J. Neural Eng. **6** (2009) 066004 (14pp) doi:10.1088/1741-2560/6/6/066004

Design and characteristics of holographic neural photo-stimulation systems

L Golan, I Reutsky, N Farah and S Shoham

Faculty of Biomedical Engineering, The Technion-I.I. T., Haifa, Israel

E-mail: sshoham@bm.technion.ac.il

Received 30 June 2009
Accepted for publication 29 September 2009
Published 16 October 2009
Online at stacks.iop.org/JNE/6/066004

Abstract
Computer-generated holography is an emerging technology for stimulation of neuronal populations with light patterns. A holographic photo-stimulation system may be designed as a powerful research tool or a compact neural interface medical device, such as an optical retinal prosthesis. We present here an overview of the main design issues including the choice of holographic device, field-of-view, resolution, physical size, generation of two- and three-dimensional patterns and their diffraction efficiency, choice of algorithms and computational effort. The performance and characteristics of a holographic pattern stimulation system with kHz frame rates are demonstrated using experimental recordings from isolated retinas.

(Some figures in this article are in colour only in the electronic version)

1. Introduction

The phenomenon of neuron excitation by light, first reported as early as the 1960s, has recently emerged as a technological alternative with significant potential benefits over electrical stimulation. Neural photo-stimulation is non-contact, enables higher spatial resolution, can be genetically targeted and used for inhibition of activity as well as excitation, and offers a solution to the problem of cross-talk generated by simultaneous electrical stimulation and recording. In the context of neural engineering, photo stimulation is becoming a key technology, as it opens up new possibilities for neural interface devices.

Light can excite neurons using a number of different mechanisms based on intrinsic tissue properties or requiring genetic or physiological manipulations. Early reports of photo-excitation [1, 2] have relied on strong illumination of stained neurons, presumably as a result of a photo-chemical effect. More recently, dye-free excitation using intense focused pulses of UV [3] and IR [4] lasers was reported, and was probably mediated by creating micro-holes in the membrane. A widely used method for experimental photo-stimulation is to photolyze caged neuro-transmitters [5], compounds which are inactive in their caged state and may be released by ultraviolet or violet light, leading to synaptic excitation. However, these compounds need to be continuously perfused into the neural tissue. In optogenetic stimulation, light-gated ion channels, such as channelrhodopsin-2 (ChR-2) [6], can be genetically introduced into mammalian neurons [7]; when illuminated, these channels enable the flow of cations through the membrane, thereby de-polarizing the cell and initiating an action potential. Different channels have been isolated, which are sensitive to different wavelengths [8], or enable neural inhibition [9]. In another development, short mid-infrared pulses have been demonstrated to intrinsically excite peripheral neurons [10], without causing long-term damage [11], apparently as a result of abrupt thermal transients created in neural tissue upon absorption of the light pulse [12]. The latter two techniques enable long-term stimulation and could potentially be used to create an optical cochlear implant [13] or an optical retinal prosthesis for individuals suffering from degenerative diseases of the outer retina [6, 14, 15].

Once neurons are rendered photo-sensitive, an appropriate system is required for the delivery of light to the tissue. Several methods such as xenon lamps [7], lasers [8], fiber optics [16] and micro-LED arrays [17] have been proposed. However, in most reports a large volume of neural tissue was non-selectively illuminated, activating large populations of neurons. In other reports, the system illumination field contained only a single cell. In both cases, the potential of photo-controlling neural populations with single-

 © 2009 IOP Publishing Ltd Printed in the UK

cell selectivity was not explored. For future neural-interface devices, such as neural prostheses, a spatially selective light-delivery technology is advantageous.

Several wide-field delivery systems with single-cell resolution have been proposed. Steering a focused laser beam using acousto-optic deflectors [18, 19] provides an accurate method; however, a minimal dwell time required at each location limits the size of neural population that can be accessed. Clearly, a technology for parallel delivery of light to many sites is desired. Digital micro-mirror devices (DMDs) are an option for projecting a light pattern onto neural tissue [14, 20]. However, DMDs manipulate the wavefront by blocking light where stimulation is undesired, and could be horribly inefficient for projecting the typically sparse wide-field stimulation patterns. Considering the high optical intensity required by the stimulation methods discussed earlier, this requires intense light sources that will be especially impractical for wearable, low-power, optical neural prosthetic devices.

The ultimate combination of parallel access with high efficiency can be found in phase modulation of the wavefront. Phase-only computer-generated holograms (CGHs) are optical devices that enable the projection of arbitrary patterns with high power efficiency [21]. With phase-only CGHs, the optical power is divided between the different focal points of the projected pattern; so theoretically, no light is lost in the modulation process. Dynamic CGHs may be realized by spatial light modulators (SLMs) which enable rapidly changing phase-only modulation in a reliable, inexpensive and potentially miniature device. Phase-only CGHs displayed on SLMs have already been used for wavefront correction [22], optical cross-connects [23], laser fabrication [24] and video display [25]. In a more biologically related context, CGHs have also been widely used to create dynamic arrays of optical traps or 'laser tweezers' [26], which are often used in biological research. Recently, CGHs have also been applied to neural photo-stimulation, activating single cells in cortical slices by the uncaging of glutamate [27]. In that work, light efficiencies of 50% and more were achieved. A similar system may be applied for artificial stimulation of the retina, which is a two-dimensional optically accessible neural tissue. Moreover, the inherent ability of holography to generate three-dimensional light patterns enables the extension of this technique to 3D tissues, where neurons are arranged at different depths. For highly scattering tissues, such as the brain, 3D holography can be combined with two-photon absorption, to achieve deeper penetration. This kind of system can also be used to combine deep two-photon imaging and stimulation of scattering tissue [28, 29]. Yet another advantage of holography is the ability for dynamic correction of changes and distortions in the optical system [30].

This manuscript describes detailed design considerations for a CGH-based photo-stimulation system, which we have encountered during the development of a prototype optical retina interface. Earlier publications in this field [27, 28] are research reports focusing on specific systems, and did not capture broad design considerations. Some of the issues discussed here are unique to the photo-stimulation application

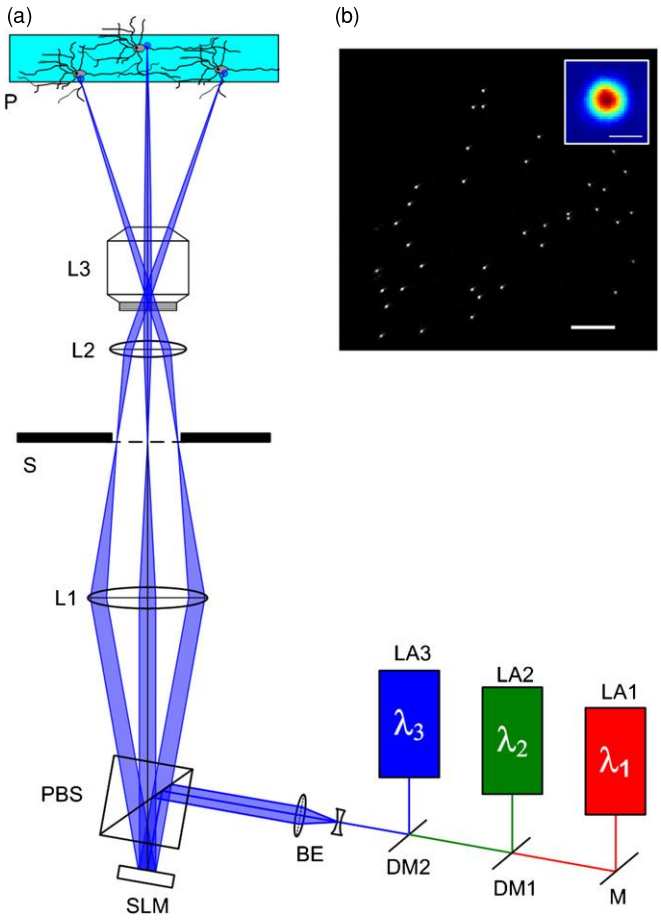

Figure 1. (a) A schematic of the holographic system. Its components: LA1–3, laser sources; M, mirror; DM1–2, dichroic mirrors; BE, beam expander; SLM, spatial light modulator; PBS, polarizing beam splitter; L1–3, lenses, of which L3 is a microscope objective; S, square slit that blocks high orders and defines FOV limits; P, neural tissue preparation. (b) An example pattern generated by the holographic system; scale bar is 100 μm. Inset, an intensity map of a single spot; scale bar is 5 μm.

and must be overcome for the successful integration of a holographic photo-stimulation system, including field-of-view limits, power efficiency and uniformity, spot shape and computation time. In section 2, we discuss the optical design of a holographic system. In section 3, we present the general considerations in hologram design, and discuss the limitations it imposes on system performance. In section 4, we review the hologram calculation algorithm, as well as the display of hologram sequences. In section 5, we present physiological results obtained with a holographic photo-stimulation system.

2. Optical design and characteristics

The layout of a multi-wavelength holographic photo-stimulation system is shown in figure 1(a). The system is essentially a Fourier holographic system [31], where the output pattern is the Fourier transformation of the modulated wavefront reflected off the SLM:

$$E_{\text{pattern}}(u, v) = \mathcal{F}\{E_{\text{hologram}}(x, y)\} \quad (1)$$

where (x, y) is the position on the hologram plane (the SLM surface), (u, v) is the so-called reconstruction plane, where the desired pattern is formed and $\mathcal{F}\{\cdot\}$ is a Fourier transform.

The Fourier lens (L3) is usually a microscope objective which ensures that the projected pattern is nearly diffraction limited. A telescope (lenses L1 and L2) is used to match the sizes of the modulated field and the objective input pupil. Typically, the SLM cannot display perfect CGHs, resulting in the creation of undesired additional diffraction orders. These orders are blocked by placing an appropriate stop at the intermediate image plane. The SLM normal is slightly tilted with respect to the optical axis in order to align the first-order beam with the optical axis while blocking the zero-order beam at the stop. The entire system can be integrated with a commercial microscope system for research purposes, or the same elements can be miniaturized and fit into a wearable neural interface device.

In the following section, we describe the primary considerations in the design of the optical system.

2.1. Choice of spatial light modulator type

The SLM may be any phase-modulating device based, among other options, on liquid crystals, MEMS or deformable mirrors. Among the available technologies, the liquid crystal on silicon (LCoS) is the most mature and readily available device, providing high pixel counts, almost no inter-pixel dead space and simple interfaces, at a reasonable cost. We will focus our discussion on this technology, keeping in mind that other technologies might prove themselves as worthy candidates in the near future. Among the critical parameters of LCoS devices for our application are diffraction efficiency and switching rate. High diffraction efficiency requires multi-level phase modulation [32], which is currently provided by nematic LC devices. However, the physics of the nematic LCs limits their switching rate, which is currently up to 100 frames per second. On the other hand, devices based on ferroelectric LCs provide very fast switching (up to kHz refresh rates), but are generally limited to binary modulation. One effect of binary modulation is an inherent upper bound on diffraction efficiency of approximately 40%. Another 40% of the power is directed to a replicate image lying on the opposite side of the optical axis, thus cutting the system addressable field in half.

For a tabletop research application, low diffraction efficiency can often be compensated by higher laser power. Fast systems with kHz refresh rates do not limit the precision of neuronal timing and also enable time multiplexing and time averaging, methods that will be discussed later on. Therefore, a fast ferroelectric SLM should be considered when such systems are desired.

When diffraction efficiency is critical, as is the case with a two-photon system, the use of multi-level nematic LC is more appropriate. Nematic devices with display rates of 60–85 Hz are available, and next generations of LCoS devices are expected to combine the benefits of both speed and diffraction efficiency.

Diffraction efficiency is also affected by fill factor—the fraction of the SLM unit cell that is occupied by the actual modulating pixel. With LCoS, electronic circuitry can be moved below the LC layer, resulting in fill factors very close to 1. However, additional parameters influence the diffraction efficiency and must be taken into account—SLM reflectivity and absorption, wavelength and temperature dependence, as well as inaccurate switching angle (for ferroelectric SLMs). The combined effect of these factors may result in an order of magnitude decrease in diffraction efficiency for some SLMs.

Another critical parameter of SLM devices is their phase modulation depth. Mathematically, a full 2π modulation is required for utilizing the benefits of phase-only modulation. Since retardation is wavelength dependent, the 2π requirement is harder for long wavelengths, such as mid-wave infrared, and easier in the visible range. As the wavelength increases, thicker LC layers are required for achieving 2π retardation, and these devices tend to switch more slowly.

2.2. Resolution and field-of-view

We now describe the spatial properties of the stimulation field. We will use the coordinates (x, y) to describe lateral position on the SLM plane, (u, v) to describe lateral position on the stimulation plane and z to describe distance along the optical axis.

The resolution of the system is determined by the objective numerical aperture (NA). The spot size can be characterized by its Airy radius:

$$\Delta u = 0.61\frac{\lambda}{\mathrm{NA}} \tag{2}$$

where λ is the stimulation wavelength. If the SLM image at the back aperture plane is smaller than the aperture, the NA is not fully utilized and the spot size will increase. On the other hand, if the SLM image is larger than the aperture, light is lost. The size of the SLM image is determined by the magnification m of the telescope formed by lenses L1 and L2 in figure 1:

$$m = \frac{f_{\mathrm{L2}}}{f_{\mathrm{L1}}}. \tag{3}$$

To match the sizes, m should be chosen to satisfy

$$mL = 2f_O\,\mathrm{NA} \tag{4}$$

where f_O is the objective focal length and L is the lateral dimension of the SLM. We note that if the SLM has a different number of pixels in the horizontal and vertical directions, L should be its short axis, and the SLM long axis is truncated by the aperture and does not contribute.

The lateral size of the area which can be accessed by the SLM is a critical factor. We adopt the term field-of-view (FOV), even though this is not an imaging system. The FOV is determined primarily by the SLM pixel pitch; smaller pixels enable larger diversion of the wavefront. The maximal lateral shift is

$$u_{\max} = \frac{\lambda}{2md}f_O \tag{5}$$

where d is the SLM pitch. This defines the FOV limits as $[-u_{\max}, u_{\max}]$ in both directions. Therefore, it is useful to define the following relative lateral coordinates:

$$\tilde{u} = \frac{u}{u_{\max}}; \qquad \tilde{v} = \frac{v}{u_{\max}}. \tag{6}$$

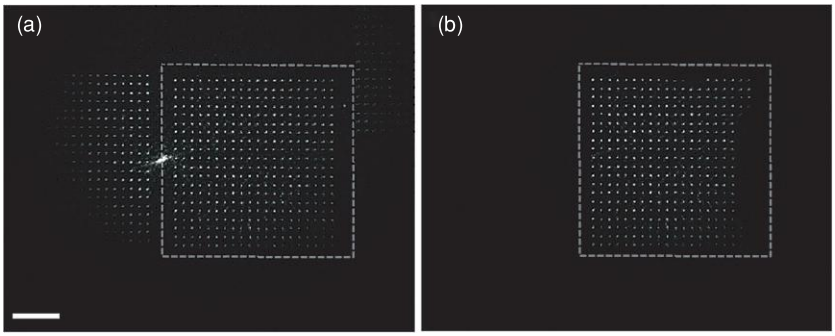

Figure 2. Limiting the FOV with a slit. (a) An unblocked output pattern of a 20 × 20 spot grid created by a binary SLM. The zero-order, replicate pattern and higher orders are visible. The dashed gray line indicates the theoretical FOV limits as defined by equation (7). Scale bar is 200 μm. (b) Inserting a square slit at the intermediate image plane blocks the undesired orders. The slit used was slightly smaller than the theoretical calculation, in order to block the zero-order 'halo' effect.

In theory, a multi-level grating may be designed to suppress the zero diffraction order, enabling use of both positive and negative angles of deflection. However, it is practically impossible to prevent some of the light, reflecting off optical surfaces or interpixel dead space, from being focused on an intense point on the optical axis. Placing a stop at the intermediate image plane at the focus of lens L1 can be used to filter the zero order. This stop may be used to block only a small region around the zero order, resulting in an irregularly shaped FOV. A different choice would be to block the entire negative half-plane, including the zero order. For a binary SLM, this also solves the problem of the duplicate pattern, while cutting the FOV in half. In this case, the resulting FOV is a rectangular with a 2:1 aspect ratio. The stop shape determines the FOV aspect ratio. A square FOV is typically comfortable, so limiting both $\tilde{u}$and $\tilde{v}$, we get the following FOV definition:

$$0 < \tilde{u} < 1; \qquad -\tfrac{1}{2} < \tilde{v} < \tfrac{1}{2}. \tag{7}$$

This particular selection of the FOV (rather than $0 < \tilde{v} < 1$) increases the average diffraction efficiency (discussed in section 3.1). The effect of this FOV stop is illustrated in figure 2.

Having defined the FOV as above, we can combine equations (2) through (5) to obtain the effective number of accessible points in the FOV:

$$\left(\frac{u_{\max}}{\Delta u}\right)^2 = \left(0.61\frac{L}{d}\right)^2 = 0.37N^2 \tag{8}$$

where N is the number of pixels along the SLM side (typically 500–2000). Note that the number of accessible points does not depend on objective focal length, because both FOV and spot size are affected by it. A convenient way to manage the trade-off between FOV and resolution is by switching between different microscope objective lenses. A higher power objective will decrease the FOV and the theoretical spot size by the same factor. Different objectives have different apertures, so the designer should choose the magnification m according to the smallest aperture.

2.3. Aperture filling

Another important consideration is the filling of the SLM aperture by the illuminating laser beam. Underfilling will ensure minimal power loss at the objective input aperture, although the objective NA will not be used to its maximum. This trade-off must be considered during system design. If we assume that the de-magnification of the telescope is chosen according to equation (4), the SLM is not a limiting factor, and the problem reduces to that of choosing the truncation ratio of the expanded Gaussian beam by the objective aperture [33]. A useful definition in this case is the truncation ratio, which is the ratio between the Gaussian beam $1/e^2$ diameter and the aperture diameter:

$$T_{\text{ratio}} = \frac{D_{\text{beam}}}{D_{\text{apperture}}}. \tag{9}$$

Decreasing the truncation ratio improves light efficiency, but increases the spot size. For photo-stimulation, light efficiency is typically more important than spot size, because the resolution of photo-stimulation is often affected by other factors such as tissue scattering, neurotransmitter diffusion, etc. Therefore, underfilling with truncation ratios of 0.5–0.7 is recommended. Making this choice, the objective effective NA is reduced, spot size increases and the constant in equation (8) should be recalculated according to the specific truncation ratio.

If the SLM image is smaller than the objective back aperture, the effective aperture becomes a multiplication of the circular objective pupil by the rectangular SLM image. One can numerically calculate the effect of this combined pupil on power and PSF in order to select the desired incident beam diameter.

2.4. Integration into commercial microscopes

When the system is based on a commercial infinity-corrected microscope, it is desirable to incorporate the photo-stimulation capability with existing microscope functions such as transmitted light imaging, fluorescence imaging and control of sample position. One possible approach is to use the fluorescence port of the microscope to insert the stimulation light. The light enters the microscope between

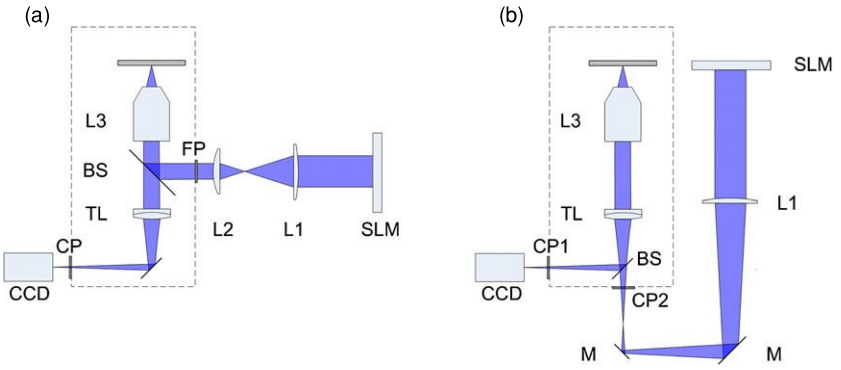

Figure 3. A comparison of the methods for microscope integration. (a) Light enters through the fluorescence port. (b) Light enters through the auxiliary camera port. Components: CCD, charge-coupled device camera; TL, tube lens; BS, beam splitter; L1–3, lenses; M, mirrors; CP, CCD port; FP, fluorescence port. The dashed line encloses the original microscope system.

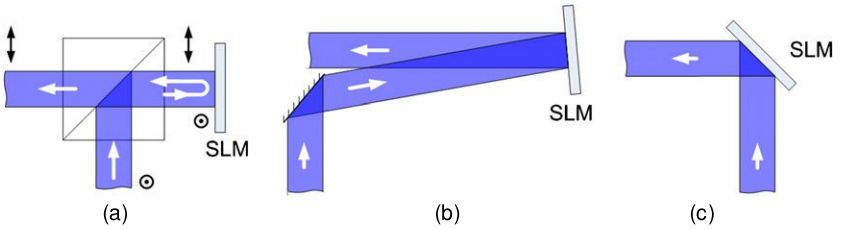

Figure 4. Different options for SLM integration: (a) polarizing beam splitter (quivers indicate polarization direction); (b) near-perpendicular incidence; (c) oblique incidence.

the objective and the so-called tube lens as described in figure 3(a). Many microscopes allow the stacking of several dichroic mirrors in this space, so fluorescence imaging may be performed in parallel to photostimulation, provided that the filter passbands are chosen correctly. This approach requires using a fluorescence cube that holds the combining mirror (a dichroic or partially reflecting mirror) and the two lenses of the imaging telescope that are positioned outside the microscope at correct distances, to create the image of the SLM on the back aperture of the objective lens.

An alternative approach is demonstrated in figure 3(b). Here, we use the microscope tube lens as the second lens of the imaging telescope. This is usually an achromatic, well-corrected lens. An auxiliary camera port of the microscope is used to insert the stimulation light. Most modern microscopes are equipped with a rotating beam splitter turret that enables the choice of splitting ratio between the auxiliary port and the ocular. Switching between different splitting ratios enables a comfortable choice between imaging, stimulation or simultaneous imaging and stimulation (at the expense of reduced light efficiency). This approach usually results in longer optical path lengths, as the tube lens has a long focal length (in the range of 165–200 mm for commercial microscopes).

2.5. Space-saving strategies

The small size of the optical system is a key factor when applying holographic projection technology in miniature wearable or implantable photo-stimulation devices. It may also be important for desktop research systems. One critical determinant of system size is the separation of the SLM incident and reflected beams. Different options are schematically described in figure 4. For binary ferroelectric LC cells, the phase modulation is achieved by a 90° rotation of the incident polarization, so a polarizing beam splitter (PBS) naturally solves the problem. This is also the case for twisted nematic (TN) LC cells whose operation involves a rotation of the incoming polarization. However, the more modern vertically aligned nematic (VAN) cells do not rotate the incoming polarization, so other solutions must be considered. One option is hitting the SLM with a near-perpendicular incidence. In this case, the modulation is almost not affected, but a minimum distance must be taken to enable separation of the two beams. This distance depends on the beam radius and the actual incidence angle. On the other hand, a large-angle oblique incidence introduces a minor distortion [34] but saves on optical path length.

Another consideration for space saving is the configuration of the de-magnifying telescope. This telescope is composed of two lenses with focal lengths f_{L1} and f_{L2}. In the 4-f configuration, the SLM is placed at the focal plane of the first lens, and its image is formed at the focal plane of the second one. The resulting total length is $2(f_{L1}+f_{L2})$. Moving the SLM closer to the first lens not only reduces the total telescope length, but also decreases the numerical aperture of the telescope [35]. This may cause vignetting, so the optical design must take into consideration the lens' sizes, the maximal diffraction angle and the incoming beam width.

2.6. An illustrative design

The concepts discussed above are general for any holographic photo-stimulation system. Nevertheless, it will be useful to demonstrate different concepts with a concrete example. We therefore present the parameters of an illustrative system, which is integrated with a commercial inverted microscope (TE-2000U, Nikon, Japan). Our system contains three lasers: red, green and blue. All of them are diode-pumped solid-state lasers, emitting at 473 nm, 532 nm and 671 nm (VA-I-N-473 & VA-I-N-532, Viasho, China and RL671T-300, Shanghai Laser, China). The beams are combined using long-pass dichroic mirrors (DMLP505 & DMLP567, Thorlabs, USA). They are expanded by a 1:4 beam expander to obtain 1.5–4 mm diameter beams (depending on the specific laser used). The SLM is a fast binary ferroelectric device with 1280 × 1024 pixels (SXGA-R3, ForthDD, Scotland) and a pitch of 13.6 μm, although only the central 512 × 512 pixels have been used for hologram display. The SLM modulates light over a wide wavelength range, although diffraction efficiency slightly varies with wavelength. Manufacturer declares a 0.92:1:0.86 relation for the specified R:G:B wavelengths. The SLM image is projected onto the objective back aperture by a 2:1 de-magnifying telescope comprising a 400 mm plano-convex lens and the microscope tube lens (f = 200 mm). Using a CFI Achromat DL 10× objective (Nikon, Japan), the system achieves a 700 μm FOV with a 3 μm FWHM spot size on the sample. The trade-off between FOV and resolution may be modified by switching to a CFI Plan Fluor 4× objective or a CFI Plan Fluor 20× objective. A demonstrative sparse pattern is shown in figure 1(b). The demonstrations shown in this work were obtained with our illustrative system at a wavelength of 532 nm, unless otherwise stated.

3. Hologram design and characteristics

The hologram displayed on the SLM may be viewed as one of the optical components of the system. However, its design is dynamic and changes each frame. The SLM limitations, such as finite pixel size and phase quantization, have an impact on the resulting pattern. The hologram design process can minimize these effects, as well as compensate for other impairments in the system, including aberrations. In this section, we review the effects of hologram and SLM characteristics on system performance.

3.1. Two-dimensional light patterns

In Fourier holography, any reconstructed pattern may be written as a sum of diffraction-limited spots. The hologram generating such a pattern can be regarded as a superposition of gratings, with each one of them diverting a fraction of the incident beam to create a beamlet with a defined direction in space. The Fourier lens then focuses each beamlet to a diffraction-limited spot on the sample. This description is especially useful for photo-stimulation patterns, which are discrete in nature and usually sparse. Any discrete pattern defined on an $N \times N$ grid may be reconstructed by the following complex amplitude hologram:

$$t(x, y) = \sum_{l=1}^{N} \sum_{n=1}^{N} e_{ln} \exp\left(2\pi \mathrm{i} \frac{lux + nvy}{\lambda f_O}\right) \tag{10}$$

where e_{ln} is the amplitude at the (l, n) point. In particular, a sparse pattern comprising only $M \ll N^2$ points at the positions (u_k, v_k), $k = 1, 2, \ldots, M$, and with amplitudes e_k, may be generated by the following complex amplitude hologram:

$$t(x, y) = \sum_{k=1}^{M} e_k \exp\left(2\pi \mathrm{i} \frac{u_k x + v_k y}{\lambda f_O}\right). \tag{11}$$

This transfer function is not necessarily phase only. However, by setting correct phase angles for e_k values, it is possible to find an approximately phase-only function of the above form. Finding the best phase angles requires calculations and is discussed in section 4. Assuming this has been done, equation (11) is still an ideal continuous function that can only be approximated by the SLM. This approximation has several well-known implications, which we will review briefly.

The quantization of phase levels reduces the diffraction efficiency and diverts light to higher, undesired orders [32]. This effect quickly becomes negligible as the number of quantization levels exceeds 16. However, for binary SLMs the maximum theoretical diffraction efficiency is 40.5%, a factor which must be accounted for when designing the system optical power budget.

The spatial discretization by the device pixels further results in a spatially varying diffraction efficiency [36, 37], effectively multiplying the output intensity by a sinc-shaped envelope, described by

$$\eta(u, v) = \mathrm{sinc}^2\left(\frac{\tilde{u}}{2}\right) \mathrm{sinc}^2\left(\frac{\tilde{v}}{2}\right) \tag{12}$$

where we have defined $\mathrm{sinc}(x) = \sin(\pi x)/\pi x$ and used the relative coordinates $(\tilde{u}, \tilde{v})$ defined in equation (6). This effect leads to non-uniform optical power density over the FOV and reduces the overall average diffraction efficiency. For example, a system with no phase quantization and an FOV defined as in equation (7) will have an average diffraction efficiency of 72% over the entire FOV. In a well-designed system, this non-uniformity must be corrected, because a constant power density is required over the entire field. A simple solution is to assign higher amplitudes (e_k values) to off-axis spots during hologram design. This results in a uniform intensity, at the cost of further reducing the average diffraction efficiency (67% in the above example). The envelope effect and its correction are demonstrated in figure 5.

The fill factor of the SLM also influences diffraction efficiency. Unmodulated light is focused to the zero order, so diffraction efficiency drops by a factor of the squared fill factor. Another effect of the fill factor is a smaller pixel, resulting in broadening of the sinc envelope, so uniformity is slightly improved. However, for modern LCoS devices the fill factor is usually larger than 90% and the effect is minor.

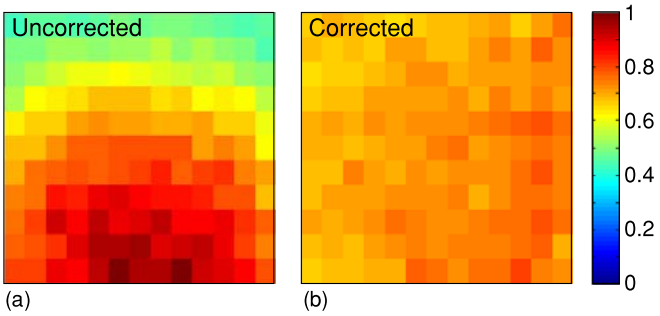

Figure 5. (a) A two-dimensional plot illustrating the non-uniformity resulting from the sinc envelope effect. Color scale indicates average intensity, in normalized units. (b) Non-uniformity corrected by applying an inverse correction factor.

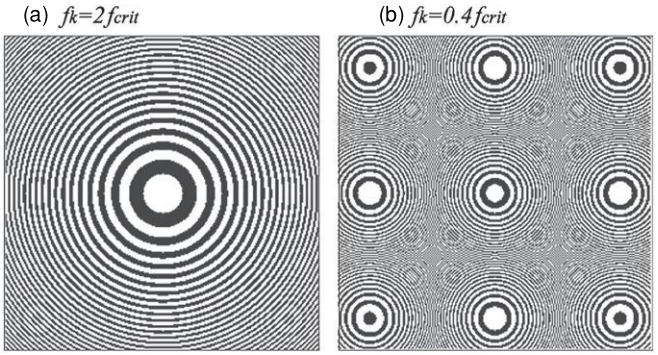

Figure 6. Phase map of a sampled binary Fresnel lens, illustrating aliasing effect. Black and white indicate phases of 0 and π, respectively. (a) When focal length is larger than critical length. (b) When focal length is shorter than critical length.

3.2. Three-dimensional light patterns

To create a three-dimensional spot pattern, a parabolic phase is added to the hologram expression:

$$t(x, y) = \sum_{k=1}^{M} e_k \exp\left[2\pi \mathrm{i}\frac{u_k x + v_k y}{\lambda(f_O + z_k)}\right] \exp\left(-2\pi \mathrm{i}\frac{x^2 + y^2}{2\lambda f_k}\right) \tag{13}$$

where z_k is the axial position of the stimulation point (with zero corresponding to the objective focal plane). This expression is a superposition of Fresnel lenses, each creating an individual beamlet with a predefined direction and curvature, resulting in an axial shift of the beamlet focus position out of the objective focal plane. f_k is the Fresnel lens focal length, calculated by the geometrical imaging formula:

$$f_k = -\frac{f_O(f_O + z_k)}{z_k} \tag{14}$$

where the minus sign indicates that a diverging Fresnel lens is required for a positive axial displacement. As the focal length decreases, the Fresnel lens contains higher spatial frequencies. The SLM is inherently limited by its ability to represent spatial frequencies, due to its finite pixel pitch. This defines a critical limit for the focal length of the Fresnel lens that can be displayed [38] as

$$f_{\mathrm{crit}} = \frac{Nd^2}{\lambda} \tag{15}$$

where N is again the number of pixels along the SLM side, as defined in section 2.2. The image of the SLM on the objective back aperture is smaller by a factor equal to the telescope magnification m, which results in a reduced effective critical focal length:

$$f'_{\mathrm{crit}} = \frac{N(md)^2}{\lambda}. \tag{16}$$

Trying to display a higher power lens on the SLM will result in aliasing, which actually creates shifted replications of the same lens (see figure 6). This diverts a large portion of the power to secondary focal spots, and deteriorates the diffraction efficiency. Therefore, it is convenient to define a dimensionless lens power:

$$\zeta_k = \frac{f'_{\mathrm{crit}}}{f_k}, \tag{17}$$

which should be smaller than 1 to avoid aliasing. The value of ζ could also be thought of as a dimensionless axial coordinate, because for small axial displacements, it changes linearly as a function of z_k:

$$\zeta_k \approx -\frac{N(md)^2}{\lambda f_O^2} z_k \qquad \text{for} \quad z_k \ll f_O. \tag{18}$$

However, for large axial displacements there is deviation from linearity.

3.3. Diffraction efficiency of 3D light patterns

Just as is the case with 2D stimulation patterns, diffraction efficiency of 3D patterns is coordinate dependent. From equation (13) we see that combined displacements in the axial and lateral directions require higher spatial frequencies in the corresponding hologram. The approximation of these patterns by a discrete SLM becomes less accurate as the stimulation point is axially shifted from the focal plane, or laterally shifted from the optical axis. We shall now present an exact expression for this position-dependent diffraction efficiency, which is useful for defining the system's accessible volume. An analytical expression for the intensity distribution of a digital hologram displayed on a pixel array and illuminated by a spherical wavefront has been derived by Haupt *et al* [39]. In our case, the wavefront is first modulated by the hologram and then by the spherical phase, which is an equivalent problem. We begin with equation (35) from [39], and after some notation conversions, we get the following expression for the total power contained within the main order of a reconstructed delta function:

$$E(u_k, v_k, z_k) = \left(\frac{f_O}{z_k}\right)^2 \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \left|T\left[\frac{u_k - u}{\lambda(f_O + z_k)}, \frac{v_k - v}{\lambda(f_O + z_k)}\right]\right|^2 \times \left|a\left(\frac{uf_O}{z_k}, \frac{vf_O}{z_k}\right)\right|^2 \mathrm{d}u\,\mathrm{d}v \tag{19}$$

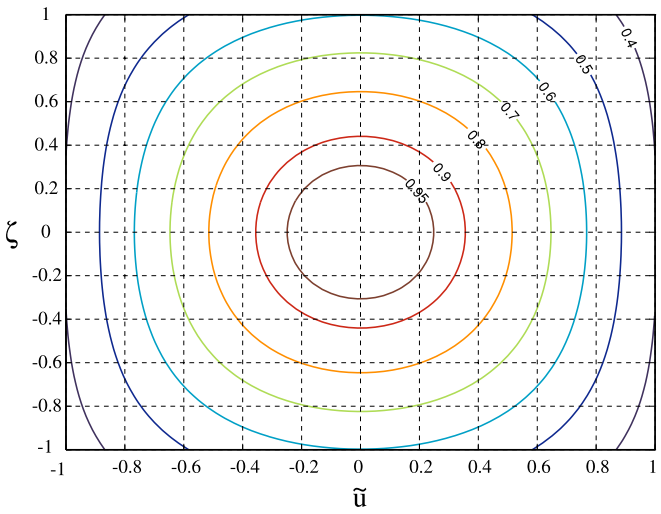

Figure 7. The diffraction efficiency of a single spot as a function of its lateral and axial locations, given in dimensionless coordinates (see the text for definitions). The distribution applies for continuous phase modulation and should be scaled accordingly for quantized phase SLMs.

where $a(x,\ y)$ is the hologram aperture function, $t(x,\ y)$ is the pixel aperture function and $T(u, v)$ is its Fraunhoffer diffraction pattern:

$$T\left[\frac{u}{\lambda(f_O + z_k)}, \frac{v}{\lambda(f_O + z_k)}\right] = \mathcal{F}\{t(x, y)\}. \tag{20}$$

In other words, the power directed toward a point in space (u_k, v_k, z_k) may be calculated by convolving the single-pixel diffraction pattern with the SLM aperture function. The diffraction efficiency may be calculated by dividing the main-order power by the incident power, assuming unit-amplitude illumination.

It is useful to calculate the diffraction efficiency for a practical case of a rectangular SLM with rectangular pixels and no inter-pixel dead space. The Fraunhoffer diffraction pattern T is then a two-dimensional sinc function. We also switch to the relative coordinates $(\tilde{u}, \tilde{v}, \zeta)$ defined in equations (6) and (17), and after some manipulations get the following expression:

$$\eta(\tilde{u}_k, \tilde{v}_k, \zeta_k) = \frac{1}{4\zeta^2}\int_{(\tilde{v}_k-\zeta_k)}^{(\tilde{v}_k+\zeta_k)} \mathrm{d}v' \int_{(\tilde{u}_k-\zeta_k)}^{(\tilde{u}_k+\zeta_k)} \mathrm{d}u' \ \mathrm{sinc}^2\left(\frac{u'}{2}\right) \mathrm{sinc}^2\left(\frac{v'}{2}\right). \tag{21}$$

One can see that when $\zeta \to 0$, the integral reduces to the two-dimensional case of equation (12). The values of the diffraction efficiency on the $(\tilde{u}, \zeta)$ plane (one axial coordinate and one lateral coordinate) are plotted in figure 7. As an example, using the parameters of our illustrative design (presented in section 2.6) with a 20× objective, and setting a threshold of 25% diffraction efficiency (including the 40.5% factor of binary modulation) yield an accessible volume 300 μm wide and 2.2 mm deep (in comparison to the 10 mm focal length of the objective). As in the case of two dimensions, the non-uniform diffraction efficiency can be compensated by applying an inverse amplitude correction during hologram design.

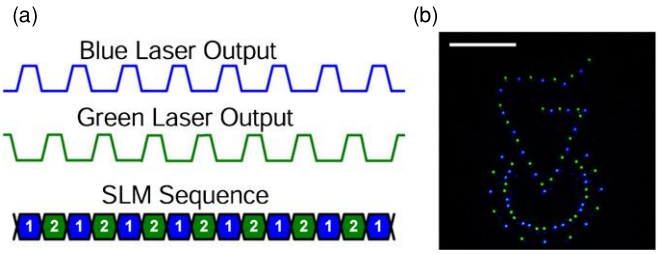

Figure 8. (a) An illustration of the proposed technique for projecting dual wavelength patterns. The lasers are modulated with a 50% duty cycle, synchronized with the SLM frames, while the SLM projects two alternating patterns with a total period of 1 ms. (b) An example of a dual wavelength pattern (our institute symbol). Scale bar is 500 μm.

When designing a photo-stimulation system for a three-dimensional volume, one should keep in mind that the beam is no longer collimated after reflecting off an SLM that displays a Fresnel lens. This fact changes the beam diameter at the different system apertures. If these apertures are not wide enough, the Gaussian beam truncation will not be negligible and the result will be loss of power. This loss can be numerically calculated for different axial shifts, in order to account for this effect and gain a more accurate prediction of the overall diffraction efficiency.

3.4. Temporal characteristics

Another characteristic of the SLM which affects system performance is its temporal behavior. The SLM response time determines how long it takes for the LC molecules to re-orient between adjacent frames. During this transition time, the phase modulation is not well defined, and most power is directed toward the blocked zero order. However, some of the power may be randomly dispersed over the FOV. If this power density is enough to cause undesirable stimulation effects, the light must be blocked during this time, either by modulating the laser or by an external modulator. The same applies to periodical times without modulation, which are used by some of the SLM electronics for internal 'house-keeping'.

Synchronizing the laser output with the SLM frame clock may also be useful in other situations. For slow SLMs, the basic frame time might be too long for a single stimulation pulse. In this case, illuminating the SLM only a fraction of the frame time solves the problem. For fast SLMs, the high refresh rate may be used to multiplex two different patterns, for example two different wavelengths. Figure 8 illustrates a dual-wavelength pattern projected by alternating two different frames on a ferroelectric SLM, while modulating two lasers at 50% duty cycle with a total period of 1 ms. This kind of dual-wavelength patterns could be used to control two different types of light-gated ion channels independently. In a research system, if the SLM frame clock is not externally available, it may be simply extracted by connecting a PC monitor in parallel with the SLM, converting its visual output to an electrical signal using a photo-diode, and using this signal as an input

to an acquisition board that generates the laser modulation signals.

4. Hologram calculation and display

In the general case, holograms cannot be calculated by direct substitution into equation (13). Trying to do so yields a complex-amplitude function that cannot be displayed on a phase-only device. Solving this problem requires a more sophisticated algorithm. In this section, we review some of the available algorithms, keeping in mind that calculation time is a critical factor for real-time systems. In addition, the speckle problem, which is closely related to the hologram calculation process, will be presented. We also discuss several issues concerning the implementation of hologram calculation and display.

4.1. Choice of algorithm

The hologram synthesis problem is a well-known and studied subject [21]. Our focus is on finding a phase-only hologram while maximizing the diffraction efficiency. There are many degrees of freedom in this problem which do not affect the resulting intensity pattern, such as the phase of the resulting pattern (phase freedom), and the amplitude outside the FOV (amplitude freedom). Many algorithms and variations have been proposed, with different levels of complexity. However, a neural prosthetic device that transforms real-time sensory data to spike trains in neural populations is expected to require at least tens of holograms per second, so a simple algorithm is a primary demand. This rules out many iterative, slowly converging algorithms that have been proposed, such as direct search [40]. On the other hand, the sparse and pseudo-random nature of photo-stimulation patterns has a positive impact on algorithm performance, and this should be taken into consideration.

One of the simplest algorithms is the superposition method [41]. In this method, one calculates the sum in equation (11) and retains only the phase angle of the result. A slightly better variant, termed SR, is obtained if a random phase is assigned to each point, i.e. the e_k values are assigned random phase angles. This method requires only substitution operations, which can be done rapidly by modern hardware as well as PCs.

A better algorithm which may also be considered is the iterative Gerchberg–Saxton (GS) algorithm [42], which is widely used in holography as well as other areas. The GS algorithm usually converges to an approximate solution which gives good results. Actually, the first iteration of the GS algorithm is equivalent to the SR algorithm, but results improve after additional iterations. The GS algorithm is computationally demanding since it requires calculations of forward and backward Fourier transformations, until reconstruction error is small enough. Fast implementations of this procedure may be realized using field programmable gate arrays (FPGAs) [43] or graphical processing units (GPUs) [44].

The choice of a specific algorithm depends heavily on the characteristics of the pattern to be reconstructed. For a wide-field system, with many neurons contained in its FOV, no more than hundreds of individual neurons will be stimulated per frame, out of 100 000 or more accessible points. Therefore, our typical pattern will be sparse. This is similar to the case encountered in holographic optical trapping. Curtis *et al* [45] have explored the performance of GS and SR algorithms on sparse patterns and found that while high-symmetry patterns such as regular lattices benefit from the GS algorithm, random-like patterns do not. Actually, for random patterns, the SR algorithm achieves close to 80% diffraction efficiency, and any improvement achieved by additional GS iterations is negligible. In a practical system, neurons rarely need to be stimulated in symmetrical structures, so the SR algorithm is clearly advantageous.

However, even when neurons are sparsely and randomly distributed over the FOV, it might be useful to create contiguous patches composed of many adjacent spots. This enables precise control over the size and shape of the stimulation spots. In other cases, one might be trying to generate a 'moving-bar' pattern sequence that allows spatial probing in physiological experiments. When the spots of a pattern are arranged in clusters, the randomness assumption must be re-tested. In what follows, we present the results of tests designed to address this issue.

The performance of algorithms is characterized not only by the average diffraction efficiency, but also by the variability of diffraction efficiency across the pattern. This variability, characterized by the standard deviation of the different point intensities, results in non-uniform power density and is to be minimized. In figure 9, the performance of the GS algorithm is presented for a high-symmetry pattern of 5 × 5 spots, a random pattern of 25 spots and a random pattern of 25 rectangular 3 × 3 patches. The results for the symmetrical and random cases coincide with the results described by Curtis *et al*, while the behavior of the randomly placed patches lies somewhere in between these cases. One should keep in mind that the performance of the SR algorithm is equivalent to the performance of the first GS iteration, so that the vertical span of each curve represents the performance gained by switching from SR to the more complex GS. This gain depends on the average patch size and number of patches, so one should consider the trade-off between power efficiency and computational complexity, depending on the characteristic pattern he is dealing with.

A notable variation of the GS algorithm that is especially suited for our application is the GSW, proposed by Dileonardo *et al* [46]. GSW provides better uniformity than the standard GS, with a minor compromise on diffraction efficiency. The improvement in uniformity is obtained by correcting the intensities of the target pattern for every iteration, according to the ratio between the desired intensities and the last result.

4.2. Spatially extended illumination patterns and the speckle problem

As we have discussed in section 3.1, some illumination patterns simply comprise a sparse set of diffraction-limited spots. However, holographic photo-stimulation systems are

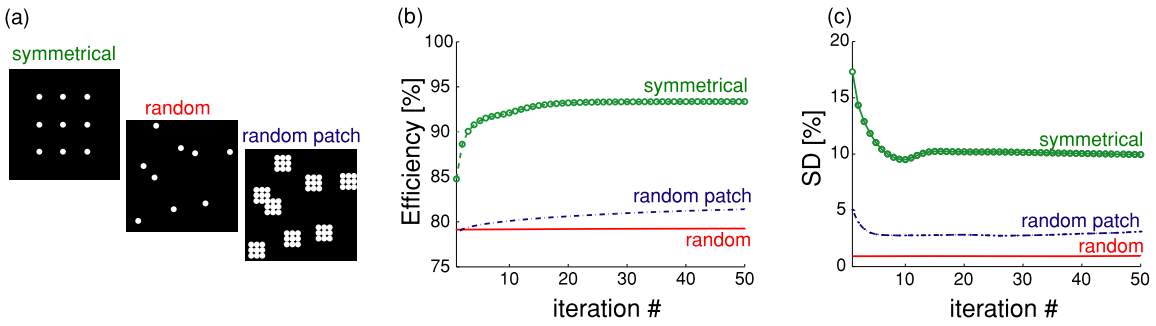

Figure 9. Comparison of the GS algorithm performance (with continuous phase modulation) for three different patterns. (a) A schematic description of the tested patterns. (b) Average diffraction efficiency. (c) Standard deviation of spot intensity.

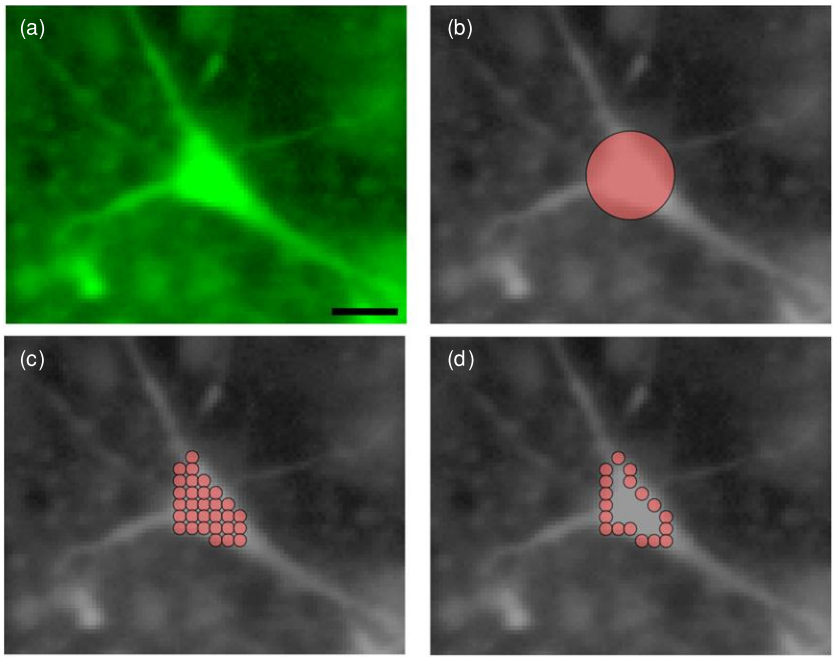

Figure 10. Illustrations of different illumination patterns. (a) A fluorescence image of a pyramidal neuron in culture. Scale bar is 10 μm. (b) Illumination with a single, large diffraction-limited spot, which does not exactly match the cell size or shape. (c) Illumination with an extended pattern, comprising adjacent, partly overlapping, small diffraction-limited spots. (d) Illumination area limited to a two-dimensional section of the cell membrane.

naturally required to generate illumination patterns of varying size and shape, not necessarily matched to the basic spot size and shape. Basic examples include matching the stimulation pattern to a dendritic branch during single-cell stimulation [27], matching the soma size (or sizes) with much smaller (high-resolution) spots when stimulating neuronal populations and creating illumination patterns that selectively and strongly excite optogenetic light-gated ion channels located on the cell's outer membrane (see figure 10). In these cases, one natural solution is to compose a contiguous illumination patch from many adjacent, partly overlapping diffraction-limited spots. However, when trying to project these patchy patterns, the resulting illumination is contaminated by speckle noise. Speckle is generated when the pattern phase freedom is used; adjacent spots have pseudo-random phases and therefore interfere randomly. For example, when two adjacent spots have opposite phases, a destructive interference will occur in their overlapping area, resulting in a dark spot in the middle of a uniform-intensity patch. The resulting distribution of intensities obeys the well-known statistics of the speckle phenomenon [47], justifying the term 'speckle noise'. A measure for the severity of this noise is speckle contrast, defined as the ratio of standard deviation to mean of the intensity [47]. Speckle noise has even greater impact in two-photon stimulation systems, because stimulation is proportional to the square of the intensity [28, 29].

A general solution to this problem involves imposing a phase-smoothing constraint into the GS algorithm [48, 49]. This requires over-sampling of the pattern, as well as post-iterations with a limited phase freedom; this increases computational load by orders of magnitude, making it impractical in our case.

Time averaging is a more practical approach, provided that the pattern changes much faster than the biological response time. Physical averaging can be done by adding a changing optical element such as a rotating diffuser [29]

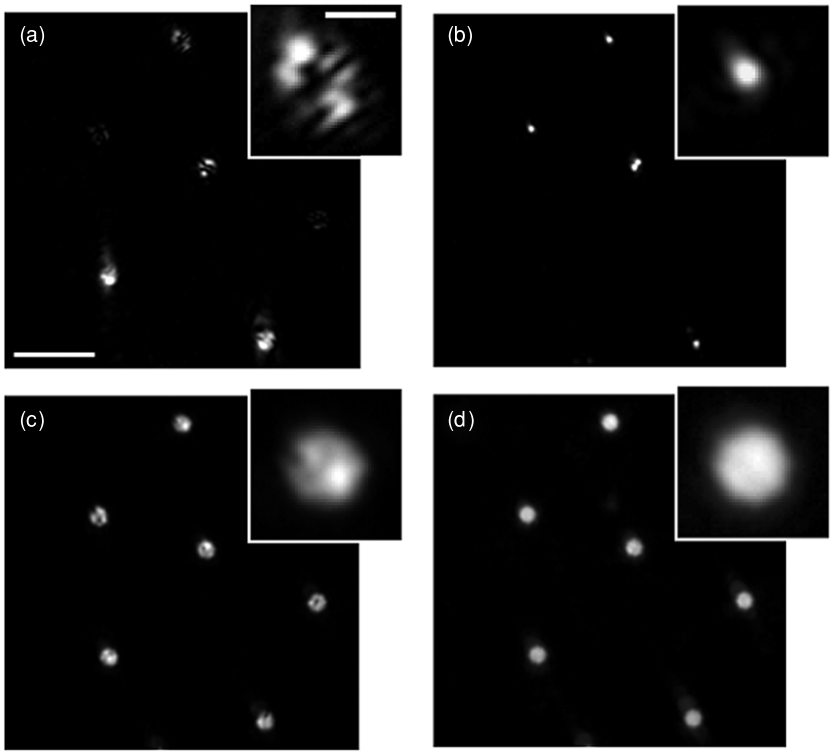

Figure 11. A stimulation pattern, containing several 5-pixel diameter circular patches, illustrating the speckle problem and the proposed solutions. In (a), the hologram is projected as is, resulting in severe speckle due to diffraction; scale bar is 50 μm. Small inset shows the top middle spot in detail; scale bar is 10 μm. One possible solution is dithering—dividing the spots into sparse subsets (24 in this case) such as the one shown in (b). Small inset shows that only a single diffraction-limited spot is included in this specific subset, so no overlap is expected. When the sparse subsets are displayed sequentially (panel c), speckle is reduced, although every subset requires the calculation of a different hologram. With shift-averaging, the basic hologram that created the pattern in (a) is repeated sequentially, with different cyclic shifts. With a correct selection of the 24 shifts, a speckle-free reconstruction is obtained (d).

that pseudo-randomly scrambles the resulting phases and reduces speckle contrast. However, this additional element adds complexity and power consumption to system design and also reduces image quality. Averaging may also be done at the hologram level [50], provided that the SLM is fast enough. If many holograms are calculated, each with a different random phase, the resulting patterns will have independent intensity fluctuations. Sequential display of the holograms will average the fluctuations, resulting in decreased speckle contrast. Sequentially displaying M holograms per pattern will reduce speckle contrast by $\sqrt{M}$, but will increase computational load by a factor of M, so this solution is far from perfect.

The speckle problem is not encountered in optical tweezer systems because the projected pattern is sparse and no real overlap exists between different spots. This leads us to offer a 'dithering' method. Instead of calculating a hologram that reconstructs full patches, we divide the spots comprising a contiguous patch into several sparse subframes, each containing a subset of the constituent spots. If the spots are divided pseudo-randomly between the subframes, high uniformity is achieved with the SR algorithm, as explained above. However, one still has to calculate many different holograms per projected pattern.

An improved approach, termed 'shift-averaging' was proposed recently by the authors [51]. It requires calculation of only a single hologram and eliminates speckle altogether. The basic idea is calculating one hologram per pattern and displaying a sequence of cyclic shifts of this single hologram. Mathematically, cyclic shifting adds a phase ramp at the Fourier plane, so this has the same effect of scrambling the phases and averaging out the fluctuations. However, an intelligent choice of a specific shift sequence eliminates the speckle. The details are given in [51].

Figure 11 illustrates the performance of dithering and shift averaging for illumination patterns containing contiguous patches.

We note that in some special cases one can potentially avoid these types of issues by using analytical rather than numerical hologram solutions. For example, intensity distributions that have been termed 'bottle-beams' [52] and are composed of two different beams which interfere destructively at the on-axis focal point, could potentially be used for selective optogenetic stimulation on a spherical shell geometry. In order to create the bottle beam, the hologram in equation (13) is multiplied by a phase-only hologram that converts the incident basic Gaussian mode to the sum of these two beams. However, the bottle beam illumination pattern will not be easily adapted to different cell shapes and orientations, and additionally, they have lower diffraction efficiencies than simple Gaussian beams.

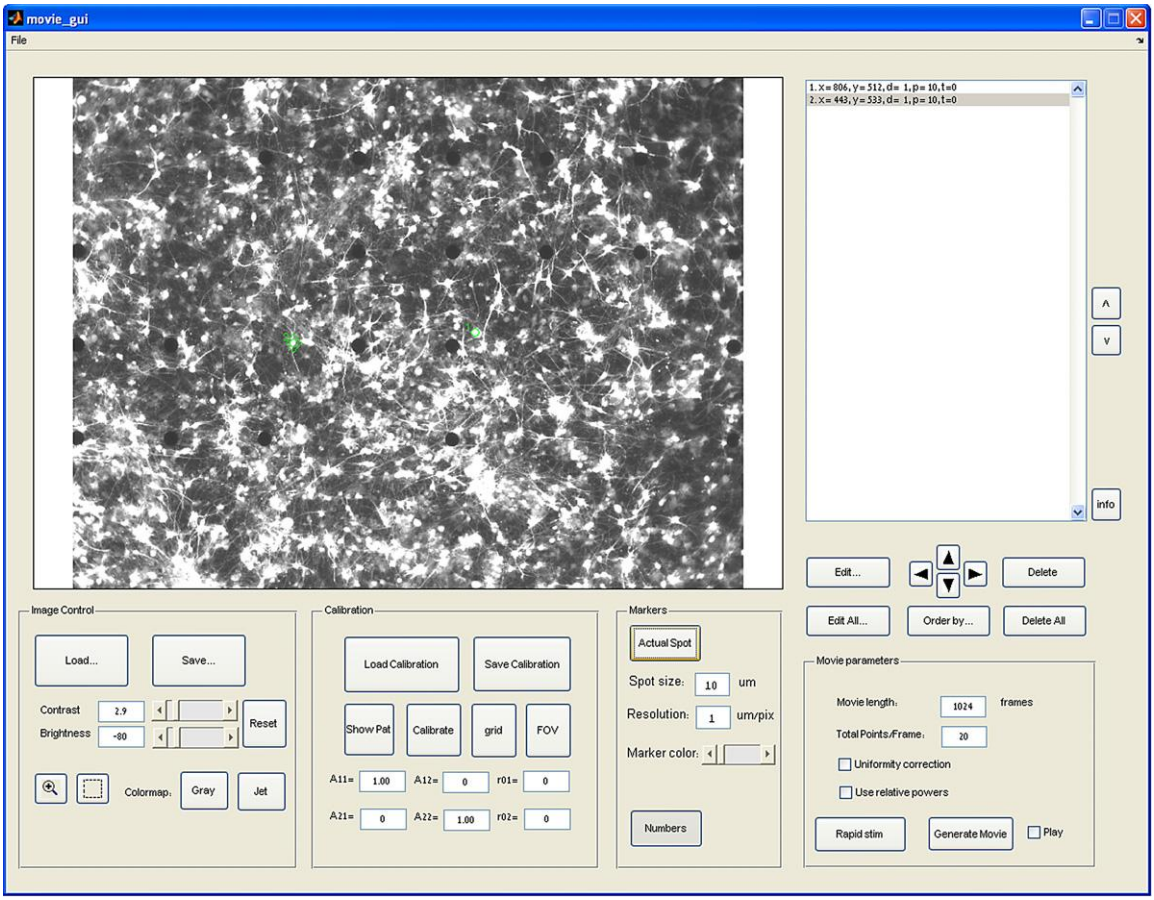

Figure 12. Graphical user interface that enables automatic generation of photo-stimulation sequences with full spatial and temporal control.

4.3. Pattern sequence generation by an automatic user interface

Similar to interactive optical trapping systems [53], the combination of stimulation and imaging in a single system enables one to identify, choose and target specific cells within a neural population [27]. After defining the required stimulation patterns in space and time, hologram sequences may be calculated and projected to generate a desired neural response. For research purposes, an automatic graphical user interface may be used to help the end user in performing all these functions.

In order to target specific cells identified in a microscope image, a transformation from the coordinates in the CCD image to the SLM coordinates is required. Neglecting optical distortions, this transformation is an affine one, that is, it contains only rotation, scaling and shifting, and therefore may be written as

$$\begin{pmatrix} u \\ v \end{pmatrix} = R \cdot \begin{pmatrix} X \\ Y \end{pmatrix} + \begin{pmatrix} X_0 \\ Y_0 \end{pmatrix} \tag{22}$$

where (X, Y) are the coordinates in the CCD image, R is a constant 2×2 matrix and (X_0, Y_0) is a constant shift. In order to calculate the constants, a calibration pattern containing several spots at known (u, v) locations is projected, its image is taken with the CCD and the (X, Y) positions are measured. Equation (22) is solved to yield the values of R and (X_0, Y_0).

In figure 12, we present a graphical user interface that enables calibration, loading microscope images, marking of individual cells, setting the parameters of their stimulation and finally calculating and displaying a hologram sequence.

4.4. Precise timing of pattern output

Once a hologram sequence has been calculated and stored, its projection requires only sending the data to the SLM. A common interface for communicating with an SLM is digital video interface (DVI), which is common for digital display devices. For high-speed binary ferroelectric devices, a single video frame sent over the interface contains several SLM frames, multiplexed into different bit planes.

However, if the data source is a PC, care should be taken to ensure that the signal timing is correct, and not affected by internal delays due to the PC operating system. These timing inaccuracies are common when standard movie-playing software is used and might distort the temporal structure of a well-designed experiment. Using low-level programming, one can get better control of video output timing, but this requires much effort and expertise. For rapid development of research systems, we have used the psychophysics toolbox extensions [54]. This MATLAB®-environment, free toolbox was developed mainly for designing psychophysical experiments, and gives very good control of display timing. Integrated into a high-level MATLAB®

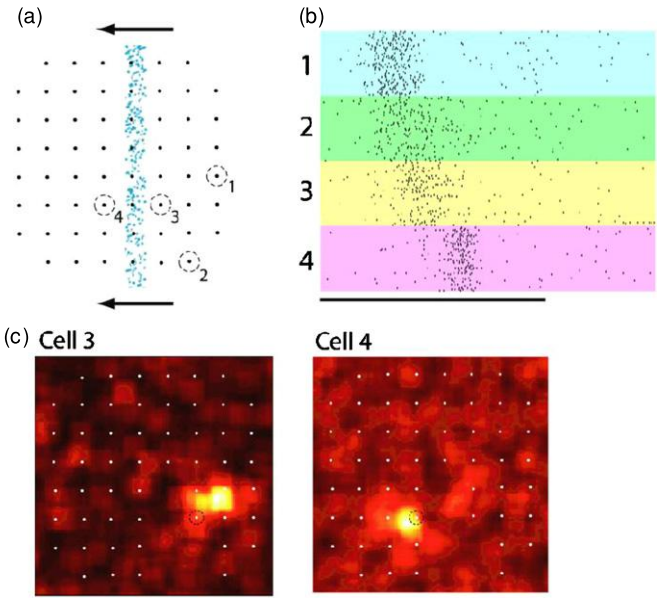

Figure 13. Neural responses of retina to holographic light patterns. (a) Neural responses to a moving bar pattern were recorded by a multi-electrode array. Four specific electrodes are marked. (b) Raster plot indicating patterned responses in accordance with bar position; the bottom line indicates the bar movement period—1.2 s. (c) Receptive field mapping of two different cells obtained by spike-triggered averaging of responses to a pseudo-randomly alternating stimulus. Dashed circles mark the recording electrodes in each case.

environment, it enables the developer to synchronize the flow of the program with the vertical retrace signal of the graphical adapter, so the projection timing of the sequences is accurate down to the single-frame level.

Synchronization of the holographic system with other equipment may be done by splitting the video output into an additional display and using a photodiode to extract the video timing signal. This signal may then be used as a trigger for data recording equipment.

5. Photo-stimulation results

We now present basic *in vitro* physiological results obtained using our example system. We performed acute experiments on rat retinas, stimulating the photo-receptors and recording from retinal ganglion cells (RGCs) using a multi-electrode array (MEA). The choice of a functioning retinal tissue—a well-studied, naturally photo-sensitive biological system, serves as a proof-of-concept test bed for the holographic projection system, before moving on to the more complex peculiarities of specific photosensitization methods.

The retinal tissue, placed on the electrode array, is centered inside the system FOV. In the first stage, 'moving-bar' patterns were used to scan the FOV for neural responses. In the second stage, randomly alternating patterns were projected, and neural responses collected over 10 min. The data were used to calculate a spike-triggered average, which estimates the spatio-temporal transfer function of the neural tissue. Figure 13 shows a typical experimental result.

6. Summary

In this article, we discussed the implementation of computer-generated holography for neural photo-stimulation systems. Using phase-only holograms, the light efficiency of sparse patterns is increased in comparison with conventional projection systems such as digital micromirror devices. We proposed an optical design scheme for such systems, and gave expressions for the field-of-view and resolution. As an example, we presented an illustrative holographic system. We also showed the basic design of holograms for reconstructing a two- or three-dimensional spot pattern, and gave the expressions for the coordinate-dependent diffraction efficiency. In addition, we discussed the speckle problem associated with multi-spot patterns and demonstrated two practical solutions: the dithering method and the shift-averaging method. Finally, we presented results of photo-stimulation with our illustrative system.

Acknowledgments

The authors acknowledge the financial support of the European Research Council (grant no 211055) and the Israel Science Foundation (1248/06). They would like to thank S Matar, T Wasserman and O Naor for critical reading of the manuscript, and two anonymous reviewers for their insightful comments. LG acknowledges financial support from the Gotwirth family. IR and NF acknowledge support from the Israeli science ministry.

References

[1] Arvanitaki A and Chalazonitis N 1961 Excitatory and inhibitory processes initiated by light and infra-red radiations in single, identifiable nerve cells *Nervous Inhibition* ed E Florey (New York: Pergamon) p 194
[2] Fork R L 1971 Laser stimulation of nerve cells in aplysia *Science* **171** 907–8
[3] Allegre G, Avrillier S and Albe-Fessard D 1994 Stimulation in the rat of a nerve fiber bundle by a short UV pulse from an excimer laser *Neurosci. Lett.* **180** 261–64
[4] Hirase H, Nikolenko V, Goldberg J H and Yuste R 2002 Multiphoton stimulation of neurons *J. Neurobiol.* **51** 237–47
[5] Adams S R and Tsien R Y 1993 Controlling cell chemistry with caged compounds *Ann. Rev. Physiol.* **55** 755
[6] Bi A *et al* 2006 Ectopic expression of a microbial-type rhodopsin restores visual responses in mice with photoreceptor degeneration *Neuron* **50** 23–33
[7] Boyden E S, Zhang F, Bamberg E, Nagel G and Deisseroth K 2005 Millisecond-timescale, genetically targeted optical control of neural activity *Nat. Neurosci.* **8** 1263–8
[8] Zhang F *et al* 2008 Red-shifted optogenetic excitation: a tool for fast neural control derived from *Volvox carteri Nat. Neurosci.* **11** 631–3
[9] Zhang F *et al* 2007 Multimodal fast optical interrogation of neural circuitry *Nature* **446** 633–9
[10] Wells J, Kao C, Jansen E D, Konrad P and Mahadevan-Jansen A 2005 Application of infrared light for *in vivo* neural stimulation *J. Biomed. Opt.* **10** 064003
[11] Wells J D *et al* 2007 Optically mediated nerve stimulation: identification of injury thresholds *Lasers Surg. Med.* **39** 513–26

[12] Wells J *et al* 2007 Biophysical mechanisms of transient optical stimulation of peripheral nerve *Biophys. J.* **93** 2567–80

[13] Izzo A D, Richter C P, Jansen E D and Walsh J T Jr 2006 Laser stimulation of the auditory nerve *Lasers Surg. Med.* **38** 745–53

[14] Farah N, Reutsky I and Shoham S 2007 Patterned optical activation of retinal ganglion cells *Conf. Proc. IEEE Eng. Med. Biol. Soc.* **2007** 6369–71

[15] Degenaar P *et al* 2009 Optobionic vision—a new genetically enhanced light on retinal prosthesis *J. Neural Eng.* **6** 035007

[16] Aravanis A M *et al* 2007 An optical neural interface: *in vivo* control of rodent motor cortex with integrated fiberoptic and optogenetic technology *J. Neural Eng.* **4** S143–56

[17] Bernstein J G *et al* 2008 Prosthetic systems for therapeutic optical activation and silencing of genetically targeted neurons *Optical Interactions with Tissue and Cells XIX* (San Jose, CA: SPIE)

[18] Saggau P, Bullen A and Patel S S 1998 Acousto-optic random-access laser scanning microscopy: fundamentals and applications to optical recording of neuronal activity *Cell Mol. Biol.* **44** 827–46

[19] Shoham S, O'Connor D H, Sarkisov D V and Wang S S 2005 Rapid neurotransmitter uncaging in spatially defined patterns *Nat. Methods* **2** 837–43

[20] Wang S *et al* 2007 All optical interface for parallel, remote, and spatiotemporal control of neuronal activity *Nano Lett.* **7** 3859–63

[21] Bryngdahl O and Wyrowski F 1990 Digital holography—computer-generated holograms *Progress in Optics* ed E Wolf (Amsterdam: Elsevier)

[22] Birch P M, Gourlay J, Love G D and Purvis A 1998 Real-time optical aberration correction with a ferroelectric liquid-crystal spatial light modulator *Appl. Opt.* **37** 2164–9

[23] Crossland W A *et al* 2000 Holographic optical switching: the 'ROSES' demonstrator *Lightwave Technol. J.* **18** 1845–54

[24] Jenness N J, Wulff K D, Johannes M S, Padgett M J, Cole D G and Clark R L 2008 Three-dimensional parallel holographic micropatterning using a spatial light modulator *Opt. Exp.* **16** 15942–8

[25] Georgiou A *et al* 2008 Liquid crystal over silicon device characteristics for holographic projection of high-definition television images *Appl. Opt.* **47** 4793–803

[26] Reicherter M, Haist T, Wagemann E U and Tiziani H J 1999 Optical particle trapping with computer-generated holograms written on a liquid-crystal display *Opt. Lett.* **24** 608–10

[27] Lutz C, Otis T S, DeSars V, Charpak S, DiGregorio D A and Emiliani V 2008 Holographic photolysis of caged neurotransmitters *Nat. Methods* **5** 821–7

[28] Nikolenko V, Watson B O, Araya R, Woodruff A, Peterka D S and Yuste R 2008 SLM microscopy: scanless two-photon imaging and photostimulation with spatial light modulators *Front Neural Circuits* **2** 5

[29] Papagiakoumou E, de Sars V, Oron D and Emiliani V 2008 Patterned two-photon illumination by spatiotemporal shaping of ultrashort pulses *Opt. Exp.* **16** 22039–47

[30] Reicherter M, Liesener J, Haist T and Tiziani H J 2003 Advantages of holographic optical tweezers *Novel Optical Instrumentation for Biomedical Applications* (Munich: SPIE)

[31] Goodman J W 2005 *Introduction to Fourier Optics* 3rd edn (Englewood, CO: Robert)

[32] Goodman J W and Silvestri A M 1970 Some effects of Fourier-domain phase quantization *IBM J. Res. Dev.* **14** 478

[33] Siegman A E 1986 *Lasers* (Sauslito, CA: University Science Books)

[34] Fukuchi N, Biqing Y, Igasaki Y, Yoshida N, Kobayashi Y and Hara T 2005 Oblique-incidence characteristics of a parallel-aligned nematic-liquid-crystal spatial light modulator *Opt. Rev.* **12** 372–7

[35] Martin-Badosa E, Montes-Usategui M, Carnicer A, Andilla J, Pleguezuelos E and Juvells I 2007 Design strategies for optimizing holographic optical tweezers set-ups *J. Opt. A: Pure Appl. Opt.* **9** S267–77

[36] Damman H 1970 Blazed synthetic phase-only holograms *Optik* **31** 95–104

[37] Tan K L *et al* 2001 Dynamic holography for optical interconnections. II. Routing holograms with predictable location and intensity of each diffraction order *J. Opt. Soc. Am.* A **18** 205–15

[38] Cottrell D M, Davis J A, Hedman T R and Lilly R A 1990 Multiple imaging phase-encoded optical elements written as programmable spatial light modulators *Appl. Opt.* **29** 2505–9

[39] Haupt C, Kolodziejczyk A and Tiziani H J 1995 Resolution and intensity distribution of output images reconstructed by sampled computer-generated holograms *Appl. Opt.* **34** 3077–86

[40] Seldowitz M A, Allebach J P and Sweeney D W 1987 Synthesis of digital holograms by direct binary search *Appl. Opt.* **26** 2788–98

[41] Liesener J, Reicherter M, Haist T and Tiziani H J 2000 Multi-functional optical tweezers using computer-generated holograms *Opt. Commun.* **185** 77–82

[42] Gerchberg R W and Saxton W O 1972 A practical algorithm for the determination of phase from image and diffraction plane pictures *Optik (Jena)* **35** 237–46

[43] Abe Y *et al* 2008 Special purpose computer system for flow visualization using holography technology *Opt. Exp.* **16** 7686–92

[44] Shimobaba T, Sato Y, Miura J, Takenouchi M and Ito T 2008 Real-time digital holographic microscopy using the graphic processing unit *Opt. Exp.* **16** 11776–81

[45] Curtis J E, Schmitz C H and Spatz J P 2005 Symmetry dependence of holograms for optical trapping *Opt. Lett.* **30** 2086–8

[46] Di Leonardo R, Ianni F and Ruocco G 2007 Computer generation of optimal holograms for optical trap arrays *Opt. Exp.* **15** 1913–22

[47] Goodman J W 1976 Some fundamental properties of speckle *J. Opt. Soc. Am.* **66** 1145–50

[48] Wyrowski F and Bryngdahl O 1989 Speckle-free reconstruction in digital holography *J. Opt. Soc. Am.* A **6** 1171–4

[49] Aagedal H, Schmid M, Beth T, Teiwes S and Wyrowski F 1996 Theory of speckles in diffractive optics and its application to beam shaping *J. Mod. Opt.* **43** 1409–21

[50] Amako J, Miura H and Sonehara T 1995 Speckle-noise reduction on kinoform reconstruction using a phase-only spatial light modulator *Appl. Opt.* **34** 3165–71

[51] Golan L and Shoham S 2009 Speckle elimination using shift-averaging in high-rate holographic projection *Opt. Exp.* **17** 1330–9

[52] Arlt J and Padgett M J 2000 Generation of a beam with a dark focus surrounded by regions of higher intensity: the optical bottle beam *Opt. Lett.* **25** 191–3

[53] Leach J *et al* 2006 Interactive approach to optical tweezers control *Appl. Opt.* **45** 897–903

[54] Brainard D H 1997 The psychophysics toolbox *Spat. Vis.* **10** 433–6