the strongest contribution. Unlike the principal components, the timeseries in T can be interpreted as a particular glomerulus signal (cp. Figures 2 and 3a).

Computational complexity of Algorithm 1 is dominated by forming the $m \times n$ residual matrix c times. After dimensionality reduction with PCA, Algorithm 1 took between three and four seconds per bee (c = 40, Intel[®] 2140 CPU, 1.6 GHz). The computational load of the entire method thus depends mainly on PCA (see Section II-B).

D. Postprocessing to remove N

We perform Algorithm 1 in PCA space, where column selection is more robust against noise and outliers. Fulllength pixel-timeseries can be extracted from A at the positions indicated by Algorithm 1. These timeseries are not mixed with the other sources, but still affected by the noise N (2). N is then averaged out in postprocessing. We could, for example, average over neighbouring pixeltimeseries within a 3×3 square around the selected one (spatial average). Instead, we employ a "temporal average", i.e. we average over those pixel-timeseries that are most similar to the selected one, the estimate for the purest signal. Given the selected $\mathbf{t}^{(r)}$, we average over all pixel-timeseries from A that are more similar to $\mathbf{t}^{(r)}$ than to any of the other timeseries in T, giving rise to the average vector $\hat{\mathbf{t}}^{(r)}$.

Algorithm 1 [T, S] = Cone_fitting $(A^{(m \times n)}, c)$

for all $r \in [0, c-1]$ do

if (r == 0) then
initialisation of p: see main text
end if

$$\begin{split} \mathbf{t}^{(r)} &= A_{Ip}^{\{r\}}; \quad \mathbf{s}_{(r)} = A^{\{r\}T} \mathbf{t}^{(r)} \\ \mathbf{s}_{(r)}^{0+} &= \text{negative_to_zero}(\mathbf{s}_{(r)}) \\ T_{Ir} &= \mathbf{t}^{(r)}; \quad S_{rJ} = \mathbf{s}_{(r)}^{0+} \\ A^{\{r+1\}} &= A^{\{r\}} - \mathbf{t}^{(r)} \mathbf{s}_{(r)}^{0+} \quad //\text{form residual matrix} \\ p &= \operatorname{argmax}_p \left| \mathbf{A}_{Ip}^{\{r+1\}} \right| \qquad //\text{index of next column} \end{split}$$

end for



Figure 3. Applying Algorithm 1 to an imaging movie. **a**) Top-10 rows (images) of matrix S. **b**) Clustering of the image plane induced by S. Positions of the selected pixel-timeseries in T are marked with grey spots. **c**) Clustering of the image plane induced by \hat{S} (after temporal averaging). **d**) Example for two timeseries $\mathbf{t}^{(r)}$ (spatial average) and **e**) $\hat{\mathbf{t}}^{(r)}$ (temporal average) on a short subsequence of the movie including an odor presentation (interval marked by black bar).

This procedure is analogous to thresholding the images in S to cancel out small coefficients. For visualisation, we compute the corresponding $\hat{\mathbf{s}}_{(r)}$: For each component $\hat{\mathbf{s}}_{(r)}$ we set all coefficients to zero which are not involved in the average $\hat{\mathbf{t}}^{(r)}$. Then, the induced clustering (Figure 3c) leads to clearly separated glomeruli. Each pure signal estimate is based on the pixels of one color. The mixed signal pixels between the glomeruli are set to white as all coefficients are zero, i.e. the signal is not close enough to any of the sources.

We show examples for $\mathbf{t}^{(r)}$ (Figure 3d, spatial average), and the corresponding $\hat{\mathbf{t}}^{(r)}$ (Figure 3e, temporal average). The temporally averaged $\hat{\mathbf{t}}^{(r)}$ are smoother, as they are averaged over a larger number of pixel-timeseries. It appears that averaging over the maximum number of pixels, without including mixed pixels, is more readily achieved with temporal averaging, which we use throughout this work.

E. Related work

Practical data analysis in this field still involves manual steps (Section I). Two algorithmical approaches have been explored: Stetter et al. [15] reconstructed honeybee imaging data by bottom-up fitting of nonlinear model functions. Strauch&Galizia [16] applied Independent Component Analysis (ICA) to decompose imaging movies. The bottom-up approach does, however, not separate glomeruli, and ICA requires non-Gaussianity of the signals (which is hard to prove for glomeruli) and does not consider non-negativity.

The convex analysis approach has previously found application on hyperspectral data in remote sensing [1]. In this field, Gruninger et al. [9] and Ifarraguerri&Chang [10] proposed convex cone algorithms related to our method. In contrast to these approaches, we do not primarily aim at unmixing signal contributions to a pixel value. In fact, we discard mixed-signal pixels, and we use the convex cone to find the purest signals which are then postprocessed.

Assuming a mixture model is what distinguishes our method from discrete clustering techniques such as *k-means*.

III. EVALUATION

A. Artificial data

We tested our method on data constructed by additive mixing of realistic source signals ($\mu = 0, \sigma = 1$, shifted to be non-negative) as in (2). The "odors" dataset consists of a series of odor responses (Figure 4a), and the "idle" dataset consists of spontaneous background activity (Figure 4b).

We constructed imaging movies by assigning source signals to spatially smooth and partially overlapping glomeruli (Figure 4c). Additionally, we applied Gaussian noise with variable standard deviation σ .

We employed a correlation score as a measure for source recovery. Let $\hat{\mathbf{t}}^{(i)}$ be a signal estimate, and let $\mathbf{u}^{(j)}$ be a source timeseries. Based on the Pearson correlation coefficient $\rho(\mathbf{x}, \mathbf{y})$ between two timeseries, we define the correlation score: corr = $\frac{1}{n} \sum_{i=1}^{n} argmax_j \rho(\hat{\mathbf{t}}^{(i)}, \mathbf{u}^{(j)})$.

In Figure 4d, we show correlation scores for various noise levels (k = 16). On both datasets, correlation scores were high for $\sigma \leq 1$, but exhibited a decline at $\sigma = 1 - 1.3$.

The clusterings in Figure 4e show that the implanted signals were detected, while mixed-signal pixels were excluded from the clusters (as for the biological data in Figure 3c).

Attempting to improve noise tolerance, we smoothened images from the movie by convolution with a Gaussian kernel (width=7). This allowed successful signal recovery even at a high noise level ($\sigma = 2$, Figures 4df). The filter attenuates the noise N, such that more pixel-timeseries are similar to the purest signal and can contribute to the average.

With additional Gaussian filtering, clusterings appear smoother, but the filtering is not necessarily required for the typical noise level in the movies (compare: without filtering in Figure 3c; with filtering: Figure 5a:Bee1).

B. Biological data

To demonstrate practical applicability, we performed our method (including the Gaussian filtering) on movies of honeybee brain activity (protocol as in Figure 1b).

In Figure 5a, we show clusterings for three different bees that clearly reveal glomerulus positions. While brain anatomy is subject to individual variation, and experimental parameters determine the subset of visible glomeruli, we can still register landmark glomeruli from the clusterings onto the anatomical reference AL [7].

Using the rank-c approximation $A_c = \widehat{T}\widehat{S}$ to the original movie A (c = 50), we give an impression of inter-trial and inter-animal variability: Figures 5b and 5c show repeated responses to the odor nonanol in Bee1. Figures 5d and 5e show nonanol responses in Bee2.

Besides brain anatomy, also odor response patterns in the AL, a plastic neural circuit that can be shaped by experience throughout lifetime, are subject to variation. Nevertheless, spatial patterns, as well as timeseries shape are largely conserved within and between bees [8] (Figures 5b-e).

An important feature of our method is movie denoising. Due to the sparseness of \hat{S} , that contains predominantly glomeruli, the rank-*c* approximations in Figures 5b-e are clearly noise-reduced with respect to unprocessed or "traditionally treated" movies (cp. Figures 1a, 5f).

IV. CONCLUSIONS

Analysis of imaging movies from the insect AL is still largely based on semi-automatic methods that require human interaction to detect glomeruli (e.g. in [6] [5] [13]).

Here, we have presented a method for automatic glomerulus detection, signal estimation and movie denoising. In the future, it will serve as the basis for automated real-time processing of imaging data, e.g. to provide visual orientation for the experimenter. Accurate signal estimation by finding the purest pixel-timeseries with the convex cone approach will be helpful for analysing subtle effects or faint signals in background activity.



Figure 4. Artificial data with 16 sources (based on real data; for each: $\mu = 0, \sigma = 1$). a) Correlation matrix and first source signal for the odors dataset. b) Correlation matrix and first source signal for the idle dataset c) above: Spatial arrangement of source signals. Pure signals are in grey, mixed signals in black. White areas contain no signal. below: With Gaussian noise ($\sigma = 1$) added. d) Results of Algorithm 1 (plus temporal averaging). Correlation scores for various noise levels/additional spatial filtering. e) Clusterings for low and medium noise levels. f) Clustering for the highest noise level.

REFERENCES

- Joseph W. Boardman. Analysis, understanding, and visualization of hyperspectral data as convex sets in n-space. In SPIE Proceedings, Imaging Spectrometry, volume 2480, 1995.
- [2] Robert Brandt, Torsten Rohlfing, Jürgen Rybak, Sabine Krofczik, Alexander Maye, Malte Westerhoff, Hans-Christian Hege, and Randolf Menzel. Three-dimensional average-shape atlas of the honeybee brain and its applications. *J Comp Neurol*, 492(1):1–19, Nov 2005.
- [3] Jon Dattoro. Convex Optimisation & Euclidean Distance Geometry. Meboo Publishing, Palo Alto, CA, 2011. ISBN: 0976401304.
- [4] Mathias Ditzen. Odor concentration and identity coding in the antennal lobe of the honeybee Apis mellifera. PhD thesis, Freie Universität Berlin, 2005.
- [5] Patricia C. Fernandez, Fernando F. Locatelli, Nicole Person-Rennell, Gregory Deleo, and Brian H. Smith. Associative Conditioning Tunes Transient Dynamics of Early Olfactory Processing. J. Neurosci., 29(33):10191–10202, 2009.
- [6] Roberto F Galán, Marcel Weidert, Randolf Menzel, Andreas V M Herz, and C. Giovanni Galizia. Sensory memory for odors is encoded in spontaneous correlated activity between olfactory glomeruli. *Neural Comput*, 18(1):10–25, Jan 2006.
- [7] C. Giovanni Galizia, Sabrina L. McIlwrath, and Randolf Menzel. A digital three-dimensional atlas of the honeybee antennal lobe based on optical sections acquired by confocal microscopy. *Cell Tissue Res*, 295(3):383–394, Mar 1999.
- [8] C. Giovanni Galizia and Randolf Menzel. The role of glomeruli in the neural representation of odors: results from optical recording studies. J. Insect Phys., 47:115–129, 2001.
- [9] John Gruninger, Anthony J. Ratkowski, and Michael L. Hoke. The sequential maximum angle convex cone (SMACC) endmember model. In SPIE Proceedings, Algorithms for Multispectral, Hyperspectral and Ultraspectral Imagery, volume 5425-1, 2004.

- [10] Agustin Ifarraguerri and Chein-I Chang. Multispectral and hyperspectral image analysis with convex cones. *IEEE Trans. Geoscience and Remote Sensing*, 37(2):756–770, 1999.
- [11] Ian T. Jolliffe. Principal Component Analysis. Springer, New York/Berlin/Heidelberg, 2nd edition, 2002. ISBN: 0387954422.
- [12] Bettina Malnic, Junzo Hirono, Takaaki Sato, and Linda B. Buck. Combinatorial receptor codes for odors. *Cell*, 96(5):713 – 723, 1999.
- [13] Lisa Rath, C. Giovanni Galizia, and Paul Szyszka. Multiple memory traces after associative learning in the honeybee antennal lobe. *European Journal of Neuroscience*, 34(2):352– 360, 2011.
- [14] Ralph Tyrell Rockafellar. Convex Analysis. Princeton University Press, Princeton, NJ, 1970. ISBN: 0691080690.
- [15] Martin Stetter, Heiner Greve, C. Giovanni Galizia, and Klaus Obermayer. Analysis of calcium imaging signals from the honeybee brain by nonlinear models. *Neuroimage*, 13(1):119– 128, Jan 2001.
- [16] Martin Strauch and C. Giovanni Galizia. Registration to a neuroanatomical reference atlas - identifying glomeruli in optical recordings of the honeybee brain. In *Proceedings of GCB*, volume 136 of *LNI*, pages 85–95. GI, 2008.
- [17] Martin Strauch and C. Giovanni Galizia. Fast PCA for processing calcium-imaging data from the brain of Drosophila melanogaster. In *Proceedings of DTMBIO*, pages 3–10. ACM, 2011.
- [18] Juyang Weng, Yilu Zhang, and Weyshiuan Hwang. Candid covariance-free incremental principal component analysis. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 25:1034–1040, 2003.

ACKNOWLEDGEMENTS

MS is associated with the DFG Research Training Group GK-1042 at the University of Konstanz.



Figure 5. a) Clusterings of the image plane for three bees. Anatomical landmarks highlighted by lines. Left and right ALs (corresponding to the left and right antenna) are mirror-symmetric and contain the same types of glomeruli. Glomeruli are labelled according to the anatomical reference AL. b)-e) left: Low-rank reconstructions $A_c = \hat{T}\hat{S}$ of the original movie A. We show excerpts (timepoints 30 to 37 out of a 110 timepoints odor block) from two movies: Beel and Bee2, each false-color coded (same color scale for all images from one bee). Both bees received the odor nonanol twice with a pause of several minutes inbetween. The black box marks the onset of odor stimulation. Timeseries in T were normalised by subtracting the mean activity before odor onset, i.e. we see relative changes in calcium level. right: Timeseries for three glomeruli (17, 29, 33) with characteristical temporal dynamics in response to the odor nonanol. b) Bee 1, first nonanol measurement c) Bee1, second nonanol measurement, d) Bee2, first nonanol measurement e) Bee2, onset plus Gaussian spatial filtering.