Compressive light-field microscopy for 3D neural activity recording: supplementary material

NICOLAS C. PÉGARD^{1,2,*}, HSIU-YUAN LIU², NICK ANTIPA², MAXIMILLIAN GERLOCK², HILLEL ADESNIK¹, AND LAURA WALLER²

¹Department of Molecular & Cell Biology, University of California, Berkeley, CA 94720

²Department of Electrical Engineering & Computer Sciences, University of California, Berkeley, CA 94720

*Corresponding author: nicolas.pegard@m4x.org

Published 12 May 2016

This document provides supplementary information to "Compressive light field microscopy for 3D neural activity recording," <http://dx.doi.org/10.1364/optica.3.000517>. © 2016 Optical Society of

America

<http://dx.doi.org/10.1364/optica.3.000517.s001>

1. LIGHT FIELD MICROSCOPE DESIGN PARAMETERS

A detailed schematic for the imaging path of our setup is shown in Fig. S1. The focal length of the micro-lenses is chosen such that the range of ray angles subtended by each micro-lens is approximately matched to that of the imaging plane of the microscope. This means that the numerical aperture of the microscope objective matches that of the micro-lens array after accounting for scaling due to magnification in the imaging system. Quantitatively, the numerical aperture of the micro-lenses NA_{ML} approximately satisfies:

$$NA_{ML} = \frac{f_{\mu}}{f_{TL}} NA = 0.014, \quad (1)$$

where $NA = 0.5$ is the numerical aperture of our objective lens (adjusting for water immersion with refractive index $n = 1.33$), $f_{\mu} = 5mm$ is the effective focal length of the microscope objective and $f_{TL} = 180mm$ is the focal length of the tube lens.

For experiments on Zebrafish, light is provided through the objective by an LED centered on the $\lambda = 480nm$ excitation wavelength. The full field illumination power was between 25mW and 30mW for our experiments, and no significant photobleaching was observed throughout experimental recordings at full video framerate. With these settings, photobleaching onset was noticeable after 30 minutes of cumulative exposure as experiments were repeated on the same sample without recovery between trials.

2. IMPROVING SPATIAL SPARSITY WITH ICA

A. Non-negative matrix factorization

The optional pre-processing step for the raw video data is an Independent Component Analysis (ICA). Our training data may be captured either separately (before or after the experiment) as a calibration step, or extracted from the experiment data. It

is convenient to use the data itself, since there will then be no active neurons in the experiment that were not active during the training. When the training data is not sparse, ICA allows it to be decomposed into a sparse set of components, provided a large enough training data set. The ICA step is necessary in areas of the brain where neural activity is too high to allow direct acquisition of spatially sparse frames. Further, greater sparsity gives better localization in our algorithm, so the ICA step may improve the quality of the result.

The number of frames to be acquired for the training dataset depends on the type of brain tissue being studied. A large enough number of frames will ensure that all neurons within the field of view appear in at least one of the recorded frames, even when neural activity is naturally sparse. In brain tissue with high background neural activity, we oppositely seek to increase image diversity in the training dataset to facilitate the separation into individual neural structures. For this, we can either increase the frame rate of the camera capturing the information (at the cost of noise performance), or use ICA to identify temporally independent components to improve spatial sparsity.

3. SOURCE IDENTIFICATION THROUGH SCATTERING

A. Spatial sparsity requirement

Each light field signature analyzed should be sparse in the sense that only a limited number of neurons are active. Our model does not require that the light field space itself be sparse (mostly zeros), but rather that the point source distribution (3D emission pattern) should be sparse. Nevertheless, localization accuracy suffers with increased optical scattering or with increased density of fluorescent sources in the volume of interest. Here, we show simulations to illustrate the interplay of the two.

By increasing the density of sources or increasing optical

Fig. S1. Experimental setup (imaging path) for light field data acquisition. We use a $40\times$, $f = 5\text{mm}$, $\text{NA} = 0.5$ microscope objective (water immersion) and a $f = 180\text{mm}$ tube lens. The field of view is a $200\mu\text{m} \times 200\mu\text{m}$ square at the sample. The recorded signal, transformed by the micro-lens array is a 2D sampling of the 4D light field. We use a $4f$ system to relay the light field sampling onto the camera sensor because the camera does not allow close proximity. Relay optics further magnify the signal to be recorded so as to assign a square of edge length $N_p = 40$ pixels for each micro-lens.

scattering, the ability to correctly detect and recover 3D positions of sources degrades. In Fig. S2a, we show simulations of 2D space-angle slices for the 4D light field as sample point source density increases, and as scattering (σ) increases. Point sources distributed randomly in a 3D volume emit spherical wave-fronts. In the imaging plane, this corresponds to 2D hyperplanes in the 4D light field domain (x, y, θ_x, θ_y) or lines in a 2D space-angle slice (see Fig. S2a). The slope of each line corresponds to the depth of the source along the optical (z) axis. Without scattering, individual sources and their 3D positions are easy to discern and our algorithm performs well. By increasing the amount of scattering, however, the light field domain blurs (see Fig. S2a). Points that are deeper have more blurring and are more tilted. This information is captured by our forward model, so we may still be able to localize these points from the 4D blobs that they generate in light field space. If, for example, only one point source exists in a light field signature, it will create a distinct blob in light field space that can be fit to the forward model in order to localize the source. Note that the 2D space-angle slices shown in this figure appear completely diffused for the high scattering case, yet our algorithm can still pick up on the structure in 4D space, which contains far more information than the 2D slices shown here. Although the light field domain incurs blurring due to scattering, structure remains intact in the 4D space long after the 2D image is completely diffused.

B. Quantifying algorithm performance

To quantify the performance of our algorithm with varying sparsity for the practical case of brain tissue imaging, we show in Fig. S2b how performance decreases as the number of sources (neural structures) increases. We choose $\sigma = 0.052$ to approximately match the amount of scattering expected in mouse brain tissue. Let N be the number of actual point sources in the simulated model, N_d the number of recovered point sources, and N_{td} the number of correctly identified point sources. We define *precision* as the ratio of true detection over total number of detected sources in the algorithm's output (N_{td}/N_d) and *recall* as the ratio of true detection over total number of point sources to be recovered (N_{td}/N). The F-score is defined as the balanced harmonic mean of precision and recall:

$$F = 2 \times \left(\frac{1}{\text{precision}} + \frac{1}{\text{recall}} \right)^{-1} = \frac{2N_{td}}{N_d + N}. \quad (2)$$

For low point density and little scattering, we can recover all the points, so both precision and recall are 1 and the F-score is 1, as expected. With more point sources, the F-score decreases

as the sparsity condition is no longer satisfied, and the retrieval of the 3D position becomes inaccurate. This plot suggests that we should aim for approximately 10 or fewer active neurons in each sparse component of our light field measurement, after ICA decomposition. In our experiment, we use the distinguishability, D , to quantify the sparsity level in a given frame and ensure that data is decomposed into sparse components that all satisfy the sparsity constraint during the ICA pre-processing step.

C. Choosing the number of individual components for NMF

The number of individual components in the training data, N_k , is determined by decomposing matrix $\mathbf{I}$ into its singular values. We write, $\mathbf{I} = \mathbf{U}\Sigma\mathbf{V}$, where $\mathbf{U}$ and $\mathbf{V}$ are unitary matrices, and Σ is a rectangular, diagonal matrix. The diagonal elements of Σ , or singular values, are ranked in decreasing order. Without noise, there should be a small number of singular values to describe the information. In the presence of noise, however, $\mathbf{I}$ is generally a full-rank matrix and the singular values, $\Sigma_{k,k}$, are all strictly positive. To identify the number of relevant independent components, N_k , despite the presence of noise, we compare the distribution of singular values for the training data to the same spectrum measured from background frames recorded without signal, but with identical noise. Results are shown in Fig S3 for a training data set of 500 frames and the effective rank is calculated by comparing the spectrum of singular values with and without signal. The two plots merge eventually, suggesting a decomposition into $N_k = 40$ components.

D. Comparison with other methods based on image reconstruction

In Figure S4, we show how sparse independent components are individually decomposed into individual light field signatures. First, neural structures are detected and localized in 3D by our forward model. Then, the raw frame or component is divided into separate signatures using a proximity rule based on the forward model simulation. In the presence of scattering, the clear advantage of our method when compared to threshold detection after volume image reconstruction is even more visible (See Supplementary Movie 2)

4. COMPUTATIONAL EFFICIENCY

The algorithms for NMF, light field signature identification and activity retrieval were implemented on a commercial desktop computer (Intel I7, 3GHz processor) with 16Gb of memory. Algorithms were written with MATLAB and Python languages, and execution was performed without GPU or parallel computing.

Fig. S2. As density of neurons or scattering increases, our localization accuracy degrades. (a) Simulated light field space-angle slices for randomly distributed point-emitters in a homogenous 3D scattering medium, for various values of density, N/V , and optical scattering, σ . Each point source generates a line in the 2D slice view, with deeper point sources having blurred and tilted lines. As both source density and scattering increase, 2D slice views become less clear; however, the spatial sparsity constraint can remain satisfied when considering the full 4D light field distribution and our scattering forward model. (b) We select a fixed scattering magnitude $\sigma = 0.052$ that approximates typical values for the somatosensory cortex of a mouse brain and repeat simulation experiments with a gradually increasing number of point sources (active neurons). In each simulation, we compute a recovery performance score, or F-score.

Fig. S3. To determine the number of independent components, we decompose the training data, $\mathbf{I}$, into its singular values. Since the data contains noise, $\mathbf{I}$ is generally a full-rank matrix with strictly positive singular values, $\Sigma_{k,k}$ (blue curve). To avoid the noise-only parts, we compare the training data SVD to that of the noise-only background (green curve) and quantify the number of singular values that are due to the neural signals - defined as the “effective rank”, which is $N_k = 40$ here.

Fig. S4. (a) - We start with an unprocessed dataframe or, as in this case, an independent component computed from a larger training data set that satisfies the spatial sparsity condition for source identification. (b) Our method decomposes the component into signatures, which are here displayed together with various random colors for illustrative purposes. (c,d) Each identified source provides an estimate of the source localization, here displayed (c) in 2D over a spatial slice of the light field domain, and (d) in 3D (red dots). For comparison, threshold based methods ((c), blue iso-value curve and (d), blue iso-value surface) clearly identify the brightest spots but fail at detecting several neural structures that are clearly visible in an (x,y) slice of the 4D light-field domain.

The typical computation time for the decomposition of 500 full-size frames into 40 independent components with NMF was approximately 10 minutes. SVD decomposition is computationally simpler than NMF and a reliable estimate of N_k was obtained in one minute. Individual neuron localization and the computation of light field signatures takes about 2 minutes per component, and the complete extraction of the dictionary of 802

neural structures was therefore completed in little over one hour.

The dictionary building is a one-time calibration step that then enables fast computation of neural activity in any subsequent acquired frame. 10s of experimental data were recorded at 100Hz on zebrafish. A quantitative measure of fluorescence activity by decomposition of the experimental data on the dictionary (802 light field signatures) was completed in 30 seconds.

It should be noted that the current implementation allows online retrieval of neural activity at 33Hz. Furthermore, using data structures that take advantage of sparsity, as well as better reconstruction models, it is possible to significantly reduce computational complexity. Faster computing hardware may be used if fast online retrieval of neural activity is required.

5. LIMITATIONS

The most practically limiting assumption of our method is that the sample is stationary throughout both the training data collection and the experiment. In the training process, sample displacements can be measured and data-frames that are recorded when the sample is not at rest can simply be ignored. During the zebrafish experiment, motion artifacts temporarily prevent accurate quantification of neural activity until the animal returns to its resting position. This issue can be mitigated in future work by digital motion correction algorithms that operate in light field space, active mechanical motion correction, or solutions that miniaturize the imaging hardware and directly attach to the brain of unrestrained animals, eliminating motion.

Surface scattering at the brain-dura matter interface and inhomogeneous scattering in oriented structures such as inner-brain layers are not currently taken into account by our forward model. These may present problems in mouse brain experiments, but can be corrected by more sophisticated forward models. For example, an inhomogeneous scattering model, either pre-measured or solved for simultaneously, could fix model mismatch of this type. Similar approaches could be taken when occlusions (e.g. blood vessels) are present.

The other assumption of our method that may break down in practice is that the temporal and spatial distribution of fluorescence activity are fully independent. This assumption remains valid as long as the characteristic time-scale of calcium diffusion within the neuron's cytoplasm is faster than the temporal resolution of the sensor. Here, at 100Hz in zebrafish, we did not detect any evidence of slow calcium diffusion. We note, however, that this may potentially become an issue in mouse brain tissue, and in the presence of strong dendritic fluorescence. Current research on functional-imaging dyes (e.g. voltage sensitive membranes [1]) indicates that fast-responding functional markers, as well as localized expression around the nucleus of the calcium markers [2] will prevent this issue in the future, and improve the quality of recordings, even at high speeds.

REFERENCES

1. C. Petersen, A. Grinvald, and B. Sakmann, "Spatiotemporal dynamics of sensory responses in layer 2/3 of rat barrel cortex measured in vivo by voltage-sensitive dye imaging combined with whole-cell voltage recordings and neuron reconstructions," *Journal of Neuroscience* **23**, 1298–1309 (2003).
2. C. K. Kim, A. Miri, L. C. Leung, A. Berndt, P. Mourrain, D. W. Tank, and R. D. Burdine, "Prolonged, brain-wide expression of nuclear-localized gcamp3 for functional circuit mapping," *Frontiers in neural circuits* **8** (2014).