University of Hertfordshire UH 25 Research Archive

Citation for published version:
J. Soelter, J. Schumacher, H. Spors, M. Schmuker, “Automatic segmentation of odor maps in the mouse olfactory bulb using regularized non-negative matrix factorization”, *NeuroImage,* Vol. 98: 279-288, April 2014.

DOI:
https://doi.org/10.1016/j.neuroimage.2014.04.041

Document Version:
This is the Published Version.

Copyright and Reuse:
© 2014 The Authors. Published by Elsevier Inc.
This is an open access article under the CC BY-NC-SA license (http://creativecommons.org/licenses/by-nc-sa/3.0/).

Enquiries
If you believe this document infringes copyright, please contact the Research & Scholarly Communications Team at rsc@herts.ac.uk

NeuroImage 98 (2014) 279–288

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

Automatic segmentation of odor maps in the mouse olfactory bulb using regularized non-negative matrix factorization

Jan Soelter [a], Jan Schumacher [b], Hartwig Spors [b], Michael Schmuker [a,c,*]

[a] Neuroinformatics & Theoretical Neuroscience, Institute of Biology, Freie Universität Berlin, 14195 Berlin, Germany
[b] Max-Planck Institute for Biophysics, Frankfurt am Main, Germany
[c] Bernstein Center for Computational Neuroscience Berlin, 10119 Berlin, Germany

ARTICLE INFO

Article history:
Accepted 12 April 2014
Available online 24 April 2014

Keywords:
Non-negative matrix factorization
Independent component analysis
Intrinsic optical signal
Olfactory bulb

ABSTRACT

Segmentation of functional parts in image series of functional activity is a common problem in neuroscience. Here we apply regularized non-negative matrix factorization (rNMF) to extract glomeruli in intrinsic optical signal (IOS) images of the olfactory bulb. Regularization allows us to incorporate prior knowledge about the spatio-temporal characteristics of glomerular signals. We demonstrate how to identify suitable regularization parameters on a surrogate dataset. With appropriate regularization segmentation by rNMF is more resilient to noise and requires fewer observations than conventional spatial independent component analysis (sICA). We validate our approach in experimental data using anatomical outlines of glomeruli obtained by 2-photon imaging of resting synapto-pHluorin fluorescence. Taken together, we show that rNMF provides a straightforward method for problem tailored source separation that enables reliable automatic segmentation of functional neural images, with particular benefit in situations with low signal-to-noise ratio as in IOS imaging.

© 2014 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC-SA license (http://creativecommons.org/licenses/by-nc-sa/3.0/).

Introduction

Measuring the activity of large neuronal ensembles is a fundamental problem in neuroscience. Functional optical imaging is a widely used tool to measure spatio-temporal responses of neuronal ensembles distributed over extended brain areas. Many neuroscientific questions require identification of regions in the imaged areas which exhibit correlated activity, for example groups of neurons which are functionally related or anatomically grouped. For an efficient work-flow it is desirable that large volumes of recorded ensemble activity are automatically disaggregated into functional/anatomical parts, with each part characterized by its associated pixels (spatial location) and their common time-course (Dorostkar et al., 2010). The automatic segmentation of neuronal activity is not only important to increase experimental throughput, but also to increase the reproducibility and reliability of the results.

For instance, this problem arises in extracting individual glomeruli and their response time-course from optical recordings of the olfactory bulb, a system extensively studied using functional imaging (for a review see Pain et al., 2011). Glomeruli are the first relay station in the olfactory pathway, with each glomerulus relaying the convergent input of a distinct type of olfactory sensory neurons (OSNs) expressing the same olfactory receptor (Firestein, 2001). Thus, the input to the glomerular ensemble represents the basic sensory representation of the olfactory world. Our understanding of this representation is still rudimentary and disputed (Ma et al., 2012; Murthy, 2011; Soucy et al., 2009). This circumstance demands large-scale studies of chemical receptive fields of glomeruli and their spatial arrangement, a task conveniently achieved using optical imaging of neuronal activity.

A well established technique for imaging the dorsal olfactory bulb is to measure reflectance at about 700 nm, the so-called intrinsic optical signal (IOS) (Rubin and Katz, 1999).

The IOS comprises both a global diffuse signal and a local signal originating in the glomeruli (Meister and Bonhoeffer, 2001). The local signal is related to OSN glutamate release and its uptake by astrocytes (Gurden et al., 2006). It has been demonstrated that the IOS corresponds well to both pre-synaptic calcium signals (Wachowiak and Cohen, 2003) and pre-synaptic exocytosis measured using synapto-pHluorin (Soucy et al., 2009). However, the intrinsic optical signal is by a factor of three to ten (calcium) up to twenty (synapto-pHluorin) weaker, implying a smaller signal-to-noise ratio which poses a significant challenge to data analysis. The advantage of IOS imaging is that it doesn't require an ion-sensitive dye, which significantly simplifies the experimental procedure.

* Corresponding author at: Neuroinformatics & Theoretical Neuroscience, Institute of Biology, Freie Universität Berlin, 14195 Berlin, Germany.
E-mail address: m.schmuker@fu-berlin.de (M. Schmuker).

http://dx.doi.org/10.1016/j.neuroimage.2014.04.041
1053-8119/© 2014 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC-SA license (http://creativecommons.org/licenses/by-nc-sa/3.0/).

A range of algorithms and techniques have been applied with the aim of identifying both the location of glomeruli and the time course of their activation from the noisy optical imaging data. In the most straightforward approach the image is de-noised and spatial regions of interest are selected as local maxima of odor induced activation, crossing a defined threshold (Bathellier et al., 2007; Belluscio and Katz, 2001; Ma et al., 2012; Meister and Bonhoeffer, 2001; Wachowiak and Cohen, 2001). The corresponding time courses are then extracted from the surrounding pixels. However, choosing the right threshold is difficult. A high threshold delivers robust segmentation of single spatial peaks in the signal, but small activations may not be detected. In contrast, neighboring strong response regions merge when using a low threshold (Dorostkar et al., 2010). Furthermore biological responses to stimuli are often superimposed onto biological and technical noise sources (Grinvald et al., 1999). The manual choice of a threshold may also complicate the reproduction of results in independent experiments even if experimental conditions vary only slightly.

Techniques for blind source separation (BSS) offer an elegant way out of this dilemma. Using BSS techniques it is possible to separate distinct sources which compose a signal, allowing the simultaneous determination of spatial extent and temporal response of the sources (*i.e.* glomeruli in our case). For example spatial independent component analysis (sICA) was successfully applied to analyze IOS odor responses in the mouse olfactory bulb (Reidl et al., 2007) and to extract glomeruli in calcium imaging experiments in the honeybee antennal lobe (Strauch and Galizia, 2012). The sICA approach relies on the spatial structures of the sources being independent, *i.e.* that they are located at different positions or have different shapes (*e.g.*, blood vessels *vs.* glomeruli). Instead of assuming spatial independence, other approaches in optical imaging segmentation have been proposed which separate sources by assuming temporal (Strauch and Galizia, 2012) or spatio-temporal (Mukamel et al., 2009) independence. But since glomeruli can exhibit largely overlapping ligand spectra and hence correlated odor responses (Ma et al., 2012), the assumption of temporal independence is violated and thus requires at least an additional step of image segmentation. Also the convex cone analysis (Strauch et al., 2012) suffers from the demand of additional post-processing steps.

We propose regularized non-negative matrix factorization (rNMF) (Cichocki and Anh-Huy, 2009) as an alternative approach to identify the spatial location and temporal activity of glomeruli in the olfactory bulb. rNMF allows incorporating *a priori* knowledge about the source characteristics through appropriate regularization terms. In our case, this knowledge comprises both the spatial arrangement of glomeruli, largely side to side with a diameter of 40–190 mm (Royet et al., 1988), and the excitatory response characteristic of OSNs to odor stimulation (Nara et al., 2011). Because the accuracy of BSS depends critically on the suitability of the assumption underlying the factorization algorithm, the incorporation of prior knowledge renders rNMF a well suited approach for glomerulus extraction.

In this study, we introduce regularization terms specific to our BSS problem of identifying glomeruli and extracting their response time-courses in IOS imaging. We elucidate how suitable parameters can be chosen in a data-driven fashion. In addition, we compare the rNMF approach to conventional sICA and examine the application domain of both approaches.

Materials and methods

Mathematical notation

Matrices are denoted in bold capital letters and bold small letters indicate column vectors. Therefore $\mathbf{a}^T\mathbf{x}$ represents the dot product and $\mathbf{a}\mathbf{x}^T$ the outer product.

Experiments

Functional intrinsic optical signal imaging

Three OMP-SpH-mice (9–22 weeks) (Bozza et al., 2004) were anaesthetized using urethane (1.5 g/kg i.p.). Anaesthetic was supplemented throughout the experiments and the body temperature was kept between 36.5 °C and 37.5 °C using a heating pad and a rectal probe. For imaging a craniotomy over one olfactory bulb was cut. The dura mater was removed and the imaging chamber was filled with agar (1.5%) and covered with a glass cover slip. The prepared skull was fixated with cement to a metal plate under the microscope. All animal care and procedures were in accordance with the animal ethics guidelines of the Max Planck Society.

Instant JChem was used for searching, managing and property prediction of odorants in a chemical database (Instant JChem 5.9.4, 2012, ChemAxon, http://www.chemaxon.com).

Odors were presented with a two armed robot (Combipal, CTC-Analytics, Zwingen, Switzerland) using the Software Chronos (Axel Semrau, Sprockhoevel, Germany). 2.5 ml of the odor headspace was injected into a constant carrier flow of filtered and humidified air (2 l/min) towards the mouse's nose. After each odor presentation the syringe used for odor transfer was flushed with nitrogen for 72 s to minimize contamination. Odor responses were recorded in the dorsal olfactory bulb for 12 s at 5 Hz using a macroscope (Pentax Zoom lens 12–48 mm, f = 1:1.0 and Nikkor 135 mm, f = 1:2.0) and an Orca-R2 camera (Hamamatsu, Japan; 1024 × 1344 pixels, field of view 1.63 mm × 1.24 mm) under illumination with red light (690 nm). Odor molecules reached the nose 2.5 ± 0.3 s after recording onset as measured by a photoionization detector (Aurora Scientific, Canada). In each animal the response to a stimulus set of 46 to 47 odors was recorded (for odor list see supplemental table T1). Each odor stimulus was repeated at least twice and stimuli were represented in a pseudo-randomized sequence. Before and after each presentation of the entire stimulus set, the pattern of blood vessels was recorded using green illumination (546 nm, 'green image') and controlled for shifts to exclude movement artifacts.

Anatomical SpH imaging

In addition to functional imaging we performed an anatomical scan in all mice. Synapto-pHluorin labeled OB glomeruli were imaged using a 2-photon laser scanning microscope (Prairie Technologies, Middleton, TN, USA), a 16× water immersion objective (N.A. 0.8, Nikon, back aperture overfilled) and a MaiTai DeepSee laser (50–170 mW, tuned to 880 nm, 80 MHz repetition rate of pulses 120 fs in length; Spectra-Physics/Newport, Santa Clara, CA, USA). Images (512 × 512 pixels) were acquired at 3 μm steps in z-direction. Multiple Z-stacks were stitched and aligned to the functional imaging data using custom written Matlab scripts. For alignment we used the blood vessel pattern obtained by illuminating the olfactory bulb using the 'green image' recorded at a wavelength of 546 nm. We manually outlined glomeruli in the Z-stacks (see Supplementary movie). To avoid a bias in the outlining procedure the glomeruli were identified with the experimenter blinded to the results of factorization.

Data preprocessing

To increase signal-to-noise ratio and reduce computational load the raw data was filtered by binning with an 8 × 8 pixel spatial and a 12 frame temporal window. Then the odor induced activation was calculated as the relative decrease of reflectance $-\Delta R/R = -(R - R_0)/R$. R_0 is the mean reflectance on the first 2 s after recording onset, well before the odors reached the nose (see above). Furthermore the data was spatially bandpass filtered with two Gaussian filters ($\sigma_{low} = 10$ pixels, $\sigma_{high} = 1$ pixels) and down-sampled by a factor of 2. The final resolution of the measurement time series was thus 64 × 84 pixels at 0.42 Hz. The concatenation of the preprocessed frames for all odors

leads to the measurement matrix $\mathbf{Y} \in R^{F \times P}$ with element $Y_{f,p}$ being the observed value of the pth pixel in the fth frame. Every fifth frame is the start of a new odor o_i, in the following denoted as f_{o_i}.

After preprocessing we estimated the remaining pixel noise as the standard deviation of pixel activation in response to the 'non-odor' Argon.

Surrogate data

To evaluate the factorization performance we created surrogate datasets resembling the main characteristics of the biological case. That is we randomly placed 40 sources side by side on a regular 9 × 9 grid in a 50 × 50 pixel image (Fig. 1a). A source s contributes to the activation of a pixel p with a decaying pixel participation $\mathbf{x}_s$ around the source's center $x_{s,p} = \exp(-0.1(p - p_s^{center})^2)$ (Fig. 1d). For a set of $i = 1, \ldots n$ surrogate stimuli o_i, we drew peak activations $a_{o_i,s}$ from a gamma distribution ($\mu = 0.2$, $\sigma = 0.28$) (Fig. 1b), resulting in a narrowly tuned response spectrum $\mathbf{spec}_s = (a_{o_1,s}, a_{o_2,s}, \ldots, a_{o_n,s})$. A temporal correlation between groups of sources, as depicted in Fig. 1b, was introduced *via* a Gaussian copula (Nelsen, 1998). Furthermore each stimulus response was expanded to a six frame time series mimicking the shape of an experimentally observed time course (Fig. 1c). The concatenation of these single stimuli responses yielded the overall time series $\mathbf{a}_s$

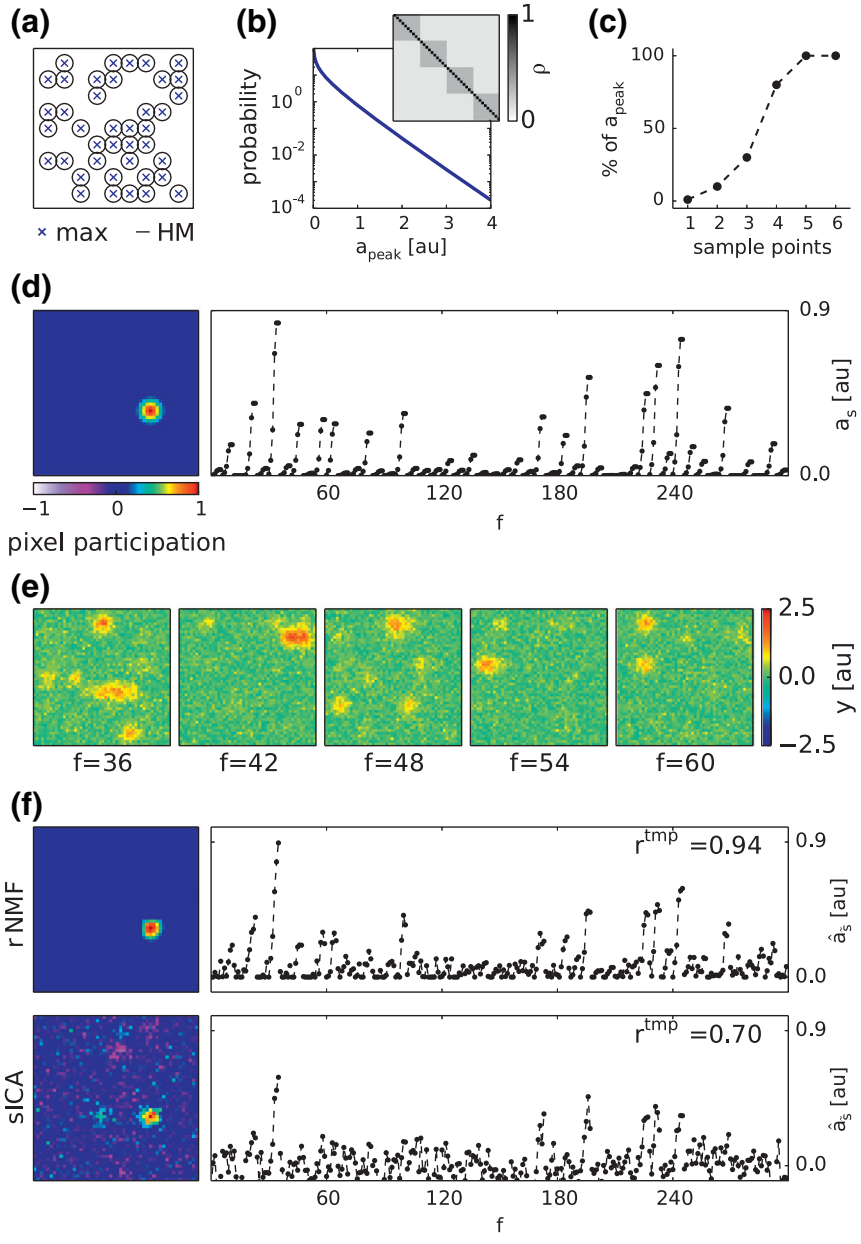

Fig. 1. Surrogate data. (a) 40 Gaussian shaped sources randomly placed on a regular 9 × 9 grid in a 50 × 50 pixel image. Crosses mark centers of pixel participation and circles indicate half maximum. (b) Distribution of per-stimulus peak activations (in arbitrary units au) (gamma distribution with $\mu = 0.2$, $\sigma = 0.28$). Inset: Gaussian copula with correlations ρ used to induce temporal correlation within four groups of sources. (c) Six-point model time-course for stimulus activation. (d) Pixel participation (left) and full activation time series of a source instance (50 odors in 300 frames f). (e) Surrogate signal for five example stimuli at peak activation, including gaussian pixel noise ($\sigma_{noise} = 0.2$). (f) Left panels: recovered source in rNMF (top) and sICA (bottom). Right panels: reconstructed source activation time series. Temporal correlation r^{tmp} to true source is given in upper left corners.

of a source s, modeling the response of a single glomerulus throughout a series of measurements. Finally the summed activation of all sources gave the overall spatio-temporal observation $\mathbf{Y} = \sum_s \mathbf{a}_s \mathbf{x}_s^T + \mathbf{N}$ additionally corrupted by gaussian pixel noise $N_{f,p} \in \mathcal{N}(0, \sigma_{\text{noise}})$ (Fig. 1e).

Matrix factorization

The objective of matrix factorization is to obtain a decomposition of observation $\mathbf{Y}$ into K components with time-courses $\hat{\mathbf{A}} \in R^{F \times K}$ and pixel participations $\hat{\mathbf{X}} \in R^{K \times P}$. Therein $\hat{a}_{f,k}$ is the activation value of component k in the fth frame and $\hat{x}_{k,p}$ the participation of the pth pixel in it. Accordingly $\hat{\mathbf{a}}_k$ is the full time-course of this component and $\hat{\mathbf{x}}_k$ the participation strength for all pixels. The reconstruction of a measurement is then given by $\hat{\mathbf{Y}} = \sum_k \hat{\mathbf{a}}_k \hat{\mathbf{x}}_k^T = \hat{\mathbf{A}}\hat{\mathbf{X}}$ and a measurement is described as its reconstruction and the unexplained residual data $\mathbf{Y} = \hat{\mathbf{A}}\hat{\mathbf{X}} + \mathbf{R}$. After factorization we rescaled $\hat{\mathbf{x}}'_k = \hat{\mathbf{x}}_k / \max(\hat{\mathbf{x}}_k)$ and $\hat{\mathbf{a}}'_k = max(\hat{\mathbf{x}}_k) * \hat{\mathbf{a}}_k$ which leaves the contribution of a component $\hat{\mathbf{a}}'_k \hat{\mathbf{x}}'^T_k = \hat{\mathbf{a}}_k \hat{\mathbf{x}}_k^T$ invariant and allows for reading the pixel-participation as relative strength of participation in the common time series.

Spatial ICA (sICA)

In sICA, factorization is performed under the objective that all pixel participations $\hat{\mathbf{x}}_k$ are mutually independent. We employed the FastICA implementation of scikit-learn (Pedregosa et al., 2011) to obtain such a factorization. First the data is whitened and reduced to K components *via* principal component analysis. The choice of K determines the number of unique components extracted by ICA. The resulting factorization $\mathbf{Y} = \mathbf{A}^{\text{PCA}}\mathbf{X}^{\text{PCA}} + \mathbf{R}$ has minimal variance in residual $\mathbf{R}$ and orthogonal pixel participations $\mathbf{x}_k^{\text{PCA}}$ of zero mean and unit variances. To obtain pixel participations $\hat{\mathbf{x}}_k$ which are not only uncorrelated but independent, FastICA estimates an unmixing matrix $\mathbf{W}$ maximizing the nongaussianity/negentropy of the 'unmixed' components $\hat{\mathbf{X}} = \mathbf{W}\mathbf{X}^{\text{PCA}}$ (Hyvärinen, 1999). This yields the final factorization $\mathbf{Y} = \hat{\mathbf{A}}\hat{\mathbf{X}} + \mathbf{R} = \mathbf{A}^{\text{PCA}} \mathbf{W}^{-1}\mathbf{W}\mathbf{X}^{\text{PCA}} + \mathbf{R}$.

Regularized NMF (rNMF)

In non-negative matrix factorization the values of factorization are restricted to be positive ($a_{f,k} \geq 0, x_{k,p} \geq 0 \; \forall \, f, k, p$). Positive pixel participations enable straightforward physiological interpretation, reading the pixel values in each component as contribution values of the extracted physiological source. Positive time-courses reflect the assumption that source responses are excitatory. To obtain such a factorization the HALS algorithm iteratively minimizes the reconstruction error $\mathbf{R} = \left|\left|\mathbf{Y} - \hat{\mathbf{A}}\hat{\mathbf{X}}\right|\right|_F^2$ under the non-negativity constraint (Cichocki and Anh-Huy, 2009), with $||||_F$ denoting the Frobenius Norm. To further constrain the factorization to known characteristics of the hidden sources the algorithm allows the imposition of additional regularization to the estimation. This is achieved by jointly minimizing the reconstruction error $\mathbf{R}$ together with a constraint C in the form $\left|\left|\mathbf{Y} - \hat{\mathbf{A}}\hat{\mathbf{X}}\right|\right|_F^2 + \alpha C\left(\hat{\mathbf{A}}, \hat{\mathbf{X}}\right)$, with α determining their trade off.

In general there is no unique solution of the NMF problem (Donoho and Stodden, 2004). Especially if two sources have a similar activation profile, a valid factorization would be a component reflecting the common ground activation and components containing the deviation.

To avoid the explanation of pixel group activations in more than one component, a good solution $\hat{\mathbf{X}}$ contains only minimal off-diagonal elements in the components' spatial overlap matrix $\hat{\mathbf{X}}\hat{\mathbf{X}}^T$, leading to the global sparseness regularization term $C_{\text{sp}}\left(\hat{\mathbf{X}}\right) = \sum_j \sum_{k \neq j} \hat{\mathbf{x}}_j \hat{\mathbf{x}}_k$ (Chen and Cichocki, 2005). This term reflects the assumption that glomerular signals recorded by IOS imaging show only weak overlap. As a further physiological constraint we took into account that neighboring pixels are likely to belong to the same source. This was reflected in a smoothness regularization, minimizing the spatial variation of each component $C_{\text{sm}} = \sum_k \sum_p \left(\hat{x}_{k,p} - \mathbf{l}_p^T \hat{\mathbf{x}}_k\right)^2$ with $\mathbf{l}_p$ being the neighborhood vector of the pth pixel reflecting its 2D connectivity (Cichocki and Anh-Huy, 2009). Taking both regularizations together, a solution is preferred where the activation of a pixel is either attributed to a single source or the mixed signal of neighboring sources.

We used the hierarchical alternating least squares (HALS) framework (Cichocki and Anh-Huy, 2009) to optimize our regularized objective function. The full algorithm is provided in Appendix A.

The HALS algorithm can be initialized with any guess of $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$. We chose a deterministic approach to obtain reproducible results and avoid any chance effects of random initialization (see Appendix B). In detail, we started with the pixel with the maximal peak activation in the signal and selected its time-course to initialize the first component. Then the participation of all pixels in this time-course was calculated and their contributions to the signal was subtracted from the data to obtain the unexplained residuals $\mathbf{R} = \mathbf{Y} - \hat{\mathbf{a}}_k \hat{\mathbf{x}}_k^T$. We repeated this procedure on the residuals $\mathbf{R}$ until we initialized all k components. This approach is similar to convex cone analysis (Strauch et al., 2012). Strauch et al. selected the pixel with the highest euclidean norm, essentially choosing the pixel with the highest variance. We instead chose the maximum norm (highest peak activation) to address the expected sources' activation characteristics of few but strong (*i.e. sparse*) activations.

Sequential compositions of rNMF and sICA

Both sICA and rNMF might provide complementary aspects of a solution. Therefore we created various sequential compositions of these approaches. The first composition we assessed was to perform one approach on the reconstructed data $\hat{\mathbf{Y}} = \hat{\mathbf{A}}\hat{\mathbf{X}} = \mathbf{Y} - \mathbf{R}$ of the other one. It implies the assumption that the residual $\mathbf{R}$ of the first factorization mainly contains noise, and removing it is equivalent to de-noising the data.

The second composition we used was to initialize rNMF by the rectified factorization $[\mathbf{X}^{\text{sICA}}]_+$, $[\mathbf{A}^{\text{sICA}}]_+$ of sICA. *Vice versa*, sICA was performed directly on the pixel participation $\mathbf{X}^{\text{rNMF}}$ of rNMF (instead of performing it on the data $\mathbf{Y}$). This procedure yielded a total of four sequential compositions of rNMF and sICA: 1) rNMF initialized with the rectified sICA components, 2) rNMF performed on the sICA reconstruction ('de-noising by sICA'), 3) sICA performed on pixel participations obtained with rNMF, and 4) sICA performed on rNMF reconstructions ('de-noising by rNMF').

Performance criteria

We employed different measures to evaluate the performance of factorization. In rNMF we estimated the efficiency of the sparseness constraint by calculating the spatial correlation (pixel-wise Pearson's r) of components. Especially the highest correlation $r_k^{\text{comp}} = \max_j r(\hat{\mathbf{x}}_k, \hat{\mathbf{x}}_j)$ highlights shared explanation of pixel groups in the factorization, which is in opposition to the assumption of locally sparse signals.

In case of the surrogate dataset where the ground truth is known we furthermore directly assessed how well the true sources were reconstructed in factorization. Therefore each source was assigned to the component to which it exhibited the highest spatial correlation $\tilde{s} = \arg\max_i r(\hat{\mathbf{x}}_i, \mathbf{x}_s)$. The recovery error of a source was then calculated as the relative mean squared error of reconstruction for all time-points and pixels $\text{MSE}_s = \sum_{f,p} (a_{f,s} x_{s,p} - \hat{a}_{f,\tilde{s}} \hat{x}_{\tilde{s},p})^2 / \sum_{f,p} (a_{f,s} x_{s,p})^2$. As $\hat{\mathbf{x}}_{\tilde{\mathbf{s}}}$ in sICA exhibits small but many non-local values (see Fig. 1f), this measure would indicate worse performance than actually could be

achieved by simple post-processing like thresholding. To ensure a fair performance comparison between sICA and rNMF, we therefore accounted for potential post-processing improvements by calculating the sICA recovery error only for the local source region $\mathrm{MSE}_s^{\mathrm{sICA}_{\mathrm{loc}}} = \sum_{f,p^{\mathrm{loc}}}\left(a_{f,s}x_{s,p^{\mathrm{loc}}} - \hat{a}_{f,\hat{s}}\hat{x}_{\hat{s},p^{\mathrm{loc}}}\right)^2 / \sum_{f,p^{\mathrm{loc}}}\left(a_{f,s}x_{s,p^{\mathrm{loc}}}\right)^2$ for pixels $p^{\mathrm{loc}} = \{p | x_{s,p} > 0.05\}$ which have a substantial participation in the original source. Instead of recovery error, we will report the counterpart source recovery $\mathrm{SR} = 1 - \mathrm{MSE}$. Furthermore we characterized coincidence of source and matched component by evaluating their mutual spatial and temporal correlation, $r_s^{\mathrm{spt}} = r(\hat{\mathbf{x}}_{\hat{s}}, \mathbf{x}_s)$ and $r_s^{\mathrm{tmp}} = r(\hat{\mathbf{a}}_{\hat{s}}, \mathbf{a}_s)$.

For the experimentally obtained IOS data the ground truth is unknown. Nevertheless we were able to utilize the exactly defined experimental design of repeated stimuli and the generally high reproducibility of odor responses on single trial basis (Spors and Grinvald, 2002) for evaluating temporal characteristics of a component. For each component k we calculated its response spectrum $\mathbf{spec}_k = (a_{o_1,k}, a_{o_2,k}, \ldots, a_{o_n,k})$ to the first and second stimulus repetition. The response to an odor o_i was the mean activation 4.8 s–12 s after odor measurement onset: $a_{o_i,k} = \sum_{f=f_{o_i}+2}^{f_{o_i}+4} \hat{a}_{f,k}/3$. The trial-to-trial correlation of response spectra $r_k^{\mathrm{t2t}} = r(\mathbf{spec}_k^{\mathrm{1st}}, \mathbf{spec}_k^{\mathrm{2nd}})$ reflects then the stimulus correlated fraction of activation in the component. Low r_k^{t2t} indicates sources with weak coupling to odor stimulation, for example blood vessels.

Source code

Source code (Python) for performing rNMF, including a GUI version, is available online at https://github.com/jansoe/FUImaging/tree/Neuroimage2014. In the "examples" subfolder of this repository, we provide IPython Notebooks to reproduce the data analysis performed in this manuscript. Supplemental file 3 contains the necessary data. Online inspection of this code is possible at http://nbviewer.ipython.org/github/jansoe/FUImaging/blob/master/examples/IOSsegmentation/regNMF.ipynb.

Results

The purpose of this study was to incorporate appropriate regularization in the NMF framework to improve automatic segmentation of glomeruli in the mouse olfactory bulb in IOS imaging. We first use surrogate data to analyze the benefits of rNMF, including the performance-critical choice of appropriate regularization. In a practical IOS imaging scenario we illustrate the capability of rNMF to infer glomerular positions and responses, and verify the results using anatomical measurements. At all stages, we compare the performance of the rNMF approach to conventional sICA.

Segmentation of surrogate data

It is generally difficult to assess factorization performance for experimentally obtained imaging data because the ground truth is unknown. To obtain nevertheless a detailed picture on the terms of performance for rNMF and sICA we constructed a parameterized surrogate dataset in which we have full control over the sources composing the signal. With this dataset we could address two important questions: First, what is the influence of method inherent parameters and how can we choose their values for a given dataset? And second, what is the application domain of both methods with respect to strength of pixel noise and number of measured stimuli? To answer these questions we constructed surrogate sources resembling the main characteristics of our biological object of research. Surrogate glomeruli are arranged side by side with overlapping spatial signal distribution (Fig. 1a). This induces a spatial correlation of 0.29, *i.e.* a small dependence, to neighboring sources. Their response spectra are narrowly tuned and groups of glomeruli exhibit correlation in their response spectra (Fig. 1b). Each glomerulus rises to peak activation with a model time-course that mimics measured response dynamics (Fig. 1c). The data to enter factorization is the concurrent observation of 40 glomeruli in response to n_{stim} stimuli (*e.g.* odors) corrupted by additional pixel noise σ_{noise} (Fig. 1e).

We started our analysis with a dataset roughly mimicking the properties of our intrinsic optical signal (IOS) imaging data with $n_{\mathrm{stim}} = 50$ stimulus observations and a noise level of $\sigma_{\mathrm{noise}} = 0.2$. Fig. 1f shows an example of a recovered source from both rNMF ($\alpha_{\mathrm{sm}} = 2$, $\alpha_{\mathrm{sp}} = 0.5$) and sICA, illustrating the general characteristics of the methods. rNMF indeed showed the desired properties of a localized, sparse and smooth pixel participation, accurately reproducing the spatial and temporal characteristics of the source. In contrast plain sICA (with no additional processing applied) generates more holistic pixel participations, containing global noise contributions besides the local source contribution. While the non-local aspects could probably be mitigated using suitable post-processing, this result points out a more noisy reconstruction of the activation courses by sICA, especially for weaker signals.

Choice of regularization parameters

The outcome of rNMF factorization depends on the choice of regularization. On the basis of problem-specific knowledge we devised two regularization terms on the basis of expected source characteristics: spatial sparseness and spatial smoothness (see Materials and methods). The relative influence of those regularization terms is governed by the parameters α_{sp} (sparseness) and α_{sm} (smoothness). We systematically evaluated their effect on factorization results on the surrogate data in order to provide a heuristic for choosing useful parameter values.

We first started with the parameter for smoothness regularization α_{sm}. With $\alpha_{\mathrm{sm}} = \alpha_{\mathrm{sp}} = 0$, that is without any regularization, pixel participations of the NMF components spread across the whole image, containing small scale structures of one pixel size. These small scale structures progressively disappeared as we increased the smoothness regularization to $\alpha_{\mathrm{sm}} = 8$ (Fig. 2a). Nonetheless, the extracted components contained contributions from different sources even with strong smoothness regularization. Therefore we introduced sparseness regularization, which controls the number of components a pixel is participating through the parameter α_{sp}. With moderate regularization ($\alpha_{\mathrm{sp}} = 0.5$), a component exactly described a single source (Fig. 2b, left panel). Setting its value too high (*e.g.* $\alpha_{\mathrm{sp}} = 4$) resulted in components covering only part of a source, as any overlap of components is prevented by the regularization (Fig. 2b, right panel).

To further quantify the effect of sparseness regularization we employed two measures. First, we calculated the recovery of the actual surrogate sources by the components, SR_k. Second, we measured shared explanation of pixel groups in multiple components *via* spatial correlation between components r_k^{comp} (see Materials and methods). This is an uninformed measure not depending on any knowledge of actual sources, and hence is also applicable when the ground truth is not known, *i.e.* in real experimental data.

We found that the mean source recovery SR_k, starting at almost zero without any sparseness regulation, increased to a maximum at about $\alpha_{\mathrm{sp}} = 0.25$ (2^{-2}) and then decreased again (Fig. 2c). In contrast, the maximal mutual correlation between components r_k^{comp} started to decrease to zero at sparseness regularization values of $\alpha_{\mathrm{sp}} > 0.125$ (2^{-3}) (Fig. 2c). The range of maximal source recovery was thus in a regime where all strong component correlations became eliminated but some residual correlation was left. This reflects the fact that neighboring sources actually have a spatial correlation of $r = 0.26$ due to their spatial overlap. The interdependency of both measures suggests the simple heuristic to choose α_{sp} in a regime where spatial component correlation starts to be significantly reduced. In the following we implemented this by choosing $\alpha_{\mathrm{sp}}^{\mathrm{opt}}$ as the first value of α_{sp} where $\max_k r_k^{\mathrm{comp}}$ drops below 0.5.

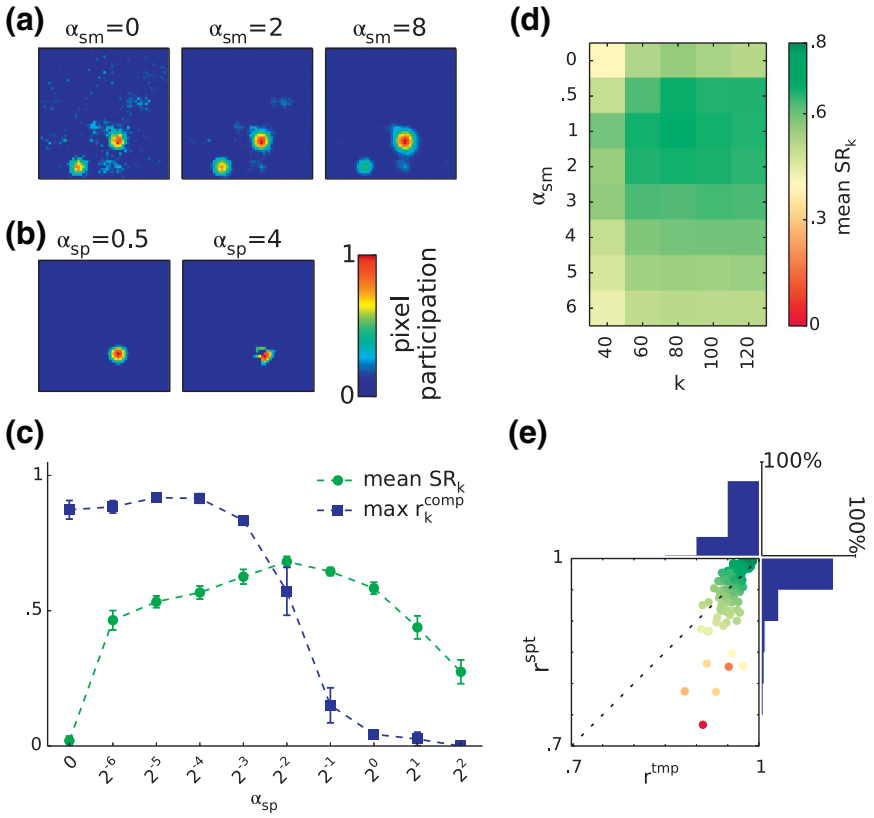

Fig. 2. Regularization parameter choice for surrogate data set ($n_{stim} = 50$, $\sigma_{noise} = 0.2$). (a) Effect of smoothness regularization parameter α_{sm} on an rNMF-extracted component ($k = 80$, $\alpha_{sp} = 0$). (b) Effect of sparseness regularization through α_{sp} on the rNMF component ($k = 80$, $\alpha_{sm} = 2$). (c) Maximal spatial component correlation r^{comp} and standard deviation for five different dataset instances (blue) respectively mean recovery SR_k (green) in dependence of α_{sp} ($k = 80$, $\alpha_{sm} = 2$). (d) Mean source recovery SR in dependence of number of components k and smoothness regularization α_{sm}. (e) Spatial and temporal correlation of matched components ($k = 80$, $\alpha_{sm} = 2$, $\alpha_{sp} = 0.5$) to 200 sources of five dataset instances. The color of each dot indicates the corresponding SR value (colorbar see (d)). Histograms on axes depict per-axis marginal distributions.

We next asked how sensitive the results are on the initial choice of smoothness regularization α_{sm} and number of components k. Setting $\alpha_{sp} = \alpha_{sp}^{opt}$ according to the above derived heuristic, the factorization performance was robust to the exact choice of k as long as it substantially exceeded the number of sources, *e.g.* by a factor of 1.5 to 3 (Fig. 2d). With respect to smoothness regularization the best results were obtained for values of α_{sm} between 0.5 and 2. This is a regime where most, but not all small scale structures are removed.

Taken together, we suggest the following heuristic to choose rNMF parameters: set the number of components k well above the number of expected sources (as a rule of thumb a factor of two), tune smoothness regularization to just remove most small scale structures and adjust sparseness regularization to just remove any strong component correlation.

With parameters corresponding to these heuristics ($k = 80$, $\alpha_{sm} = 2$, $\alpha_{sp} = \alpha_{sp}^{opt} = 0.5$) we obtained near optimal factorizations of the surrogate data with very good source recovery. This especially implied an extremely high temporal accuracy with 99.5% of components exhibiting a temporal correlation to their corresponding source higher than 0.9. Fig. 2e shows that for components with moderate source recovery values the temporal correlation to the source was always larger than 0.85.

Comparison of rNMF and sICA

Having demonstrated the benefits of appropriate regularization in rNMF, we next compare the performance of rNMF to that of conventional sICA on the surrogate data set. In sICA both spatial and temporal reconstruction of sources were more noisy, as already pointed out in the example above (*cf.* Fig. 1f). Due to the many but small non-local pixel participation values, the performance measure we used for rNMF might underestimate sICA's recovery performance compared to what could be achieved by simple post-processing like thresholding. To account for such post-processing and enable a fair comparison we only computed source recovery for localized sICA components ($sICA^{loc}$, see Materials and methods). Nonetheless $sICA^{loc}$ components showed significantly lower source recovery than rNMF components (Fig. 3a) ($p = 4 * 10^{-41}$, Kolmogorov–Smirnov test). This result was independent of the number of components k chosen to initialize the method (Fig. 3b). In particular, rNMF outperformed sICA regarding the reconstruction of the true time course of the sources, as measured by the temporal correlation r^{tmp} (Fig. 3c).

In general matrix factorization consists of the objective function to be optimized and the optimization procedure thereof. In FastICA the optimization procedure is further sub-divided in a first step of data reduction to k components *via* PCA and subsequent optimization of the component's independence. To better understand the particular influence of these steps we performed various sequential combinations of rNMF and sICA (see Materials and methods). Initializing rNMF with sICA components had no effect on recovery performance (Fig. 3d, sICA ini). In contrast, recovery was impaired when rNMF was applied on the data reconstruction from the sICA factorization (Fig. 3d, sICA dat). Since the sICA dat condition implies removing the part of the data that sICA reconstruction did not

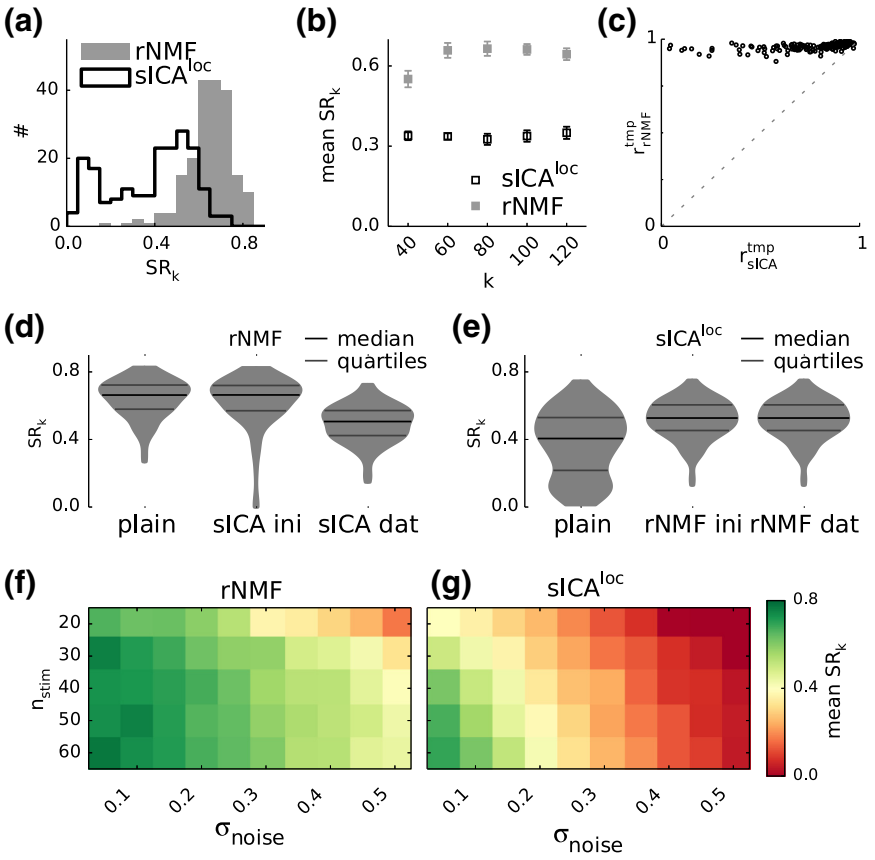

Fig. 3. Comparison of rNMF and sICA performance in source reconstruction. Data is pooled for five independently drawn surrogate datasets. (a) Histogram of source recovery values SR_k for rNMF (solid gray, $k = 80$, $\alpha_{sm} = 2$, $\alpha_{sp} = 0.5$) and localized sICA components $sICA^{loc}$ (black line, $k = 80$). (b) Dependence of mean source recovery SR_k on the number of components k in the surrogate data. (c) Temporal correlation of rNMF components $r^{tmp}{}_{rNMF}$ *vs.* $r^{tmp}{}_{sICA}$ of sICA components. (d,e) Violin plots of distribution on source recovery values SR_k for (d) plain rNMF, rNMF initialized with sICA components (sICA ini) and rNMF performed on sICA data reconstruction (sICA dat) and (e) for plain sICA, sICA performed on pixel participation of rNMF components (rNMF init) and sICA performed on rNMF data reconstruction (rNMF dat). (f,g) Mean source recovery SR_k in (f) rNMF and (g) $sICA^{loc}$ for different amounts of pixel noise σ_{noise} and number of stimuli used for dataset generation n_{stim}.

explain (the residual), this effect shows that in sICA part of the sparse signal is lost in the noise due to the variance driven dimensionality reduction of PCA. We also evaluated the reverse procedures, *i.e.* initializing sICA with rNMF pixel participations (Fig. 3e, rNMF ini), and performing sICA on rNMF reconstructions (Fig. 3e, rNMF dat). We found that both procedures improved the performance of plain sICA, showing that de-noising the data with rNMF is beneficial for sICA segmentation. On the other hand, the resulting sICA segmentation is worse than the results of the prior rNMF step, *i.e.* before applying sICA. This indicates that the objective function of rNMF better suits the source characteristics than the spatial independence objective of sICA.

Finally we asked how dataset statistics influenced the performance of rNMF and sICA.

In particular we asked to which extent the methods were affected by noise and the number of stimuli used for dataset generation (n_{stim}). Both rNMF and sICA performed well in low noise regimes and when the number of stimuli was larger than the number of sources (Figs. 3f, g). However, rNMF was more resilient to pixel noise and also coped better with the 'overcomplete' case when the number of stimuli n_{stim} was smaller than the number of sources (which was set to 40 in all examples). This result points out that the rNMF method is particularly suited to studying odor maps in the olfactory bulb with wide-field IOS imaging, when the expected number of sources (glomeruli) is in the same range or greater than the number of stimuli (unique odorants) used in one recording session.

In-vivo imaging data

The surrogate data provided insight into the application domain of the methods and yielded a heuristic for parameter choice. But experimentally obtained *in-vivo* IOS data is more complex than the surrogate data. With respect to spatial characteristics, glomeruli have varying sizes and shapes and also the exact spatial signal distribution is unknown. Furthermore the temporal characteristics of glomeruli do not obey an unique probability distribution but vary strongly in overall peak activation strength and odor selectivity (Nara et al., 2011).

Since IOS imaging alone doesn't reveal the ground truth regarding glomerulus locations, we performed 2-photon synaptopHluorin imaging of the resting fluorescence and obtained outlines of anatomical glomerulus positions (Fig. 4a). While the anatomical outlines naturally can't provide full information on the actual contribution of individual glomeruli to each pixel in the IOS recording, they do provide reliable information about glomerular position. The process of manual outlining and alignment of functional and anatomical measurements (Fig. 4b) may introduce a small error, complicating a 'hard' assessment of reconstruction quality. Nevertheless these outlines allowed for visual inspection of the match between anatomical glomerulus positions and extracted components.

We thus asked whether our general observations on the surrogate dataset still held for the IOS data. To answer this question, we

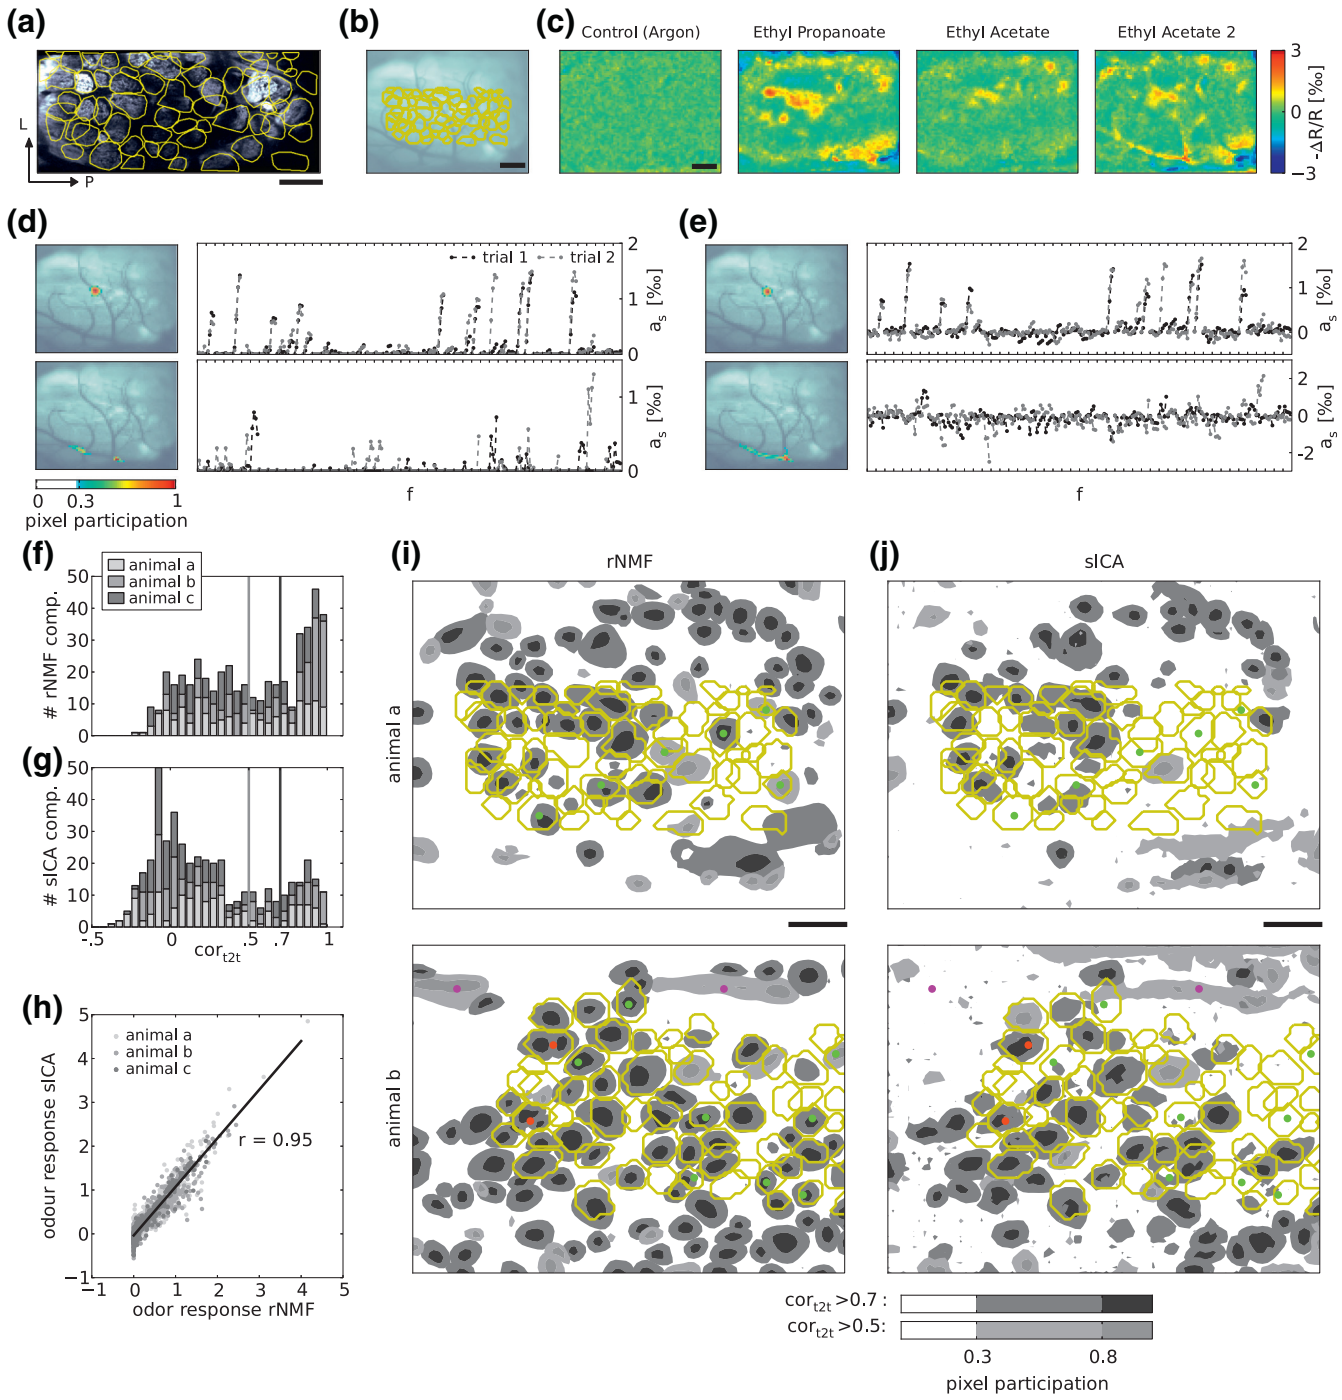

Fig. 4. Factorization of IOS mouse olfactory bulb imaging. (a) Z-frame of the synapto-pHluorin resting fluorescence image stack and manually outlined glomeruli (yellow contours, scale bar: 200 μm). (b) Glomerulus outlines (yellow) aligned to green image of the dorsal OB. (c) Examples of odor response maps obtained with IOS; mean $-\Delta R/R$ 6 s after odor delivery. (d) Two exemplary extracted rNMF components. Top: glomerulus-like component with localized activity (left panel) and high trial-to-trial correlation. Bottom: blood vessel-like component with low trial-to-trial correlation (right panel). (e) Similar components extracted by sICA. (f,g) Stacked histograms of observed trial-to-trial correlations of (f) rNMF and (g) sICA components. Vertical lines indicate thresholds for contour plots in (i,j). (h) Odor responses of reliably extracted glomeruli ($r^{t2t} > 0.7$) obtained by sICA *vs.* responses of matched glomeruli obtained by rNMF. (i) Contour plot of extracted pixel participation compared to anatomical glomerulus outlines (yellow) for rNMF in two animals. Black/dark gray contours indicate components with trial-to-trial correlation ($r_k^{t2t} > 0.7$), light gray contours correspond to components with $r_k^{t2t} > 0.5$. Green marks highlight glomeruli that only appeared in rNMF. Red marks highlight glomeruli better recovered in sICA. Magenta marks highlight components with $0.5 < r^{t2t} < 0.7$ that do not resemble glomeruli. (j) Same as (i) for sICA.

performed sICA and rNMF for three measurements (termed animal a, b, and c) according to our deduced heuristic. Both the pixel noise ($\sigma_{\text{noise}} = 0.15$ (animal a), $\sigma_{\text{noise}} = 0.15$ (animal b), $\sigma_{\text{noise}} = 0.19$ (animal c), estimated from the response to the 'non-odor' Argon) as well as the range of observed activations were in the range of the surrogate dataset (Fig. 4c).

We chose $k = 150$ to initialize rNMF and sICA, well above the estimated number of around 60 activated glomeruli.

In addition we set $\alpha_{sm} = 2$ in rNMF to avoid one-pixel scale structures and $\alpha_{sp} = 0.5$ to prohibit any component correlation above 0.5. In both rNMF and sICA we found some components that resembled glomeruli regarding their spatial extent and the activation course, while others looked like parts of blood vessels or measurement artifacts ('*noise*') (Figs. 4d,e). To distinguish putative glomerulus components from non-glomeruli we exploited the trial-to-trial correlation r_k^{t2t} of the components' odor spectra. The distributions of r_k^{t2t} values were roughly bimodal for both methods with a population of highly stimulus dependent components (Figs. 4f,g). Those stimulus dependent components exhibited good overlap with the anatomical outlines (Figs. 4i,j and Supplemental Fig. S1), taking into account the ambiguity of the outlines along the z-axis.

We made three observations: First, components with $r^{t2t} > 0.7$ corresponded well to anatomical glomerulus outlines. The same holds true for components with $0.5 < r^{t2t} < 0.7$ except in a few cases where components rather resembled blood vessels or other artifacts (magenta marks in Figs. 4i,j). Second, many glomeruli were detected by both methods, but some were solely extracted by rNMF (green marks in Figs. 4i,j) or exhibited higher trial-to-trial correlation r^{t2t} in rNMF than in sICA. Third, in rare cases we observed that two rNMF components were located within one anatomical outline (red marks in Figs. 4i,j) indicating different temporal activation of these parts, potentially due to underlying blood vessels or as an artifact of bandpass filtering.

In general these observations matched the results obtained using the surrogate dataset. rNMF yielded higher source recovery than sICA with a more precise temporal reconstruction, *i.e.* a higher trial-to-trial correlation. These results were robust over a range of parameter values (see Supplemental Fig. S2).

Finally, we asked whether the confinement of rNMF to positive activations had an effect on the extracted components compared to those extracted by sICA. To asses this we matched the most reliable components extracted with sICA ($r^{t2t} > 0.7$) to the respective rNMF component with which it exhibited the highest spatial correlation. The odor response spectra of the matched components were highly similar for both methods (Fig. 4h, $r = 0.95$). Furthermore, sICA extracted only a few and small negative signals which are not present in rNMF. Although this could indicate that our assumption of observing purely excitatory responses with IOS imaging has been mildly impaired, negative values in sICA might just as well be artifacts of the initial highpass filtering.

Discussion

In this study we demonstrated how to introduce prior knowledge about source characteristics through regularization in the NMF framework. We showed that regularized NMF outperformed standard sICA as a blind source separation approach for the automatic segmentation of glomeruli from IOS images of the olfactory bulb. We achieved this by reducing the approach's 'blindness' through incorporating knowledge about the spatial continuity and spatial separation of glomeruli. As a result of the regularization, the rNMF approach was more resilient to pixel noise and required fewer independent observations than sICA to reliably extract the locations and odor response spectra of individual glomeruli. Since the number of measurable stimuli is often restricted by experimental constraints, the reduced demand in stimulus number makes rNMF an interesting choice for other techniques with high signal-to-noise ratio, *e.g.* calcium sensitive dye imaging (Spors et al., 2006).

The combination of smoothness and sparseness regularization promoted factorizations in which observations were explained by neighboring pixels. In our case the almost two dimensional arrangement of glomeruli on the surface of the olfactory bulb justifies this approach. Nevertheless this assumption holds also in many 3D imaging scenarios like functional 2-photon imaging or fMRI. Moreover, tuning the introduced regularizations to the expected distribution of activity in the imaged tissue allows for a fine-grained adjustment of the methods' sensitivity in a broad range of applications.

We have demonstrated that the rNMF results remained stable over a wide range of parameters. But we also showed that regularization is absolutely necessary in NMF to obtain good results. Given a sufficiently large set of stimuli and/or low noise levels, sICA will also yield robust and fast segmentation, with no regularization parameters to tune. Hence, there is a tradeoff between the benefits of rNMF and the additional effort required for appropriate regularization.

In general sICA and NMF are closely related (in particular sparseness regulated NMF Hoyer, 2002, 2004). Therefore instead of alternating the basic NMF approach, one might also modify standard ICA to obtain a problem specific objective function. Indeed there exist numerous ICA variations addressing specific aspects of our regularization, *e.g.* non-negative ICA (Plumbley, 2003), sparse ICA (Babaie-Zadeh et al., 2006) spatio-temporal tomographic non-negative ICA (Valdés-Sosa et al., 2009) or overcomplete ICA (Lee et al., 1999). However, the sICA assumption of spatial independence is violated in our data set because the sources show partial overlap. The advantage of rNMF is that it allows the gradual adjustment of the spatial correlation of the segmented components to the expected spatial correlation of the true sources.

In general rNMF stands out by its straightforward implementation of constraints from prior knowledge. It allows for the addition or replacement of further modifications like temporal decorrelation or temporal smoothness and hence provides an opportunity to tailor factorization to known source statistics. While the assumption of non-negative responses in our data is in agreement with the literature (Nara et al., 2011) and our own results (see the *In-vivo* imaging data section), it is also straightforward to relax this assumption in the HALS algorithm by dismissing the rectification step (Cichocki and Anh-Huy, 2009). This would allow for negative source signals to be covered while still keeping the interpretability of non-negative pixel participations. Such an approach may prove useful to apply rNMF in cases where inhibitory signals are expected, such as using calcium imaging to measure odor maps in the insect antennal lobe (Sachse and Galizia, 2002).

The great flexibility in generating problem tailored factorization makes rNMF a promising approach for automated analysis in many functional imaging situations. In our case it facilitates an automatic and reliable high throughput investigation of chemical receptive fields of glomeruli in the mouse olfactory bulb.

Supplementary data to this article can be found online at http://dx.doi.org/10.1016/j.neuroimage.2014.04.041.

Acknowledgments

This work was supported by grants of Deutsche Forschungsgemeinschaft (MS: DFG SCHM2474/1-1, SCHM2474/1-2 (SPP 1392); HS: FOR 643, SP1134/1-1, SP1134/2-1 (SPP1392)). We thank Alexander Lehmann for sharing Matlab functions and Stephan Gabler for developing a graphical interface to the analysis method. We also thank Anna D'Errico, Martin Vogel, Thomas Rost and Joachim Haenicke for helpful discussions on the manuscript. We thank Jürgen Reichert and his team for excellent technical support, Christiane Kalmbach for electronic support, and Bolek Zapiec for IT support. We thank Peter Mombaerts for generous financial support, for providing infrastructure and resources for mouse breeding and genotyping, and for scientific discussions. We thank Prof. M. Nawrot for his general support and insightful comments.

Appendix A. Regularized NMF algorithm

Symbols are defined in the Materials and methods section. Additionally matrix $\mathbf{L}$ contains the neighborhood vector $\mathbf{l}_p$ of pixel p in it's pth row and $\hat{\mathbf{x}}_k^{\text{prev}}$ denotes the estimation of $\hat{\mathbf{x}}_k$ in the previous iteration.

Algorithm 1 regularized HALS

Initialize $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$

$\mathbf{R} = \mathbf{Y} - \hat{\mathbf{A}}\hat{\mathbf{X}}$

repeat

  for $k = 1$ to K **do**

    $\mathbf{R} \leftarrow \mathbf{R} + \hat{\mathbf{a}}_k \hat{\mathbf{x}}_k^T$

    $\hat{\mathbf{x}}_k \leftarrow \frac{1}{1+\alpha_{\text{sm}}}[\mathbf{R}^T\hat{\mathbf{a}}_k - \alpha_{\text{sp}}\sum_{j\neq k}\hat{\mathbf{x}}_j + \alpha_{\text{sm}}\mathbf{L}\hat{\mathbf{x}}_k^{\text{prev}}]_+$

    $\hat{\mathbf{a}}_k \leftarrow [\mathbf{R}^T\hat{\mathbf{x}}_k]_+$

    $\hat{\mathbf{a}}_k \leftarrow \frac{\hat{\mathbf{a}}_k}{||\hat{\mathbf{a}}_k||_2^2}$

    $\mathbf{R} \leftarrow \mathbf{R} - \hat{\mathbf{a}}_k \hat{\mathbf{x}}_k^T$

  end for

until convergence criterion is reached.

Appendix B. Initialization

Initialization procedure for $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$. $\mathbf{r}_p$ denotes the pth row of matrix $\mathbf{R}$ and therefore the residual time-course of pixel p.

Algorithm 2 Init

$\mathbf{R} = \mathbf{Y}$

for $k = 1$ to K **do**

  $\tilde{p} \leftarrow \text{argmax}_p \|\mathbf{r}_p\|_\infty$

  $\hat{\mathbf{a}}_k \leftarrow \mathbf{r}_{\tilde{p}}$

  $\hat{\mathbf{a}}_k \leftarrow \frac{\hat{\mathbf{a}}_k}{||\hat{\mathbf{a}}_k||_2^2}$

  $\hat{\mathbf{x}}_k \leftarrow [\mathbf{R}^T\hat{\mathbf{a}}_k]_+$

  $\mathbf{R} \leftarrow \mathbf{R} - \hat{\mathbf{a}}_k \hat{\mathbf{x}}_k^T$

end for

References

Babaie-Zadeh, M., Jutten, C., Mansour, A., 2006. Sparse ICA via cluster-wise PCA. Neurocomputing 69 (13–15), 1458–1466.

Bathellier, B., Van De Ville, D., Blu, T., Unser, M., Carleton, A., 2007. Wavelet-based multiresolution statistics for optical imaging signals: application to automated detection of odour activated glomeruli in the mouse olfactory bulb. NeuroImage 34 (3), 1020–1035.

Belluscio, L., Katz, L.C., 2001. Symmetry, stereotypy, and topography of odorant representations in mouse olfactory bulbs. J. Neurosci. 21 (6), 2113–2122.

Bozza, T., McGann, J.P., Mombaerts, P., Wachowiak, M., 2004. In vivo imaging of neuronal activity by targeted expression of a genetically encoded probe in the mouse. Neuron 42 (1), 9–21.

Chen, Z., Cichocki, A., 2005. Nonnegative matrix factorization with temporal smoothness and/or spatial decorrelation constraints. Technical Report, Laboratory for Advanced Brain Signal Processing, RIKEN, pp. 1–10.

Cichocki, A., Anh-Huy, P., 2009. Fast local algorithms for large scale nonnegative matrix and tensor factorizations. IEICE Trans. Fundam. Electron. Commun. Comput. Sci. E92-A (3), 708–721.

Donoho, D., Stodden, V., 2004. When does non-negative matrix factorization give a correct decomposition into parts? In: Thrun, S., Saul, L., Schölkopf, B. (Eds.), Advances in Neural Information Processing Systems, 16. MIT Press, Cambridge, MA.

Dorostkar, M.M., Dreosti, E., Odermatt, B., Lagnado, L., 2010. Computational processing of optical measurements of neuronal and synaptic activity in networks. J. Neurosci. Methods 188 (1), 141–150.

Firestein, S., 2001. How the olfactory system makes sense of scents. Nature 413 (6852), 211–218.

Grinvald, A., Shoham, D., Shmuel, A., Glaser, D., Vanzetta, I., Shtoyerman, E., Slovin, H., Wijnbergen, C., Hildesheim, R., Arieli, A., 1999. In-vivo optical imaging of cortical architecture and dynamics. In: Windhorst, U., Johansson, H.K. (Eds.), Modern Techniques in Neuroscience Research. Springer, Berlin Heidelberg, pp. 893–969.

Gurden, H., Uchida, N., Mainen, Z.F., 2006. Sensory-evoked intrinsic optical signals in the olfactory bulb are coupled to glutamate release and uptake. Neuron 52 (2), 335–345.

Hoyer, P.O., 2002. Non-negative sparse coding. Proceedings of the 12th IEEE Workshop on Neural Networks for Signal Processing. IEEE. IEEE, pp. 557–565.

Hoyer, P.O., 2004. Non-negative matrix factorization with sparseness constraints. J. Mach. Learn. Res. 5, 1457–1469.

Hyvärinen, A., 1999. Fast and robust fixed-point algorithms for independent component analysis. IEEE Trans. Neural Netw. 10 (3), 626–634.

Lee, T.-W., Lewicki, M., Girolami, M., Sejnowski, T., 1999. Blind source separation of more sources than mixtures using overcomplete representations. IEEE Signal Process. Lett. 6 (4), 87–90.

Ma, L., Qiu, Q., Gradwohl, S., Scott, A., Yu, E.Q., Alexander, R., Wiegraebe, W., Yu, C.R., 2012. Distributed representation of chemical features and tunotopic organization of glomeruli in the mouse olfactory bulb. Proc. Natl. Acad. Sci. U. S. A. 109 (14), 5481–5486.

Meister, M., Bonhoeffer, T., 2001. Tuning and topography in an odor map on the rat olfactory bulb. J. Neurosci. 21 (4), 1351–1360.

Mukamel, E.A., Nimmerjahn, A., Schnitzer, M.J., 2009. Automated analysis of cellular signals from large-scale calcium imaging data. Neuron 63 (6), 747–760.

Murthy, V.N., 2011. Olfactory maps in the brain. Annu. Rev. Neurosci. 34, 233–258.

Nara, K., Saraiva, L.R., Ye, X., Buck, L.B., 2011. A large-scale analysis of odor coding in the olfactory epithelium. J. Neurosci. 31 (25), 9179–9191.

Nelsen, R.B., 1998. An Introduction to Copulas, 1st edition. Springer, New York.

Pain, F., L'heureux, B., Gurden, H., 2011. Visualizing odor representation in the brain: a review of imaging techniques for the mapping of sensory activity in the olfactory glomeruli. Cell. Mol. Life Sci. 68 (16), 2689–2709.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: machine learning in Python. J. Mach. Learn. Res. 12 (Oct), 2825–2830.

Plumbley, M.D., 2003. Algorithms for nonnegative independent component analysis. IEEE Trans. Neural Netw. 14 (3), 534–543.

Reidl, J., Starke, J., Omer, D.B., Grinvald, A., Spors, H., 2007. Independent component analysis of high-resolution imaging data identifies distinct functional domains. NeuroImage 34 (1), 94–108.

Royet, J.P., Souchier, C., Jourdan, F., Ploye, H., 1988. Morphometric study of the glomerular population in the mouse olfactory bulb: numerical density and size distribution along the rostrocaudal axis. J. Comp. Neurol. 270 (4), 559–568.

Rubin, B.D., Katz, L.C., 1999. Optical imaging of odorant representations in the mammalian olfactory bulb. Neuron 23 (3), 499–511.

Sachse, S., Galizia, C.G., 2002. Role of inhibition for temporal and spatial odor representation in olfactory output neurons: a calcium imaging study. J. Neurophysiol. 87 (2), 1106–1117.

Soucy, E.R., Albeanu, D.F., Fantana, A.L., Murthy, V.N., Meister, M., 2009. Precision and diversity in an odor map on the olfactory bulb. Nat. Neurosci. 12 (2), 210–220.

Spors, H., Grinvald, A., 2002. Spatio-temporal dynamics of odor representations in the mammalian olfactory bulb. Neuron 34 (2), 301–315.

Spors, H., Wachowiak, M., Cohen, L.B., Friedrich, R.W., 2006. Temporal dynamics and latency patterns of receptor neuron input to the olfactory bulb. J. Neurosci. 26 (4), 1247–1259.

Strauch, M., Galizia, C.G., 2012. Fast PCA for processing calcium-imaging data from the brain of *Drosophila melanogaster*. BMC Med. Inform. Decis. Mak. 12 (Suppl. 1), S2.

Strauch, M., Rein, J., Galizia, C.G., 2012. Signal extraction from movies of honeybee brain activity by convex analysis. Computational Advances in Bio and Medical Sciences (ICCABS), 2012 IEEE 2nd International Conference on. IEEE, pp. 1–6.

Valdés-Sosa, P.A., Vega-Hernández, M., Sánchez-Bornot, J.M., Martínez-Montes, E., Bobes, M.A., 2009. EEG source imaging with spatio-temporal tomographic nonnegative independent component analysis. Hum. Brain Mapp. 30 (6), 1898–1910.

Wachowiak, M., Cohen, L.B., 2001. Representation of odorants by receptor neuron input to the mouse olfactory bulb. Neuron 32 (4), 723–735.

Wachowiak, M., Cohen, L.B., 2003. Correspondence between odorant-evoked patterns of receptor neuron input and intrinsic optical signals in the mouse olfactory bulb. J. Neurophysiol. 89 (3), 1623–1639.