



Optics Letters

Time-multiplexed light field synthesis via factored Wigner distribution function

STEPHEN HAMANN,^{1,†} LIANG SHI,^{1,2,†} OLAV SOLGAARD,¹ AND GORDON WETZSTEIN^{1,*}

¹Department of Electrical Engineering, Stanford University, Stanford, California 94305, USA

²Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA

*Corresponding author: gordon.wetzstein@stanford.edu

Received 6 November 2017; revised 27 December 2017; accepted 4 January 2018; posted 4 January 2018 (Doc. ID 312874); published 31 January 2018

An optimization algorithm for preparing display-ready holographic elements (hogels) to synthesize a light field is outlined, and proof of concept is experimentally demonstrated. This method allows for higher-rank factorization, which can be used for time-multiplexing multiple frames for improved image quality, using phase-only and fully complex modulation with a single spatial light modulator. © 2018 Optical Society of America

OCIS codes: (090.2870) Holographic display; (090.1760) Computer holography; (100.6890) Three-dimensional image processing.

<https://doi.org/10.1364/OL.43.000599>

The relationship between the Wigner distribution function (WDF) and the light field has been discussed by Zhang and Levoy in 2009 [1] and by Goodman in 2013 [2]. The WDF, like the light field, provides a joint representation of space and spatial frequency (phase space), but additionally models the wave effects of diffraction and interference. These discussions reveal a compelling path for converting between light field recordings and complex wavefronts which completely describe the scene. In this Letter, we provide an algorithm to find an arbitrary-rank phase-only, amplitude-only or complex 2D pattern whose WDF best replicates a desired real-valued input light field, such as a white-light camera recording, through an alternating direction method of multipliers (ADMM) optimization. This method allows for higher-rank factorization; time-multiplexing multiple frames per image result in a better simulated peak signal-to-noise ratio (PSNR) on the reconstructed angular distribution and is beneficial to speckle reduction [3]. Explicit optimization of the WDF provides an alternative path to the holographic element (“hogel”) [4] generation using the Gerchberg–Saxton algorithm [5] or simulated annealing [6].

This algorithm is demonstrated experimentally for a number of chosen distributions using a display prototype with off-the-shelf optical components and a Holoeye Pluto VIS LCoS phase modulator. Both phase-only and combined phase and amplitude (complex) modulation are tested. Complex modulation is achieved by using phase grating “super-pixels” [7].

For simplicity, we limit our discussion to a 1D signal and its 2D WDF throughout the Letter, but extension to the higher-dimensional cases is straightforward. The WDF of a plane wave is defined as

$$W_{\text{plane}}(x, u) = \delta(u - u_0), \quad (1)$$

where u_0 characterizes the unique propagation angle $\theta = \sin^{-1}(u_0\lambda)$ at a wavelength of light λ .

A layer of microstructures modulates the complex amplitude of incident light multiplicatively and is generally modeled as a complex transmittance function $t(x) = a(x) \exp(i\phi(x))$, where $a(x)$ denotes the amplitude attenuation and $\phi(x)$ denotes the phase delay at spatial location x . The mutual correlation between any two points on the layer is given by the mutual intensity function [8]:

$$J(x, f_u) = \left\langle t\left(x + \frac{f_u}{2}\right) t^*\left(x - \frac{f_u}{2}\right) \right\rangle, \quad (2)$$

where the spatial offset f_u is the Fourier conjugate variable of spatial frequency u . In the far field, spatial frequency relates to angle as $\theta \approx u\lambda$ under the small angle approximation. The $\langle \cdot \rangle$ is a time-averaged ensemble and can be removed for coherent illumination. The WDF of the layer is defined as the Fourier transform of its mutual intensity along f_u :

$$W_{\text{Layer}}(x, u) = \int J(x, f_u) \exp(-2\pi i f_u u) df_u \quad (3)$$

The modulation between the incident light and the layer can be modeled as a convolution of their WDFs along the angular axis. For an incident plane wave, the output WDF is given by

$$W_{\text{out}}(x, u) = W_{\text{Layer}}(x, u) \otimes W_{\text{plane}}(x, u), \quad (4)$$

where \otimes is a convolution operator along the spatial frequency axis: $W_1 \otimes W_2 = \int W_1(x, u - u') W_2(x, u') du'$. Assuming normal incidence, the output WDF is equivalent to the WDF of the layer:

$$W_{\text{out}}(x, u) = W_{\text{Layer}}(x, u) \otimes \delta(u) = W_{\text{Layer}}(x, u). \quad (5)$$

Therefore, synthesizing a target light field using a single modulation layer under normally incident plane wave illumination is equivalent to optimizing the layer’s complex

transmittance function such that the resulting WDF best resembles the target.

Solving this problem globally for a high-resolution modulation layer requires storing and processing trillions of angular samples and, thus, is computationally intractable. Similar to Ye *et al.* [6] and Shi *et al.* [9], we employ a “holographic element” (hogel) model to limit the problem into a local area. Specifically, the entire layer is partitioned into non-overlapping local patches so that each is treated and optimized independently for a local light distribution. Though each hogel consists of multiple pixels, we treat one hogel as the minimally observable feature and are only interested in the average angular variation over its surface area.

Conceptually, a hogel functions as a lenslet in a light field camera. Computationally, this model splits the global problem into sub-problems that can be solved in parallel. The hogel size imposes a trade-off: a larger hogel yields higher angular resolution, but a lower spatial resolution.

We introduce the algorithm in the discrete case to allow for numerical optimization. Let $\mathbf{t} \in \mathbb{C}^N$ be the discrete complex transmittance function of a hogel and $\mathbf{h} \in \mathbb{R}_{\geq 0}^N$ be a target 1D light field. With coherent illumination across the hogel, the discrete mutual intensity function $\mathbf{J}_t \in \mathbb{C}^{N \times N}$ is

$$\mathbf{J}_t[x, f_u] = \mathbf{t} \left[x + \frac{f_u}{2} \right] \mathbf{t}^* \left[x - \frac{f_u}{2} \right], \quad (6)$$

where x and f_u are indices considered for the discrete case. The discrete WDF $\mathbf{W}_t \in \mathbb{R}^{N \times N}$ is followed by

$$\mathbf{W}_t[x, u] = \mathcal{F}_u(\mathbf{J}_t[x, f_u]), \quad (7)$$

where \mathcal{F}_u is the discrete 1D Fourier transform operator along f_u .

Due to the limitations of visual acuity, the human eye “averages” the variation over a hogel’s surface area. Mathematically, we define a projection matrix $\mathbf{P} \in \mathbb{R}^{N \times N^2}$ which sums each frequency component u of $W_t[x, u]$ over all spatial locations x :

$$(\mathbf{P} \text{vec}(\mathbf{W}_t))[u] = \sum_{x=1}^N \mathbf{W}_t[x, u], \quad (8)$$

where $\text{vec}(\cdot)$ is a linear operator that reshapes a matrix into a column vector by stacking up all the matrix columns. The light field synthesis problem is subsequently converted to the following optimization problem:

$$\underset{\mathbf{t}}{\text{minimize}} \quad \|\mathbf{h} - \mathbf{P} \text{vec}(\mathbf{W}_t)\|_2^2. \quad (9)$$

We solve this problem using ADMM [10], and Eq. (9) is reformulated as

$$\underset{\mathbf{t}}{\text{minimize}} \quad \|\text{vec}(\mathbf{W}_t) - \mathbf{I}\|_2^2, \quad (10)$$

$$\text{subject to } \mathbf{P}\mathbf{I} = \mathbf{h}, \quad (11)$$

where $\mathbf{I} \in \mathbb{R}^{N^2}$ is an intermediate WDF (vectorized) which connects the new objective and the constraint, known as the splitting variable. Following the general ADMM strategy, we formulate the augmented Lagrangian of Eq. (10):

$$L_\rho(\mathbf{t}, \mathbf{I}, \mathbf{y}) = \|\text{vec}(\mathbf{W}_t) - \mathbf{I}\|_2^2 + \mathbf{y}^T(\mathbf{P}\mathbf{I} - \mathbf{h}) + \frac{\rho}{2} \|\mathbf{P}\mathbf{I} - \mathbf{h}\|_2^2, \quad (12)$$

where \mathbf{y} is the Lagrange multiplier (dual variable) and ρ is the penalty term for violation of the constraint. Using the scaled

form of the augmented Lagrangian, the following iterative updates rules can be derived:

$$\mathbf{I} \leftarrow \underset{\mathbf{I}}{\text{argmin}} \quad \|\text{vec}(\mathbf{W}_t) - \mathbf{I}\|_2^2 + \frac{\rho}{2} \|\mathbf{P}\mathbf{I} - \mathbf{h} + \mathbf{u}\|_2^2, \quad (13)$$

$$\mathbf{t} \leftarrow \underset{\mathbf{t}}{\text{argmin}} \quad \|\text{vec}(\mathbf{W}_t) - \mathbf{I}\|_2^2, \quad (14)$$

$$\mathbf{u} \leftarrow \mathbf{u} + (\mathbf{P}\mathbf{I} - \mathbf{h}), \quad (15)$$

where $\mathbf{u} = (1/\rho)\mathbf{y}$ is used to simplify the notation of Eq. (12).

The \mathbf{I} -update is quadratic, thus convex, and can be iteratively solved by gradient-based methods. The \mathbf{t} -update seeks the choice of transmittance function \mathbf{t} whose WDF best approximates the updated intermediate WDF \mathbf{I} . Let $\text{ivec}(\cdot)$ be a linear operator that reshapes the vector back to its matrix form. Assuming a transmittance function $\hat{\mathbf{t}}$ exists whose WDF exactly equals $\text{ivec}(\mathbf{I})$, its mutual intensity will be given by

$$\mathbf{J}_{\hat{\mathbf{t}}}[x, f_u] = \mathcal{F}_u^{-1}(\text{ivec}(\mathbf{I})), \quad (16)$$

where \mathcal{F}_u^{-1} is the 1D discrete inverse Fourier transform operator along u . Though $\hat{\mathbf{t}}$ may not physically exist because \mathbf{I} only serves as a minimizer of Eq. (13), we can still advance the \mathbf{t} -update by finding \mathbf{t} whose mutual intensity best resembles $\mathbf{J}_{\hat{\mathbf{t}}}$. To solve this new problem, we *partially* reconstruct the outer product $\mathbf{O}_{\hat{\mathbf{t}}} \in \mathbb{C}^{N \times N}$ of $\hat{\mathbf{t}}$ by filling its entries using $\mathbf{J}_{\hat{\mathbf{t}}}$, which consists of a subset of entries in $\mathbf{O}_{\hat{\mathbf{t}}}$. We define a binary weight matrix $\mathbf{M} \in \mathbb{B}^{N \times N}$ which equals 1 at indices where $\mathbf{O}_{\hat{\mathbf{t}}}$ finds a correspondence in $\mathbf{J}_{\hat{\mathbf{t}}}$ and 0 otherwise. This transforms the \mathbf{t} -update into a weighted complex-valued low-rank matrix factorization problem:

$$\underset{\mathbf{t}}{\mathbf{t} \leftarrow \text{argmin}} \quad \|\mathbf{M} \circ (\mathbf{O}_{\hat{\mathbf{t}}} - \mathbf{t}\mathbf{t}^H)\|_{\mathcal{F}}^2, \quad (17)$$

where H is the conjugate transpose operator, \circ denotes the Hadamard product for element-wise multiplication, and the Frobenius norm $\|\cdot\|_{\mathcal{F}}^2$ measures the sum of squared differences of all matrix elements. Equation (17) is minimized iteratively using gradient descent:

$$\mathbf{t}^{(q)} = \mathbf{t}^{(q-1)} + \alpha_t \mathbf{M} \circ (\mathbf{O}_{\hat{\mathbf{t}}} + \mathbf{O}_{\hat{\mathbf{t}}}^H - 2\mathbf{t}\mathbf{t}^H)\mathbf{t}, \quad (18)$$

where α_t is the step length at a typical scale of $1e^{-3}$.

The reformulated \mathbf{t} -subproblem allows for an arbitrary-rank factorization without the necessity to alter the update formula. By replacing $\mathbf{t} \in \mathbb{C}^N$ to $\mathbf{t} \in \mathbb{C}^{N \times K}$ for a given rank K , a partially coherent light field can be synthesized by time-multiplexing K frames. On the other hand, expanding $\mathbf{t} = |\mathbf{t}| \exp(i\theta)$ allows for phase-only or amplitude-only optimization and, thus, facilitates the generation of an amplitude-only or phase-only modulation layer. As an example, we provide the update formula for phase-only factorization in the following:

$$\theta^{(q)} = \theta^{(q-1)} + \alpha_\theta [(\mathbf{M} \circ \overline{\mathbf{O}_{\hat{\mathbf{t}}} - \mathbf{t}\mathbf{t}^H}) \circ \mathbf{t} - (\mathbf{M} \circ (\mathbf{O}_{\hat{\mathbf{t}}} - \mathbf{t}\mathbf{t}^H)) \circ \bar{\mathbf{t}}], \quad (19)$$

where α_θ is the phase step length at a typical scale of $1e^{-3}$ and $\bar{\cdot}$ is the conjugate operator. The amplitude update formula can be derived similarly.

We implemented the proposed algorithm in MATLAB and sped up the WDF calculation and \mathbf{t} -update by implementing it as a mex module. On an Intel i7-6800K 3.4 GHz processor with 16GB RAM, one ADMM iteration on a 13×13 hogel takes 0.1 s. GPU implementation may provide

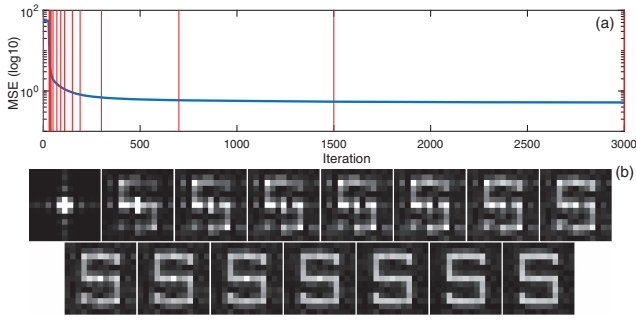


Fig. 1. (a) Convergence of the mean square error between synthesized “S” (phase-only) and target “S” in \log_{10} space. (b) Synthesized “S” at iterations marked by the red lines.

orders-of-magnitude speed improvement over our current implementation.

To assess the proposed algorithm, we simulate the phase-only results for a set of chosen angular profiles (13×13) and demonstrate proof-of-concept experimental results on a liquid crystal on silicon (LCoS) device.

Figure 1 shows the convergence of the proposed algorithm for an “S” shaped angular profile and intermediate results over the iterations, where the “S” distribution is gradually approached. We observe that for an arbitrary angular profile, 1000 iterations typically suffice to reach the convergence. Figure 2(a) visualizes the reconstruction of four different angular distributions. The visual quality consistently improves as the rank increases. The artifacts seen in the rank-1 results are barely visible in most of the rank-3 results. Figure 2(b) plots the PSNR versus the rank. The PSNR asymptotically improves for all the chosen angular profiles as a function of the rank, with little improvement seen past rank 5.

The optical setup in Fig. 3 is designed to view the light field angular response of calculated phase patterns in a focal conjugate plane. The lightwave emitted from a 532 nm laser diode is expanded by two lenses and followed by a nonpolarizing beam splitter to enable normal incidence on the spatial light modulator. A pulsed laser with appropriate coherence length across the hogel could also be used. The LCoS phase-only modulator displays the patterns. A 4F system of two lenses with a Fourier filter is used to pass the zeroth order for phase-only modulation or the first order for phase and amplitude complex modulation. Two more lenses are used to image the conjugate plane of the

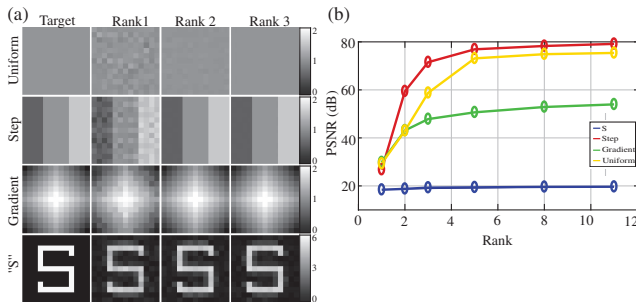


Fig. 2. (a) Synthesized (phase-only) “Uniform,” “Step,” “Gradient,” and “S” angular distributions over different ranks. (b) PSNR of synthesized results versus rank for “Uniform,” “Step,” “Gradient,” and “S” angular distributions.




Fig. 3. Optical setup for synthesizing the light field.

4F system onto a CMOS camera sensor, and the last lens aperture is restricted to allow for desired angular discrimination by moving the camera. Alternatively, infinite conjugated projection optics can be used for far-field viewing.

For phase and amplitude modulation, we define a super-pixel of three repetitions of high and low phase lines that are six LCoS pixels wide, using the double phase method [11,12]. The difference between the high and low phase creates phase gratings that diffract light into the ± 1 st and higher-order diffraction modes, enabling amplitude modulation [13]. The difference in phases controls the effective amplitude of the super-pixel. The mean of the high and low phases is the effective phase of the super-pixel. Thus, this super-pixel modulates the effective phase and amplitude of a single hogel entry. Given a desired amplitude A and phase θ_{sp} for the m th entry of a complex hogel, high and low phases θ_1 and θ_2 are found by using the following formula:

$$\theta_1[m] = \theta_{sp}[m] + \arcsin(A[m]/2), \quad (20)$$

$$\theta_2[m] = \theta_{sp}[m] - \arcsin(A[m]/2). \quad (21)$$

For full amplitude modulation combined with 2π phase modulation, the LCoS must be calibrated for a 0 to 2.5π phase stroke, as illustrated in Fig. 4.

For experimental verification, we design a pattern of three regions: an “S,” a “U,” and a background, bounded by colors in Fig. 5. Each region consists of a repeated 13×13 target hogel, enlarged on the right side of Fig. 5 and highlighted by the corresponding region’s color. Three color dots on each hogel represent three different angular perspectives that are experimentally captured, as shown in Fig. 6. Starting from the brown dot perspective, the “U” and the background is dark, and only the “S” is visible. When the camera is moved to the blue dot




Fig. 4. Example of Eqs. (20) and (21) phase grating modulation for three super-pixels. In this example, the amplitude of the first super-pixel equals the amplitude of the second, and the phase of the first equals the phase of the third.




Fig. 5. Target light field for hogel calculation. Three regions consisting of hogels with different angular responses are highlighted in red, yellow, and green. The top right is the enlarged elementary hogel for the “S” (red), the middle right is for the “U” (yellow), and the bottom right is for the background (green).




Fig. 6. Photographed angular views. Each row is color coded to match the angular view indicated on the angular patterns in Fig. 5. Rank 3 time-multiplexed images are constructed by averaging the results of three component frames.

perspective, all of the elements become visible. At the purple dot perspective, the “S” and the background disappear, leaving only the “U” visible. The right two columns in Fig. 6 are obtained with time-multiplexing. Compared to the rank 1 phase-only result, there is a marked improvement in the brown dot view for time-multiplexing and complex modulation. The contrast between the average grayscale value of the “S” pixels to the “U” pixels is 1.3 for rank 1 phase-only, 1.9 for rank 3 phase-only, and 5.4 for rank 3 complex modulation. However, we note that the complex modulation has a different resolution than the phase-only due to the use of super-pixels. In addition, as the phase of the super-pixel is averaged, complex modulation benefits from the simplicity of the test patterns.

To further confirm that the captured light field is equivalent to the calculated one, we photograph the Fourier planes for rank 1 and rank 3 “S” patterns and compare them to the calculated Fourier response of the optimized layer, as shown in Fig. 7. Note that the expected Fourier response, calculated here as a 2D FFT of the phase pattern generated by the algorithm, is not the same as the synthesized hogels generated by Eq. (9). When using the reflective phase-only LCoS, the Fourier response of a single “S” hogel will be overwhelmed by the Fourier response of the passive pixels. To address that, we repeat the “S” hogel to fill the entire display. This repetition




Fig. 7. Rank 1 and Rank 3 expected and measured the Fourier response of the repeated phase pattern for the “S” hogel from the last row of Fig. 2(a). The missing darkened pixels seen in the “S” shape from the calculated Fourier responses are seen in corresponding locations in the measured responses.

causes the Fourier response to narrow into peaks. The response is photographed by removing the camera lens from the conjugate setup (Fig. 3) and moving the sensor forward to the marked Fourier plane. As shown, the photographed results match the expected responses, most noticeably with darkened pixels in the S shape itself, validating that we are indeed observing the calculated light field. We have demonstrated that ADMM optimization of the WDF is capable of generating phase-only and complex hogels for displaying real-valued light fields. The simulation shows considerable improvement in PSNR under higher rank factorization, though the amount of improvement is highly dependent on the input target angular distribution. The simulation is experimentally verified on an SLM-based display prototype, where complex modulation is achieved by using the double phase method based “super-pixel.” Rank 3 time-multiplexing is prohibitive for the 60 Hz Pluto SLM, but may be an inexpensive trade-off on faster modulators. The described algorithm is directly applicable to coherent illumination displays for computer-generated or white-light camera recorded light fields.

[†]The indicated authors acknowledge equal contributions by sharing first authorship.

REFERENCES

- Z. Zhang and M. Levoy, in *IEEE International Conference on Computational Photography (ICCP)* (2009), pp. 1–10.
- J. W. Goodman, in *Fringe 2013* (Springer, 2014), pp. 3–15.
- J. Amako, H. Miura, and T. Sonehara, *Appl. Opt.* **34**, 3165 (1995).
- M. Lucente, “Diffraction-specific Fringe computation for electroholography,” Ph.D. thesis (Massachusetts Institute of Technology, 1994).
- R. Gerchberg and W. Saxton, *Optik* **35**, 237 (1972).
- G. Ye, S. Jolly, V. M. Bove, Jr., Q. Dai, R. Raskar, and G. Wetzstein, *ACM Trans. Graph.* **33**, 191 (2014).
- E. G. van Putten, I. M. Vellekoop, and A. P. Mosk, *Appl. Opt.* **47**, 2076 (2008).
- J. W. Goodman, *Statistical Optics* (Wiley, 2015).
- L. Shi, F.-C. Huang, W. Lopes, W. Matusik, and D. Luebke, *ACM Trans. Graph.* **36**, 236 (2017).
- S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, *Found. Trends Mach. Learn.* **3**, 1 (2011).
- C. K. Hsueh and A. A. Sawchuk, *Appl. Opt.* **17**, 3874 (1978).
- A. Maimone, A. Georgiou, and J. Kollin, *ACM Trans. Graph.* **36**, 1 (2017).
- O. Solgaard, F. S. A. Sandejas, and D. M. Bloom, *Opt. Lett.* **17**, 688 (1992).