

SCIENTIFIC REPORTS



OPEN

Compressive Volumetric Light-Field Excitation

David C. Schedl & Oliver Bimber

We explain how volumetric light-field excitation can be converted to a process that entirely avoids 3D reconstruction, deconvolution, and calibration of optical elements while taking scattering in the probe better into account. For spatially static probes, this is achieved by an efficient (one-time) light-transport sampling and light-field factorization. Individual probe particles (and arbitrary combinations thereof) can subsequently be excited in a dynamically controlled way while still supporting volumetric reconstruction of the entire probe in real-time based on a single light-field recording.

Volumetric recording of fluorescent activity is an essential tool in modern microscopy (e.g., for optogenetics^{1–7}). Achieving fast readouts remains challenging, since the majority of microscopy techniques record volumes by scanning (e.g., confocal microscopy, light-sheet scanning⁸, and one- or two-photon SLM-based scanning techniques^{9,10}).

Light-Field Microscopy (LFM)^{7,11,12}, in contrast, captures the 4D incident light (e.g. by means of a microlens array (MLA) placed at the intermediate image plane of the imaging path), and thus supports the computation of a focal stack based on a single sensor recording. Optical phenomena in thick tissue such as scattering and optical aberrations prevent the application of traditional imaging techniques. Light-field recordings, however, encode these effects.

Pegard *et al.*¹³ applied LFM imaging to distinguish and localize 3D neuronal activity by non-negative matrix factorization¹⁴ during a training phase (a light-field video recorded at constant illumination during which the neurons fired randomly) to separate overlapping light-field signatures. Decomposition works as long as only few neurons fire simultaneously and at least once independently from other neurons in the training period (i.e., neural activity must be sparse in the temporal and spatial domains). Training generates a database of light-field signatures that allows fast readouts of neuronal activity while completely avoiding any image formation process.

We have explained in previous work¹⁵ how to concentrate light simultaneously at multiple selected volumetric positions by means of a 4D illumination light field that is computed directly from the appearance of the probe. We refer to this as *volumetric light-field excitation* (VLE), which may find applications in the field of optogenetics. In contrast to holographic projection with phase modulators^{9,16,17}, light-field projection with spatial-light modulators (e.g. in combination with MLAs or realized with multiple interplaying SLMs^{18,19}) does neither limit the excitation area or the illumination pattern²⁰, nor suffers from spatially varying diffraction efficiency and the presence of zero-order diffraction spots, ghosting, and intensity fluctuations (speckles)^{3,21}. With one sensor recording, we capture a 4D light field of the probe under full illumination and apply synthetic aperture rendering^{22,23} to compute a 3D focal stack or individual perspective images from it. The focal stack undergoes 3D deconvolution²⁴ to produce a defocus-free z-stack. Within the z-stack, we select parts of the probe (3D-segmented fluorescent microspheres of typical neuron sizes of model organisms, such as *C. elegans* and zebrafish larvae) that are to be excited. From the selection in the z-stack, we then determine a 4D light-field mask that is projected simultaneously into the probe such that the light concentrated at the volumetric positions of the selection is maximized while the illumination of other regions is minimized. This approach, however, has several essential drawbacks: It requires precise 3D reconstruction of the probe from a focal stack. As deconvolution is ill-posed, this is not always possible and can lead to reconstruction errors. Furthermore, it requires precise calibration of all optical elements in the imaging and illumination paths of the microscope, and ignores scattering in the probe.

In this article, we explain how VLE can be converted to a process that entirely avoids 3D reconstruction, deconvolution, and calibration of optical elements while taking scattering in the probe better into account. For spatially static probes, this is achieved by an efficient (one-time) light-transport sampling and light-field factorization. Individual probe particles (and arbitrary combinations thereof) can subsequently be excited in a dynamically controlled way while still supporting volumetric reconstruction of the entire probe in real-time based on a single light-field recording.

Faculty of Engineering and Natural Sciences, Johannes Kepler University, Linz, 4040, Austria. Correspondence and requests for materials should be addressed to O.B. (email: oliver.bimber@jku.at)

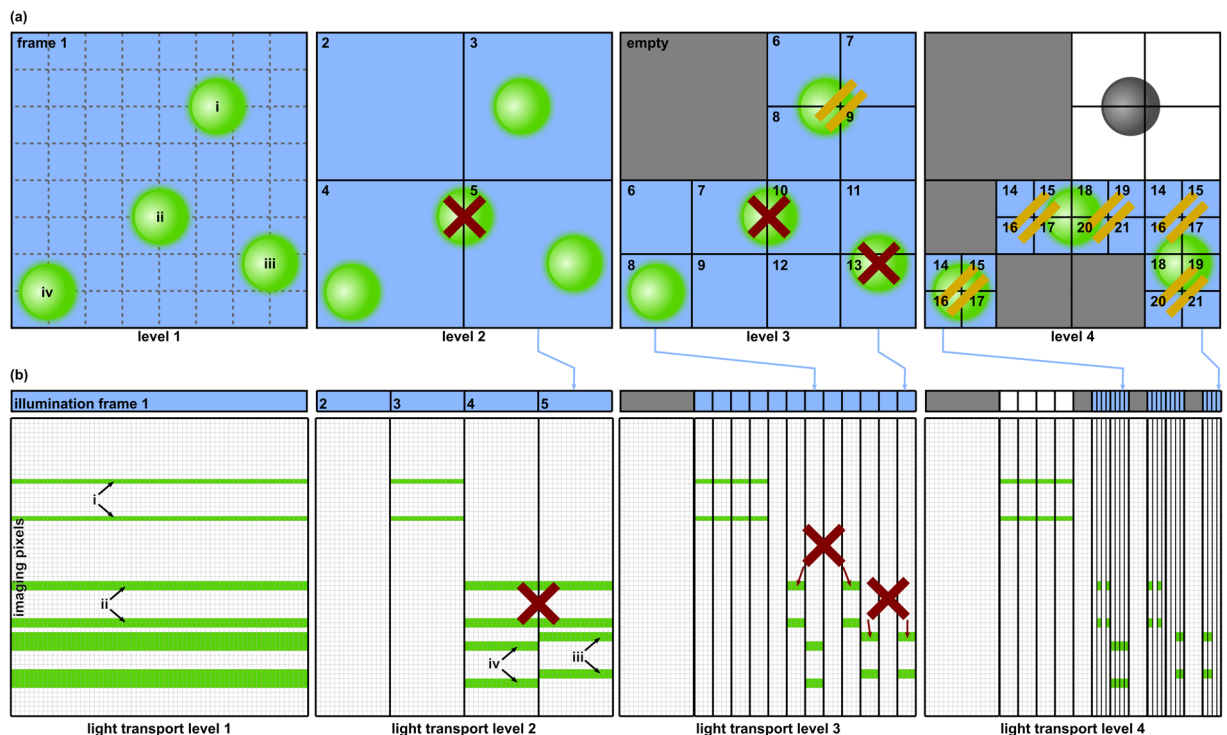


Figure 1. Our algorithm illustrated by a simplified 2D example with four isotropic particles (i to iv). **(a)** Top view of hierarchical illumination levels. We start by recording a frame under full illumination that is subdivided recursively. Non-conflicting frames recorded in parallel are denoted by the same frame number. Illumination rays that cause blank recordings are rejected (indicated by grey color). **(b)** Estimated versions of the light-transport matrix (64×64) after scanning through various levels. Conflicts are detected, and our isotropy hypothesis is applied at each level. Rows in the matrix correspond to recorded imaging rays, while columns are related to illumination rays. Emission signatures of particles (i.e., rows in the light transport) are indicated in green. Conflicts are marked by crosses, while isotropic similarities are marked by two diagonal bars. Illumination of particle i is stopped after 3 subdivision levels due to isotropic emission. Overall, the algorithm stops after scanning four hierarchical levels in the example presented.

Compressive VLE

Equation 1 outlines the principle of *compressive volumetric light-field excitation* (CVLE):

$$[T][\Phi] = [Y] = [\Psi][X] \quad (1)$$

Let Φ be a sequence of illumination light fields (each column of the matrix represents one illumination light field), and Y be the corresponding sequence of imaging light fields (each column y_i represents the imaging light field recorded for the corresponding illumination light field in ϕ_i), then the light-transport matrix T represents the entire transport of light from ϕ through the probe to y in the sense of $y = T\phi$.

If T were known, then, for a given desired imaging light field y_p , we could determine the corresponding illumination light field ϕ_i with $\phi_i = T^{-1}y_i$ (or by solving $y_i = T\phi_i$ for ϕ_i). Determining T by brute-force calibration (i.e., Φ is identity), however, is infeasible, as the number of measurements would have to match the resolution of the illumination light field (i.e. in the range of several thousands to more than one million).

Nevertheless, if we can assume that the emission of excited particles in the probe is isotropic, T contains a significant amount of redundancy. Compressive sensing principles²⁵ can be applied to reduce the number of measurements to obtain a representation of T that suits our isotropy assumption.

If Φ is a sampling sequence, Y the corresponding imaging light fields and the probe isotropic, then there might exist a basis Ψ in which Y can be expressed with few coefficients X such that $Y = \Psi X$. Note that matrices Ψ and X are much smaller than Y (i.e., contain fewer coefficients) and are thus compressed. The challenge in compressive sensing is to choose a proper sampling Φ such that Y is compressible. Furthermore, the number of samples in Φ can be (far) smaller than required by the Nyquist-Shannon sampling theorem (i.e., Φ is identity).

In fact, Ψ and X can be determined directly from Y by means of an efficient independent component analysis (ICA), as demonstrated by Pegard *et al.*¹³ for light-field microscopy imaging. This could be traditional non-negative matrix factorization¹⁴, one of its variants¹³, or a modern dictionary learning technique^{26,27}. In contrast to the Pegard *et al.* approach¹³, however, our sampling is adaptively controllable by choosing the sequence of illumination light fields in Φ .




Figure 2. Our algorithm applied to a 4D light-field simulation of four particles. **(a)** Example imaging light-field frames (MLA recordings) for scanning levels 1 to 8. **(b)** The ground truth and updated estimates of the light-transport matrix with corresponding mean squared errors (MSE). Note that unused and empty illumination and imaging rays (i.e., empty columns and rows) are not shown in the matrix. Furthermore, we down-sampled the imaging dimension of the light transport for better visibility (from 2,881,200 to 2352 rows). The light-transport sampling processing for simulated and measured data is also shown in the supplementary video.

Once sparse representations of T and Ψ have been found, we can finally determine an illumination light field ϕ_i for a desired imaging light field y_i , as explained above (i.e., with $\phi_i = T^{-1}y_i$ or by solving $y_i = T\phi_i$ for ϕ_i), composing the desired imaging light field y_i by a linear combination of the columns in Ψ .

A 3D reconstruction for the selection of the probe parts to be illuminated as in Schedl and Bimber's¹⁵, which requires deconvolution and precise optical calibration, is not necessary. Furthermore, Ψ also contains the information about scattering in the probe, which can now be considered.

Note that T and Ψ have to be determined only once per probe as part of a training phase. Such training by recording a light-field video of the probe under constant illumination is also required in the approaches of Pegard *et al.*¹³ and Pnevmatikakis *et al.*²⁸. After this training period, however, computation of volumetric light-field excitation patterns (i.e., solving for ϕ_i) as well as 3D imaging and visualization of the excited probe by means of light-field imaging and synthetic aperture rendering^{22,23} or model-driven 3D estimations of probe particles^{13,29} can still be achieved in real time (i.e., at the speed of the camera exposure required).

Light-Transport Sampling

Acquiring a robust estimate of the light-transport (T) more efficiently than with brute-force scanning is essential to computing accurate illumination light fields. We propose a parallel algorithm that acquires T for 4D light fields adaptively—similarly to the hierarchical 2D image sampling proposed by Sen *et al.*³⁰. Furthermore, we reduce the number of recordings by applying an isotropy hypothesis that determines whether finer illumination levels are needed.

For better illustration, we explain our algorithm with the help of the simplified 2D example in Fig. 1. Note, however, that in practice sampling is carried out in the 4D ray space's (spatial and directional) domains of illumination and imaging light fields.

We start by recording a frame (y_1) under full illumination (i.e., all illumination rays in ϕ_1 are on). This leads to a first estimate of the light transport by filling all columns of T with y_1 (reshaped as a column vector), where rows in T correspond to recorded imaging rays, while columns are related to illumination rays. In the next step, ϕ_1 is subdivided – in our 2D example in Fig. 1 into four equal sub-illuminations ϕ_2 to ϕ_5 . In hierarchical tree notation, the subdivided illuminations are lower-level branches of the full illumination. We then record frames




Figure 3. Simulated CVLE results (ground truth; after NMF; after reprojecting factorized signatures): **(a)** MLA image of color-coded particle light fields after reprojecting their factorized signatures. **(b)** Close-ups of each particle signature. **(c)** Close-ups of overlapping signatures in **(a)**. **(d)** Center light-field views with corresponding space-angle (epi) slices. Space-angle representation along dashed lines in center view indicates correct handling of occlusion cases (crossings of rays). The ACD between ground truth and final results (reprojected factorized signatures) is 0.0033. Simulated results are also shown in the supplementary video.

with the subdivided illuminations and estimate a finer version of T (i.e., we update the corresponding columns in T with the more accurate measurements). Sub-illuminations that cause blank imaging light fields (i.e., that do not excite any particles) are rejected. Figure 1 shows such blank regions at levels 2, 3 and 4. The algorithm continues by recursively subdividing the illumination light field, recording imaging light fields, and updating the light transport matrix T with each new measurement. For each intermediate estimate of T , we determine conflicts (i.e., imaging rays that are active for multiple sub-illuminations of the same tree level). In Fig. 1, these conflicts are indicated in the light-transport matrices by rows that contain non-zero coefficients in multiple columns across different sub-illuminations. Independent (i.e., non-conflicting) sub-illuminations can be scanned in parallel. This reduces the number of scans for estimating T significantly. In Fig. 1, frames are recorded in parallel at the subdivision levels 3 and 4. In the best case (i.e., without conflicts), our algorithm requires $4\log_4(m) + 1$ scans (where m is the total number of illumination rays), while brute-force scanning requires m scans. For our toy example in Fig. 1, this difference would amount to 13 to 64 scans ($\approx \times 5$ speed-up/80% compression). Due to conflicts, however, 21 scans are required in our example ($\approx \times 3$ speedup/67% compression). Note that the speed-up is significantly higher for realistic resolutions of the illumination light field, as we will show later. To achieve an additional speed-up, we consider the isotropic emission of the probe and avoid subdivisions down to the last (single-ray) level: We hypothesize that isotropically emitting particles cause the same imaging light-field signature (pattern) when hit by different illumination rays. Thus, we expect only a single particle to be illuminated

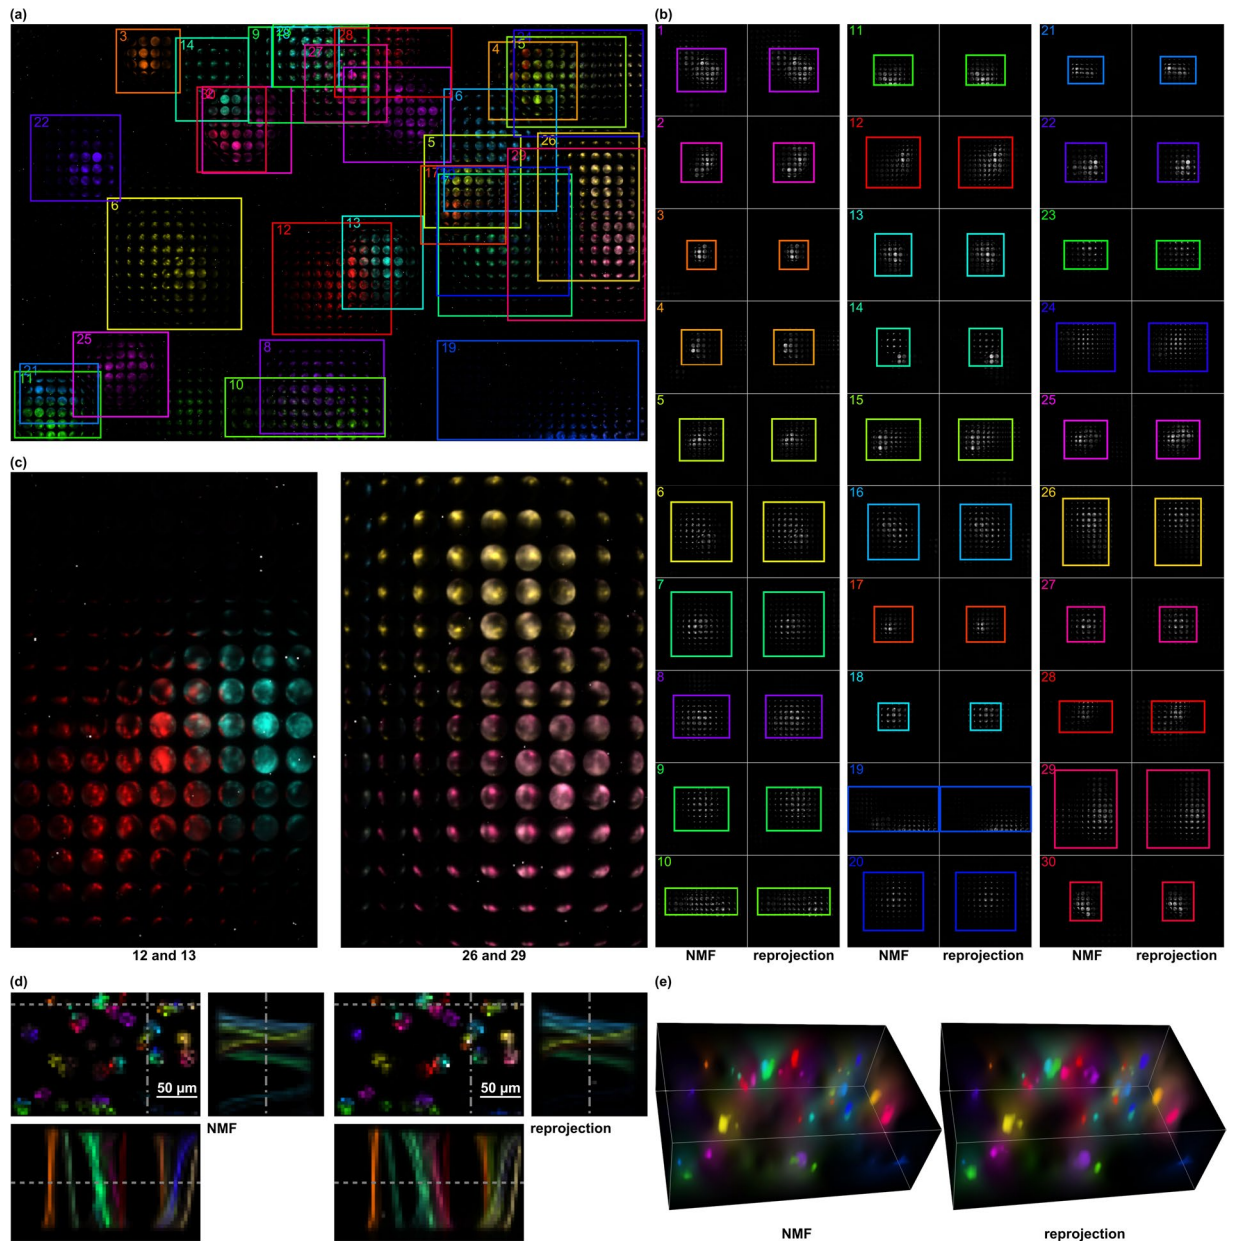


Figure 4. Measured CVLE results (after NMF; after reprojecting factorized signatures): **(a)** MLA image of color-coded particle light-fields after reprojecting their factorized signatures. **(b)** Close-ups of each particle signature. **(c)** Close-ups of overlapping signatures in **(a)**. **(d)** Center light-field views with corresponding space-angle (epi) slices. Space-angle representation along dashed lines in center view indicates correct handling of occlusion cases (crossings of rays). **(e)** Volumetric renderings (perspective projection along axial direction of center view). The ACD between NMF and final results (reprojected factorized signatures) is 0.1177. Note, that a ground truth does not exist in this case. Measured results are also shown in the supplementary video.

if the imaging rays of all non-blank children are similar to the imaging rays of their parent. We make no further subdivisions in this case.

Figure 2 illustrates an example similar to that in Fig. 1, but with our algorithm applied to realistic 4D light-field simulations (first subdividing the spatial domain and then the directional domain). In this case, 179 frames must be scanned for 8085 illumination rays to estimate a light-transport matrix that is close to the ground truth (i.e., the result of brute-force sampling). The speed-up is $\approx \times 45$ ($\approx 97.8\%$ compression).

Light-Field Factorization

We assume that the probe consists of k particles with isotropic emission when excited that are scanned by l illumination patterns as part of the light-transport sampling process. The recorded sequence of imaging light fields (y_1 to y_l) contains multiple contributions from these particles. We decompose Y into the independent components Ψ (i.e., light-field signatures of individual isotropic particles) such that

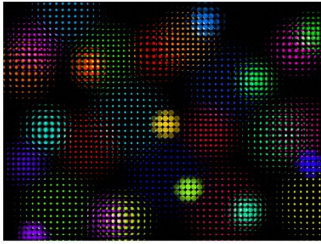
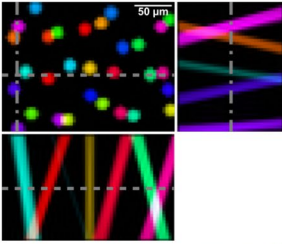
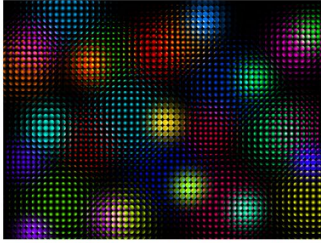
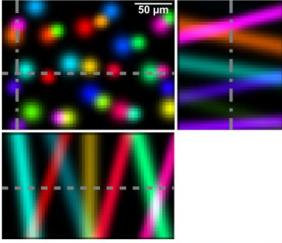
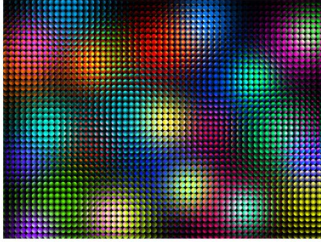
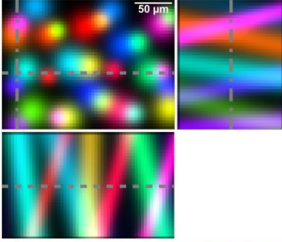
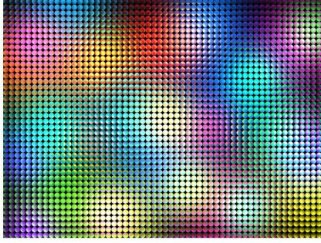
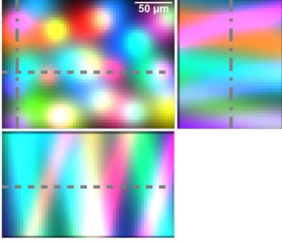
| σ | ground-truth MLA images | ground-truth center views and space-angle slices | error ACD | frames (speed-up / compr.) |
|----------|--|---|-----------|---------------------------------|
| 0.01 |  |  | 0.022 | 863 ($\times 9.4 / 89\%$) |
| 0.10 |  |  | 0.030 | 739 ($\times 10.9 / 91\%$) |
| 0.20 |  |  | 0.045 | 649 ($\times 12.5 / 92\%$) |
| 0.30 |  |  | 0.073 | 597 ($\times 13.5 / 93\%$) |

Figure 5. Simulated CVLE results with different amounts of scattering (σ). With increasing scattering, the signatures of individual particles become more isotropic, as shown in the MLA images and the center views with corresponding space-angle slices. The ACD between ground truth and reprojection represents only the error in the light-transport estimation, as we assume a perfect light-field factorization in this example. Results for simulated scattering are also shown in the supplementary video.

$$Y = \Psi X, \quad (2)$$

where the coefficients X represent the contributions of each particle's signature in Y .

While y_1 (recorded at full illumination) contains contributions from all particles (resulting in dense x_1), recordings under sparser illumination light fields (e.g., y_i) will result in less excited particles and therefore in sparse coefficients. By constraining $\Psi \geq 0$, $X \geq 0$ (which is valid for isotropic probes), the problem in Equation 2 can be solved by non-negative matrix factorization (NMF).

While NMF is an NP-hard problem, heuristics exist that are guaranteed to converge³¹. In our implementation we apply an active-set approach to alternating non-negative least squares³².

Results and Discussion

Figure 3 presents a light-field simulation of a more complex probe (dimensions: $100 \mu\text{m}$ axial, $175 \mu\text{m} \times 234 \mu\text{m}$ lateral) with 30 particles ($20 \mu\text{m}$ diameter fluorescent microbeads). The simulated light-field segments had (spatial \times directional) resolutions of $42 \times 56 \times 35 \times 35$ (imaging) and $11 \times 15 \times 7 \times 7$ (illumination). A total of 845 scans was required by our algorithm to estimate the light-transport matrix (speed-up of $\approx \times 10/90\%$ compression compared to brute-force scanning). Since a ground truth exists for the simulation, we can determine an average cosine distance (ACD) error of 0.0033 when comparing the ground truth light-field signatures with the light-field signatures that result from reprojecting the factorized signatures (i.e., from exciting single particles only). The latter is achieved by projecting the light-field illumination ϕ_i computed from $y_i = T\phi_i$, where y_i is the factorized light-field signature of a single particle. Figure 4 illustrates measurement results acquired in experiments




Figure 6. Optical layout of our light-field microscopy prototype. The 4D illumination pattern is generated by a spatial light modulator (SLM) and concentrated on the probe by a microlens array (MLA2) and the objective (OBJ). Emissive light from the probe is propagated on a camera sensor (CAM) via the objective (OBJ) and another microlens array (MLA1). The beam splitter (BS) and corresponding filters restrict the light in the optical paths to certain wavelengths. Example illumination and imaging light fields are shown on SLM and CAM.

using the prototype described in the methods section and under similar conditions as in the simulations shown above: around 30 ($10\ \mu\text{m}$ to $20\ \mu\text{m}$) fluorescent microbeads embedded in a polydimethylsiloxane carrier, scanned in a section of $100\ \mu\text{m}$ axial, $150\ \mu\text{m} \times 230\ \mu\text{m}$ lateral, with light-field resolutions (spatial \times directional) of $36 \times 55 \times 35 \times 35$ (imaging) and $11 \times 15 \times 5 \times 5$ (illumination). In total, our algorithm required 561 scans to estimate the light-transport matrix (speed-up $\times 7.4/86\%$ compression). Since a ground truth does not exist in this case, we can determine an ACD of 0.1177 only when comparing the factorized light-field signatures with the light-field signatures that result from reprojecting the factorized signatures (i.e., from exciting single particles only). Note that in contrast to the simulation, this can only indicate the quality of the light-transport matrix, but ignores the error introduced by the factorization.

The results presented above reveal that, up to a remaining precision error arising during light-transport sampling and light-field factorization, our approach supports simultaneous, controlled concentration of light at multiple volumetric positions by means of a 4D illumination light field. It avoids 3D reconstruction of the probe, deconvolution, and calibration of optical elements, while speeding up light-transport sampling by up to one order of magnitude (85–90% compression).

Since scattering is included in the light-transport matrix, it is taken into account when solving $y_i = T\phi_i$ for ϕ_i . Given that the factorized probe signatures y_i can be determined, scattering during the illumination is considered in ϕ_i .

Figure 5 presents simulation results for a probe with increasing amount of scattering. We model scattering as homogeneous and Gaussian-distributed in angle²⁹, where a higher scatter coefficient σ indicates more scattering. Additionally, we assume perfect factorization results. With increasing scattering, the reprojection error rises and the number of frames required for measuring the light-transport matrix decreases. The latter is the result of more isotropy from scattering in the light transport that causes our scanning to stop early.

Our approach has several limitations. First, NMF requires a manual estimate of the number of particles present in the probe. Evaluating more advanced factorization techniques, such as model-driven approaches^{13,29}, that

support automatic estimation and provide better factorization results will be part of our future work. Currently, our approach (like other factorization-based techniques^{13,28}) is limited to spatially static (i.e., fixed) probes. A dynamic update of T for tracked particles of a moving probe and experiments with living organisms will be investigated in the future. Furthermore, our hierarchical light-transport sampling results in low amounts of light at finer illumination levels, and thus leads to low signal-to-noise ratios or long camera exposure times. At present, we increase the exposure time with finer sampling levels. More advanced illumination strategies (e.g., similar to Hadamard patterns) and high-dynamic range camera systems will help to overcome this issue. The limited spatial and directional resolutions of our prototype can be overcome by applying higher resolution SLMs, imaging sensors, and MLAs.

Method

In our experiments we used a light-field microscope prototype¹² as shown in Fig. 6 and explained in more detail by Schedl and Bimber¹⁵. We equipped a $60\times/1.2\text{NA}$ objective with a $250\text{ }\mu\text{m}$ imaging (MLA1) and $300\text{ }\mu\text{m}$ illumination MLA (MLA2) of focal lengths $6250\text{ }\mu\text{m}$ and $7500\text{ }\mu\text{m}$, respectively. We used a Texas Instruments Digital Light Processor at a resolution of 340×256 (pixel size $41.1\text{ }\mu\text{m}$) as spatial light modulator (SLM). Note that we used the SLM at reduced resolution when compared to Schedl and Bimber¹⁵. The light source was an $120\text{ }\mu\text{m}$ metal halide lamp (X-Cite $120\text{ }\mu\text{m}$, EXFO) and we placed the corresponding filters in a standard Nikon microscopic filter cube. For recording (CAM) we used a Retiga 4000 R monochrome camera with 2048×2048 resolution (pixel size $7.4\text{ }\mu\text{m}$). In front of the imaging and illumination MLAs we placed 1:1 relay lenses; two Nikon AF Nikkor with focal length 50 mm mounted nose-to-nose and a telecentric relay lens from Brilliant Technologies, respectively.

The fluorescent microbeads used in the experiment shown in Fig. 4 were $10\text{ }\mu\text{m}$ to $20\text{ }\mu\text{m}$ Green Polyethylene microspheres (material density 0.99 to 1.01 g cm^{-3} ; peak excitation 470 nm ; peak emission 505 nm ; distributor cospheric). The polydimethylsiloxane carrier in which they were embedded was Elastosil RT 604 from Wacker (material density 0.97 g cm^{-3}).

For solving Eqn. 2 (NMF) we used the Matlab implementation of the non-negative least-squares active-set algorithm from³² with 200 iterations. When reprojecting individual signatures (solving for ϕ_i in $y_i = T\phi_i$) we used 500 iterations of the simultaneous algebraic reconstruction technique (SART)³³ and binarized the 4D illumination pattern. The source code for our simulated CVLE experiments (i.e., Figs 3 and 5) is available at <https://github.com/JKU-ICG/CVLE>.

References

- Lima, S. Q. & Miesenböck, G. Remote control of behavior through genetically targeted photostimulation of neurons. *Cell* **121**, 141–152 (2005).
- Boyden, E. S., Zhang, F., Bamberg, E., Nagel, G. & Deisseroth, K. Millisecond-timescale, genetically targeted optical control of neural activity. *Nat. Neurosci.* **8**, 1263–1268 (2005).
- Vaziri, A. & Emiliani, V. Reshaping the optical dimension in optogenetics. *Curr. Opin. Neurobiol.* **22**, 128–137 (2012).
- Packer, A. M., Roska, B. & Häusser, M. Targeting neurons and photons for optogenetics. *Nat. Neurosci.* **16**, 805–815 (2013).
- Deisseroth, K. & Schnitzer, M. J. Engineering approaches to illuminating brain structure and dynamics. *Neuron* **80**, 568–577 (2013).
- Zalocusky, K. A., Fenno, L. E. & Deisseroth, K. Current challenges in optogenetics. In Hegemann, P. & Sigris, S. (eds.) *Optogenetics*, 23–33 (De Gruyter, 2013).
- Cruz Perez, C. *et al.* Calcium neuroimaging in behaving zebrafish larvae using a turn-key light field camera. *Journal of Biomedical Optics* **20**, 096009, <https://doi.org/10.1117/1.JBO.20.9.096009> (2015).
- Bouchard, M. B. *et al.* Swept confocally-aligned planar excitation (scapec) microscopy for high-speed volumetric imaging of behaving organisms. *Nat. Photonics* **9**, 113–119 (2015).
- Quirin, S., Peterka, D. S. & Yuste, R. Instantaneous three-dimensional sensing using spatial light modulator illumination with extended depth of field imaging. *Opt. Express* **21**, 16007–16021 (2013).
- Yang, W. *et al.* Simultaneous multi-plane imaging of neural circuits. *Neuron* **89**, 269–284 (2016).
- Levoy, M., Ng, R., Adams, A., Footer, M. & Horowitz, M. Light field microscopy. *ACM Trans. Graph.* **25**, 924–934 (2006).
- Levoy, M., Zhang, Z. & McDowall, I. Recording and controlling the 4d light field in a microscope using microlens arrays. *J. Microsc.* **235**, 144–162 (2009).
- Pégar, N. C. *et al.* Compressive light-field microscopy for 3d neural activity recording. *Optica* **3**, 517–524, <http://www.osapublishing.org/optica/abstract.cfm?URI=optica-3-5-517> (2016).
- Lee, D. D. & Seung, H. S. Algorithms for non-negative matrix factorization. In *In NIPS*, 556–562 (MIT Press, 2000).
- Schedl, D. C. & Bimber, O. Volumetric light-field excitation. *Scientific Reports* **6** (2016).
- Packer, A. M. *et al.* Two-photon optogenetics of dendritic spines and neural circuits. *Nat. Methods* **9**, 1202–1205 (2012).
- Paluch-Siegler, S. *et al.* All-optical bidirectional neural interfacing using hybrid multiphoton holographic optogenetic stimulation. *Neurophotonics* **2**, 031208–031208 (2015).
- Rückler, F. *et al.* Micro mirror arrays as high-resolution spatial light modulators for photoactivation and optogenetics. In *SPIE BiOS*, 85860U–85860U (International Society for Optics and Photonics, 2013).
- Rückler, F., Berndt, D., Heber, J. & Shorte, S. Photoactivation and optogenetics with micro mirror enhanced illumination. In *SPIE Photonics Europe*, 913017–913017 (International Society for Optics and Photonics, 2014).
- Zhang, Z., Barbastathis, G. & Levoy, M. Limitations of coherent computer generated holograms. In *Digital Holography and Three-Dimensional Imaging*, DTuB5 (Optical Society of America, 2011).
- Golan, L., Reutsky, L., Farah, N. & Shoham, S. Design and characteristics of holographic neural photo-stimulation systems. *J. Neural Eng.* **6**, 066004 (2009).
- Isaksen, A., McMillan, L. & Gortler, S. J. Dynamically reparameterized light fields. In *SIGGRAPH*, 297–306 (ACM, 2000).
- Levoy, M. *et al.* Synthetic aperture confocal imaging. *ACM Trans. Graph.* **23**, 825–834 (2004).
- Holmes, T. J. *et al.* Light microscopic images reconstructed by maximum likelihood deconvolution. In Pawley, J. B. (ed.) *Handbook of biological confocal microscopy*, 389–402 (Springer, 1995).
- Tosic, I. & Frossard, P. Dictionary learning. *IEEE Signal Processing Magazine* **28**, 27–38 (2011).
- Olshausen, B. A. & Field, D. J. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision research* **37**, 3311–3325 (1997).
- Aharon, M., Elad, M. & Bruckstein, A. K-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on signal processing* **54**, 4311–4322 (2006).

28. Pnevmatikakis, E. *et al.* Simultaneous denoising, deconvolution, and demixing of calcium imaging data. *Neuron* **89**, 285–299, <http://www.sciencedirect.com/science/article/pii/S0896627315010843> (2016).
29. Liu, H.-Y. *et al.* 3d imaging in volumetric scattering media using phase-space measurements. *Optics express* **23**, 14461–14471 (2015).
30. Sen, P. *et al.* Dual photography. *ACM Trans. Graph.* **24**, 745–755 (2005).
31. Vavasis, S. A. On the complexity of nonnegative matrix factorization. *SIAM Journal on Optimization* **20**, 1364–1377 (2009).
32. Li, Y. & Ngom, A. The non-negative matrix factorization toolbox for biological data mining. *Source code for biology and medicine* **8**, 10 (2013).
33. Andersen, A. & Kak, A. C. Simultaneous algebraic reconstruction technique (sart): a superior implementation of the art algorithm. *Ultrason. Imaging* **6**, 81–94 (1984).

Acknowledgements

We thank Marc Levoy and Gordon Wetzstein of Stanford University, Enrico Geissler of Carl Zeiss AG, and Laura Waller of UC Berkeley for their support and for providing optical equipment.

Author Contributions

D.S. implemented the prototypes (soft- and hardware), performed the physical experiments, and collected measurements; O.B. developed the overall concepts, supervised the experiments and the implementations, and wrote the paper; both authors reviewed the manuscript.

Additional Information

Supplementary information accompanies this paper at <https://doi.org/10.1038/s41598-017-13136-2>.

Competing Interests: The authors declare that they have no competing interests.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2017