New Journal of Physics

Simulation of superresolution holography for optical tweezers

Laura C Thomson, Yannick Boissel[1]**, Graeme Whyte**[2]**, Eric Yao and Johannes Courtial**[3]

SUPA, Department of Physics and Astronomy, University of Glasgow, Glasgow G12 8QQ, UK
E-mail: j.courtial@physics.gla.ac.uk

Abstract. Optical tweezers manipulate microscopic particles using foci of light beams. Their performance is therefore limited by diffraction. Using computer simulations of a model system, we investigate the application of superresolution holography for two-dimensional (2D) light shaping in optical tweezers, which can beat the diffraction limit. We use the direct-search and Gerchberg algorithms to shape the center of a light beam into one or two bright spots; we do not constrain the remainder of the beam. We demonstrate that superresolution algorithms can significantly improve the normalized stiffness of an optical trap and the minimum separation at which neighboring traps can be resolved. We also test if such algorithms can be used interactively, as is desirable in optical tweezers.

[1] Present address: Observatoire de Paris-Meudon, 5 place Jules Janssen, 92190 Meudon, France.

[2] Present address: Department of Chemical Engineering, University of Cambridge, Cambridge CB2 3RA, UK.

[3] Author to whom any correspondence should be addressed.

IOP Institute of Physics Φ DEUTSCHE PHYSIKALISCHE GESELLSCHAFT

Contents

1. Introduction

Superresolution—shaping a signal to have features with a (spatial or temporal) frequency that is higher than what would be expected to be possible according to the frequency spectrum—works by leaving one part of the signal unconstrained, which in turn allows the other part to be controlled more tightly. Berry [1] illustrates that part of a frequency-bandwidth-limited signal can be shaped arbitrarily by explaining how to construct a signal whose bandwidth is a mere 1 Hz but that nevertheless reproduces Beethoven's ninth symphony exactly. Locally, the signal contains frequency components not contained in the signal as a whole. These new frequency components are called superoscillations.

In optics, superoscillations appear close to a beam focus [2], in a speckle pattern [3] and at the core of an optical vortex [4]. It is possible to engineer superoscillations; computer algorithms that calculate hologram patterns that shape light beyond the diffraction limit have been known for several decades (e.g. [5, 6]); new algorithms are still being developed (e.g. [7]–[9]). Quantum-mechanical wave functions have been specifically constructed to have pronounced superoscillations [10].

An obvious area of application for superoscillations is microscopy. Much work has gone into minimizing the (bright or dark) spot size of a light beam for use in microscopy [11], two-photon polymerization [12], and for the manipulation of molecules [13], for example by manipulating the distribution of polarization near the focus. The three-dimensional (3D) shape of the focus of a light beam is also important in optical tweezers, which trap and manipulate microscopic particles at the focus [14]. Manipulation of the polarization near the focus could also be applied to optical tweezers, but to the best of our knowledge this has never been done

and we do not investigate it here. One particular type of superresolution that recently has been applied to optical tweezers [15] uses the fact that materials with a negative refractive index can in principle perform perfect imaging, including even an amplification of evanescent waves [16]. This is not related to superresolution holography and is also not investigated here.

Superresolution holography has been suggested as potentially useful when applied to optical tweezers [17]–[19], but an application to optical tweezers has never been investigated in any detail. In this paper, we investigate whether superresolution holography could be useful in optical tweezers. We use a model system with a relatively small number of SLM pixels, and restrict ourselves to superresolution light shaping in one plane. Within this plane, we shape the beam only in an 'area of interest', which can then be shaped to smaller detail than the diffraction limit would suggest. We characterize 2D light distributions that can be achieved by two different superresolution algorithms in terms of quantities important for optical tweezers, namely minimum distance between two neighboring traps ('trap resolution') and trap stiffness. We use spatial-frequency constraints typical for holographic optical tweezers [20], and our implementations of the superresolution algorithms reflect the importance of the power in an optical trap by controlling the fraction of the overall power being shaped. We also investigate whether such algorithms are suitable for interactive use on a computer with modest computing power.

The paper is organized as follows. In section 2, we briefly review optical tweezers and the quantities of interest in optical micro-manipulation. With an application to optical tweezers in mind, section 3 outlines the ideas behind superresolution holography. The superresolution algorithms we have implemented, tailored to optical tweezers, are outlined in section 4. Section 5 describes the relationship between the fraction of power being shaped—the parameter we use to characterize the level of superresolution—and the trap depth. Section 6 describes benchmark values for trap resolution and relative stiffness. Sections 7 and 8, respectively, describe the improvements in relative trap stiffness and resolution we achieved with our algorithms. Section 9 investigates 3D superresolution. Section 10 tests the generality of the results gained from our model system. Section 11 discusses the suitability of superresolution algorithms for interactive use, before we conclude and discuss possible future work in section 12.

2. Optical tweezers

Optical tweezers [21, 22] use a tightly focused laser beam to trap microscopic dielectric objects, for example blood cells, at the laser beam's focus. Over the last few years, the use of computer-controlled phase holograms in the form of spatial light modulators (SLMs) has revolutionized the field: SLMs can do everything holograms can do, that is they can split up a light beam into multiple foci at arbitrary 3D positions and control the properties (such as brightness and shape) of each individual focus (i.e. trap). Most importantly, they can do all of this under real-time computer control at interactive speeds [20]. Figure 1 shows a simplified layout of typical holographic optical tweezers.

Some of the quantities important in optical tweezers are trap depth, stiffness and trap resolution. In the case of complete darkness surrounding the trap (i.e. the focus), trap depth is simply proportional to the intensity at the focus. We measure here the inverse of trap resolution, namely the minimum separation between two side-by-side traps such that the traps are still resolvable according to the Rayleigh criterion [23].
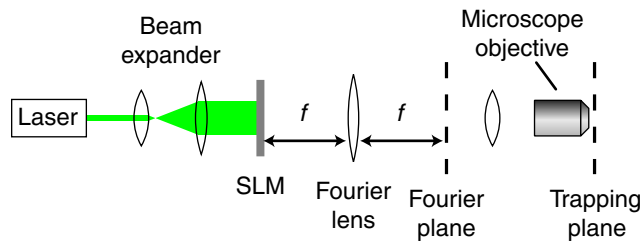
Figure 1. Layout of typical holographic optical tweezers. The beam from a laser is expanded before it hits a SLM. For simplicity, the SLM is drawn here as working in transmission, although most SLMs used in optical tweezers work in reflection. The first lens behind the SLM acts like a Fourier lens (focal length f): it Fourier-transforms the field in the SLM plane, which is placed in the Fourier lens's front focal plane, into its back focal plane—the Fourier plane. The remaining optics (lens(es) and a microscope objective) image the Fourier plane into the trapping plane.

Trap stiffness [24] requires a slightly longer discussion. A high stiffness, that is the ability to hold objects as still as possible, is important for precision measurements. It describes the 'spring constant' of the trap—the variation in the restoring force with distance from the trap center. In the limit of infinitely small particles, the trap stiffness is simply the 2nd derivative of intensity at the peak. The stiffness of a trap scales with intensity, and it depends on the shape of the peak: a narrow peak has a higher stiffness than a wide peak of the same peak intensity. We find (see section 8) that our superresolution-holography algorithms reduce a trap's peak intensity more than they decrease its width, so they can only reduce the stiffness of a trap from that achieved using non-superresolution holography. If, however, the absolute stiffness is less important than the shape of the trap, for example when a trap must be stiff without being too bright, superresolution holography can help. We characterise a trap's shape in terms of the ratio of the trap's stiffness and peak intensity, which we call its relative stiffness.

In the following, we investigate the trade-off between trap depth on the one hand and trap resolution and relative stiffness on the other hand. This complements earlier work on the precision with which optical traps can be steered holographically [25]. Such work is becoming increasingly important as optical tweezers are inching towards sub-micron, i.e. nano, manipulation.

3. Superresolution for optical tweezers

We approach superresolution in optical tweezers by first considering a familiar trade-off between power and trap resolution. The laser used in optical tweezers (figure 1) usually outputs a Gaussian beam; the size, w, of this Gaussian as it illuminates the SLM (of side length s—see figure 2) is important. One common way is to illuminate the SLM with a fairly narrow beam. The advantage of this is that essentially all the power in the beam hits the SLM. The disadvantage is that very little power passes through the edges of the SLM, which correspond to high spatial frequencies in the Fourier plane. Another way to illuminate the SLM is with a fairly wide beam. The advantage of that technique is that the edges of the SLM are fully utilized; the disadvantage is that a large fraction of the power is lost as it never hits the SLM. Full use of
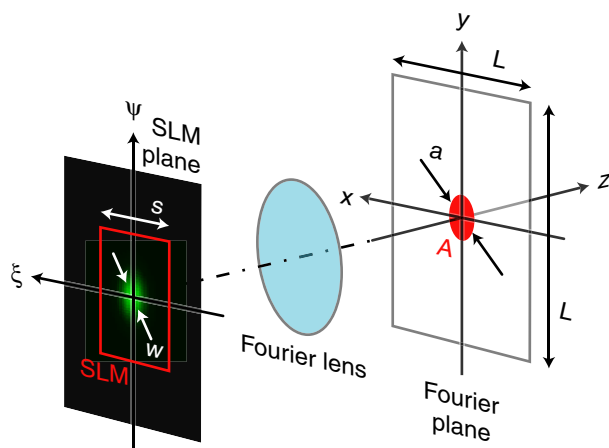
Figure 2. Geometry and illumination of a SLM by a Gaussian beam and of the area of interest in the Fourier plane. As the Gaussian beam hits the SLM it is of width w. Within the Fourier plane, we are interested here only in the field in a circular area A. The power in the beam that does not hit the SLM is lost, as is the power of the beam outside the area of interest. Our simulations are for a square SLM of side length s; the diameter of A is a. In the Fourier plane, the grid of data points represents a square of side length L. We use the Cartesian coordinates (ξ, ψ) in the SLM plane, and (x, y, z) in and around the Fourier plane.

high-spatial frequencies means better shaping and therefore greater stiffness and trap resolution, so wasting more light (by having a larger Gaussian) means better shaping of the light that is not wasted.

There is a similar trade-off in the trapping plane (see figure 2): losing more light by having a larger fraction of the light outside an 'area of interest', A, means better shaping of the light in A. Here, we ask what trap separations (section 7) and relative stiffnesses (stiffness divided by trap intensity; section 8) can be achieved as a function of the fraction p of power in the beam behind the SLM that passes through the area of interest, A.

It is perhaps worth at this point discussing a few possible choices for the area of interest, A. It can be chosen arbitrarily, depending on application and algorithm, but perhaps the most straightforward choice is to restrict trapping to a finite area and set A to this area plus a small margin around it. As explained in the following section, this is essentially what we chose to do in our model system, where a small disc-shaped area (see figure 4) covers all trap positions we require for the demonstrations in this paper. Another interesting choice for the area of interest is the sum of discs, one around each trap position. These discs could be chosen to be only just big enough to prevent trapped particles from interacting with the light surrounding A. Alternatively, the discs could be chosen to be so big that—for a defined limited trapping area—the disc around each trap also includes all other trap positions. As discussed in section 11, a pre-calculated hologram for a single superresolved trap could then be combined with standard tweezers algorithms to create arrangements of superresolved traps.

We also briefly investigate the effect of the size of the illumination beam in the SLM plane by running our relative-stiffness simulations with two illuminations: a narrow Gaussian with a waist size $w = 0.4s$ (where s is the side length of the SLM—see figure 2) and a wide Gaussian

with $w = 4s$, both truncated to be zero outside the SLM (the intensity at the corner pixels is 95% of the intensity in the center). The former is chosen such that most of the power in the Gaussian beam—$\approx$98%—hits the SLM, at the expense of not much light hitting the corners of the SLM; the latter is chosen to illuminate the entire SLM, at the expense of only a small fraction of the power—$\approx$3.4%—hitting the SLM.

4. Algorithms

We have implemented two tweezers-specific superresolution algorithms: a variation of the direct-search (DS) algorithm [26] with superresolution, and the Gerchberg algorithm [5]. The former finds better holograms, the latter is faster.

Both algorithms are iterative and find a phase distribution, $\phi_S(\xi_i, \psi_j)$, in the SLM plane, S; this phase distribution is the phase-hologram pattern that can then be displayed on the SLM. Under illumination with a given intensity distribution $I_S(\xi_i, \psi_j)$ (and plane phase), this produces an intensity distribution in the Fourier plane, F, that approximates a target intensity distribution $I_F(x_k, y_l)$ in the area of interest, A. Both algorithms control the 'superresolution strength' with the parameter P, which specifies the target fraction of the beam power in A. (Note that P does not measure superresolution, but controls it. How much superresolution the algorithms achieve is measured later, in sections 7 and 8.)

Our implementations of both algorithms are written in LabVIEW [27]. They represent transverse beam cross-sections on a grid of 64×64 complex numbers. In the SLM plane, the central 15×15 numbers represent the SLM pixels. The remaining elements in the 64×64 array representing the SLM plane are all zero; this zero-padding in the SLM plane increases the model resolution in its Fourier plane—the trapping plane—by a factor $64/15 = 4.3$. This, in turn, means that we can resolve features in the trapping plane that are up to 4.3 times smaller than the structure that corresponds to the plane waves with the largest transverse-spatial-frequency difference beating together, which is the 4-corner-pixel benchmark derived in section 6. In the trapping plane the area of interest, A, is a circular area of diameter 16 elements.

The number of SLM pixels in our model system is very small: whereas a typical SLM has of the order of 1000×1000 pixels, it is represented here by 15×15 pixels. On the other hand, the number of pixels is larger than the number of electrodes of many deformable mirrors; in [28], for example, the number of electrodes is 37. We chose the pixel number to be small compared to typical SLM-pixel-numbers for convenience, namely to speed up our calculations. More specifically, we perform our simulations for a model system with a pixel number that is big enough to allow nontrivial intensity shaping (two bright spots), but small enough to calculate the number of holograms required to calculate the results within a few days. We believe that the improvement factors we achieve in our model system are indicative of those that could be achieved in holographic optical tweezers (see section 10).

Our model can be translated into physical dimensions as follows. Because of the properties of the fast Fourier transform, which we use to transform between the planes, the wavelength λ of the light, the focal length f of the Fourier lens, the side length L of the square area in the Fourier plane represented by the 64×64 grid of complex numbers, and the center-to-center separation between neighboring elements in the SLM plane, Δ, which is also the SLM's pixel pitch, are related through the equation

$$\lambda/L = \sin\left(\tan^{-1}(\Delta/f)\right), \tag{1}$$

IOP Institute of Physics Φ DEUTSCHE PHYSIKALISCHE GESELLSCHAFT

which simplifies to

$$\Delta L \approx \lambda f, \tag{2}$$

for $\Delta \ll f$. For example, for an SLM with a pixel pitch of $\Delta = 10\,\mu$m, a wavelength $\lambda = 633$ nm, and a represented square of side length $L = 1$ mm, this corresponds to a Fourier lens of focal length $f \approx 16$ mm.

4.1. DS algorithm

The DS algorithm [26] repeatedly makes a random change to a hologram pattern, evaluates whether or not it represents an improvement, and accordingly keeps or discards it. In our implementation of the DS algorithm, the hologram pattern—the SLM pixels—consists of the 15×15 elements representing the SLM; the random change to the hologram pattern consists of randomizing the phase of a randomly chosen pixel; and the evaluation of the hologram comprises calculating the corresponding field in the Fourier plane and evaluating the similarity of the intensity distribution in the area of interest to the target intensity distribution. The details of the $(n+1)$th iteration of the algorithm are given below.

The phase distribution in the SLM plane returned by the previous (nth) iteration is $\phi_S^{(n)}(\xi_i, \psi_j)$, where ξ_i and ψ_j ($i, j = 1, \ldots, 15$) are the coordinates of the SLM pixels. (If $n = 0$, i.e. this is the first iteration, we arbitrarily set $\phi_S^{(0)}(\xi_i, \psi_j) = \varphi_{i,j}$, where $\varphi_{i,j}$ is the azimuthal angle associated with coordinate (ξ_i, ψ_j).) The algorithm picks a random SLM pixel and sets its phase to a random value between 0 and 2π, resulting in the altered phase distribution $\phi_{S,\text{test}}^{(n)}(\xi_i, \psi_j)$. This altered phase distribution is then combined with the illumination intensity distribution across the SLM, $I_S(\xi_i, \psi_j)$, to give the complex field there, namely $u_S^{(n)}(\xi_i, \psi_j) = (I_S(\xi_i, \psi_j))^{1/2}\exp(\mathrm{i}\phi_{S,\text{test}}^{(n)}(\xi_i, \psi_j))$.

This field is then zero-padded to 64×64 complex numbers to represent a larger area in the SLM plane, and fast-Fourier-transformed to give the corresponding field in the Fourier plane, $u_F^{(n)}(x_k, y_l)$, where x_k and y_l ($k, l = 1, \ldots, 64$) are the coordinates of the points represented in the Fourier plane. The algorithm then calculates an 'error' that is associated with this field. As we are only interested in the intensity distribution corresponding to $u_F^{(n)}$, $I_F^{(n)}(x_k, y_l) = |u_F^{(n)}(x_k, y_l)|^2$, we define the error as the squares of the differences between $I_F^{(n)}(x_k, y_l)$ and the target intensity distribution, $I_F^t(x_k, y_l)$, summed over all elements in the area of interest, A:

$$\epsilon_{\text{test}}^{(n)} = \sum_{(x_k, y_l) \in A} \left(I_F^{(n)}(x_k, y_l) - I_F^t(x_k, y_l) \right)^2. \tag{3}$$

Note that the target intensity distribution in A is scaled such that a fraction P of the power of the beam is in A, i.e.

$$\sum_{(x_k, y_l) \in A} I_F^t(x_k, y_l) = P \sum_{(x_k, y_l)} I_F(x_k, y_l). \tag{4}$$

Setting the parameter P to 1 turns the algorithm into a standard, no superresolution, algorithm; the closer P is set to zero, the stronger the superresolution effect the algorithm produces. If the intensity in A takes on exactly the target shape, and if the power in A is P times the total power in the beam, then $\epsilon_{\text{test}}^{(n)} = 0$. The more different $I_F^{(n)}$ is from I_F^t, the greater the value of $\epsilon_{\text{test}}^{(n)}$.

If an iteration reduces the error, i.e. if $\epsilon_{\text{test}}^{(n+1)} < \epsilon^{(n)}$, the altered phase distribution and resulting error are kept, otherwise they are discarded.

The rate at which iterations reduce the error falls off as the error converges towards a (local) minimum. We stop the algorithm after the error has remained constant for approximately 1000 iterations. This could easily be automated, but in our case, the operator determines when this point is reached by observing our program's error output.

4.2. Gerchberg algorithm

In the hope that it would enable interactive use, we have also implemented the Gerchberg algorithm [5], a variation of the well-known Gerchberg–Saxton algorithm [29], for use in optical tweezers. In this section, we briefly outline the Gerchberg–Saxton algorithm, describe how it is modified to become the Gerchberg algorithm, and then describe our implementation of the Gerchberg algorithm.

The Gerchberg–Saxton algorithm can be outlined as follows. Like before, the field is of interest in two planes, in our case the SLM plane and its Fourier plane. The intensity distributions in both planes are given: in the case of the SLM plane the intensity distribution is that of the illuminating laser beam, in the case of the Fourier plane it is the target intensity distribution. The Gerchberg–Saxton algorithm starts off with an arbitrary field in plane 1, replaces the intensity distribution with the one given for that plane (while not altering the phase distribution), calculates from this plane-1 field the corresponding field in plane 2, replaces the intensity distribution there with the appropriate given one, calculates from this altered plane-2 field the corresponding plane-1 field, replaces the intensity distribution with the given one again, etc. It can be shown [30] that the intensity distribution in the Fourier plane before replacement converges towards an approximation of the target intensity distribution. The phase distribution in the SLM plane therefore converges towards a useful phase-hologram pattern.

The Gerchberg algorithm is very similar to the Gerchberg–Saxton algorithm. The only difference is that each iteration does not replace the intensity distribution across the whole of plane 2, but only in a part of it, in our case the area of interest.

The details of our implementation of the Gerchberg algorithm are as follows. Like the DS algorithm, the Gerchberg algorithm starts with the phase distribution in the SLM plane calculated by the previous iteration (or, in the case of the 1st iteration, with an arbitrary phase distribution); for the $(n + 1)$th iteration, which we discuss here, the initial phase distribution is $\phi_S^{(n)}(\xi_i, \psi_j)$. The algorithm combines this phase distribution with the intensity distribution in the same plane, $I_S(\xi_i, \psi_j)$, to calculate the complex field there, namely $u_S^{(n)}(\xi_i, \psi_j) = (I_S(\xi_i, \psi_j))^{1/2} \exp(\mathrm{i}\phi_S^{(n)}(\xi_i, \psi_j))$. This field is then discrete-Fourier-transformed to give the corresponding field in the Fourier plane. The algorithm retains the phase of the field in the whole plane, and the intensity outside the area A. For the intensity distribution in A it substitutes the target intensity distribution, $I_F(x_k, y_l)$, to construct a new Fourier-plane field $u_F^{(n+1)}(x_k, y_l)$. In our tweezers-specific implementation of the Gerchberg algorithm, we scale the target intensity distribution in A such that it contains a given fraction P of the power in the entire plane F. The resulting field plane F is then inverse-Fourier-transformed to get the corresponding field in plane S, of which only the phase, $\phi_S^{(n+1)}(\xi_k, \psi_l)$, is retained. The process of replacement of the intensity distribution with that of the illumination laser, Fourier transform, replacement of the intensity distribution in the area of interest with the target one, and inverse Fourier transform is repeated until the process converges. Our program tests for convergence by comparing the separation between traps or relative stiffness for the latest iteration and 10 000 iterations ago; if the values are the same to within 0.01%, the program considers the algorithm to have converged.

IOP Institute of Physics Φ DEUTSCHE PHYSIKALISCHE GESELLSCHAFT
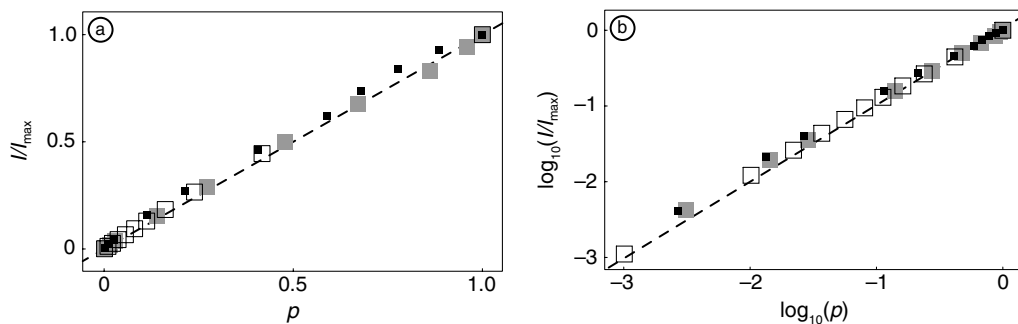
Figure 3. Relationship between trap intensity and power fraction in the area of interest. The graphs show the trap intensity, I, normalized with respect to I_{max}, the highest intensity we found for any value of p, as a function of p; (a) shows a linear–linear plot, (b) a log–log plot (base 10), which is very sensitive to small deviations from linearity for smaller values. The solid squares represent data points calculated using the DS algorithm, the empty squares with the Gerchberg algorithm; the smaller squares used narrow-Gaussian ($w = 0.4s$) illumination, the larger squares (gray in the case of filled-in squares) used wide-Gaussian ($w = 4s$) illumination. (No narrow-Gaussian-illumination Gerchberg-algorithm data were calculated.)

If N is the number of iterations after which this is the case, $\phi_S^{(N)}(\xi, \psi)$ is the phase-hologram pattern found by the algorithm.

5. Trap depth and power fraction in area of interest

As described in sections 3 and 4, we deliberately do not attempt to shape the beam in part of the Fourier plane, thereby gaining trap resolution in the area in which we do shape the beam—the area of interest. Once the area of interest has been selected, we control the fraction of the power in the beam to lie within the area of interest by setting the target fraction of the beam power in the area of interest, P. Note that in the target-plane intensity distributions $I_F(x_k, y_l)$ produced by our algorithms, the fraction p of power that lies in the area of interest, A, where

$$p = \frac{\sum_{(x_k, y_l) \in A} I_F(x_k, y_l)}{\sum_{x_k, y_l} I_F(x_k, y_l)} \tag{5}$$

can be significantly different from the target power fraction, P. We use this actual power fraction in A as a measure of 'inverse superresolution strength'.

The aim of this section is to demonstrate that, at least in all the examples we studied, the fraction of power in the area of interest, p, is roughly proportional to the intensity of the shaped peaks, and therefore approximately to the depth of the corresponding traps. (The latter approximation is exact for the case of peaks surrounded by complete darkness.)

Figure 3 shows the relationship between p and peak intensity for the DS and Gerchberg algorithms shaping the beam in the area of interest into a single peak, and using both narrow- and wide-Gaussian illumination. The figure demonstrates that this relationship is approximately linear in all cases we studied. We therefore believe that the values for the fraction of the power

Table 1. Benchmark values for relative stiffness, S_{xy}, and closest peak-to-peak separation, d, that can be achieved with standard holography. The fraction of the power in the area of interest, p, is also shown.

Illumination intensity	S_{xy} (elements^{-2})	d (elements)	p
Narrow Gaussian	0.24	–	0.99
Wide Gaussian	0.33	–	0.91
Uniform	0.33	–	0.90
4 corner pixels	0.80	4.6	0.048

in the area of interest, p, in all the graphs in this paper can also be interpreted as approximate relative peak intensity and consequently trap depth.

6. Benchmark values for trap resolution and relative stiffness

In this section, we calculate values for peak-to-peak separation and relative stiffness that can be achieved without running superresolution algorithms. These values will later serve as benchmarks for assessing the performance of our superresolution algorithms.

Perhaps the simplest benchmarks for the relative stiffness are the values that result from uniform phase across the SLM, as this is how a trap at the focus would be created in standard, non-superresolution and holography algorithms. The first three rows of table 1 show the values of the relative stiffness, S_{xy}, resulting from uniform phase across the SLM for illumination with a narrow Gaussian (width $w = 0.4s$, where s is the side length of the SLM—see figure 2), a wide Gaussian ($w = 4s$), and uniform intensity. It can be seen that the wide-Gaussian illumination, which, together with the narrow-Gaussian illumination, we will use as standard illumination throughout this paper, is very similar to uniform-intensity illumination.

The previous benchmarks were all for relative stiffness, as they correspond to light distributions consisting essentially of one peak. The simplest way to create multiple peaks in 1D is to beat two spatial frequencies together; to achieve the closest spacing of the peaks, the spatial frequencies should be as far apart as possible. We do just this in both the x- and y-directions: by illuminating the SLM's four corner pixels (equal brightness) and leaving all other elements dark we create in the SLM's Fourier plane a rectangular 2D array of peaks. A quick calculation reveals that, as neighboring corner pixels are separated by 14 elements in the x- or y-direction (the side length of the SLM in our model system is 15 elements), the corresponding number of peaks across the 64×64 elements of the discrete Fourier transform (DFT) is 14, so the peak-to-peak separation is $(64/14) = 4.6$ elements. Reassuringly, this is also the value found by our algorithm, shown in the fourth row of table 1, which also shows the relative stiffness of the peaks in this pattern to be very high.

It is worth noting two points about the corner-pixel illumination benchmark. Firstly, corner-pixel illumination can be seen as some form of superresolution shaping of one or two peaks as it does not actually create just the desired one or two peaks, but a whole array. This array spreads outside the area of interest, so the corresponding value of p is very low compared to the other benchmarks (right-hand column in table 1). Secondly, in order to compare like with like when measuring our superresolution algorithms against this benchmark, it should be taken
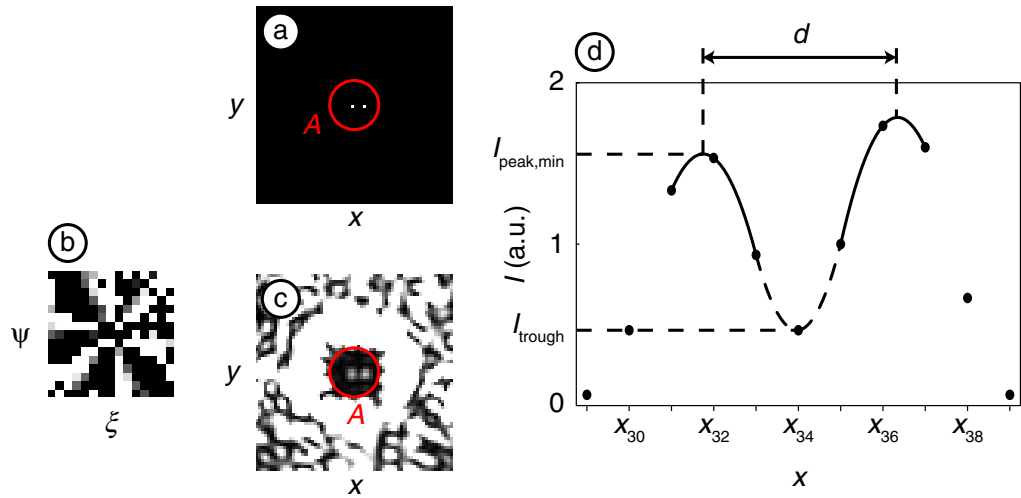
Figure 4. Example of superresolution. (a) Target intensity distribution in the area of interest, A; (b) gray-scale representation ($0 =$ black, $2\pi =$ white) of the phase hologram returned by the DS algorithm; (c) corresponding intensity distribution in the Fourier plane; (d) cross-section along the x-axis with fitted parabolas. The area of interest, A, is indicated by a red circle in (a) and (c). The intensity in (c) is saturated to make the peaks in A visible. This example was calculated using the DS algorithm for an illuminating-Gaussian-width of $w = 4s$, a target power fraction in A of $P = 0.001$, and a spacing of the bright elements in the target intensity distribution of $D = 4$ elements. The resulting actual fraction of the power in A was $p \approx 0.0026$, the actual peak separation was $d = 4.722$ elements.

into account that the power in the four corner pixels forms only a fraction of the power in our standard illuminations: if it was created from the narrow-Gaussian illumination (for example by putting a phase gradient across all other pixels such that the light gets focused to a point outside the area of interest, which would be a pattern that could be found by our superresolution algorithms), only a fraction $p_{c,1} = 0.000\,54$ of the illumination light hits the four corner pixels; if it was created from the wide-Gaussian illumination, the fraction of the power in the four corner pixels is $p_{c,2} = 0.017$. The fraction of the power that hits the four corner pixels and also gets diffracted into the area of interest is therefore the product of the corresponding value of p_c and the value of $p = 0.048$ from table 1, which gives $p = 0.000\,026$ ($\log_{10}(p) = -4.6$) for narrow-Gaussian illumination and $p = 0.000\,82$ ($\log_{10}(p) = -3.1$) for wide-Gaussian illumination.

7. Trap resolution

One limitation of optical tweezers, which is becoming increasingly important as micro-manipulation movestowards smaller scales, is their trap resolution: how close to each other can optical traps be positioned? We consider here the simplest case of two optical traps separated in the x-direction.

The simplest light distribution for achieving this consists of two bright foci at the target trapping positions, surrounded by complete darkness. Our target intensity distribution in the area of interest consists therefore of two bright elements (of equal brightness), separated by a

distance D, and surrounded by darkness. Figure 4(a) shows an example of such a target intensity distribution.

Figures 4(b) and (c), respectively, show the hologram found by the DS algorithm for the above example and the corresponding intensity distribution in the Fourier plane, calculated for a specific target fraction of the power in the area of interest of $P = 0.001$ and illumination of the SLM with a Gaussian of width $w = 4s$.

In order to establish the actual separation between the traps in the resulting intensity distribution—and whether or not they are actually resolvable—we evaluate this intensity distribution numerically. As discussed in section 4, zero-padding in the SLM plane by a factor of $64/15 \approx 4.3$ has already increased the resolution with which the field in the trapping plane is represented by a factor of 4.3 beyond that required to detect the 4-corner-pixel benchmark value (section 6), and also well beyond any structure size we detect later in this section. We fit parabolas to the center trough and the peaks on either side of the trough by fitting a quadratic function to the three data points closest to the trough or peak (figure 4(d)); the horizontal and vertical positions of the vertices of these parabolas allow us to establish with greater precision and accuracy the peak separation, d, and the intensities at the bottom of the trough, I_{trough}, and at the top of the smaller of the peaks, $I_{\text{peak,min}}$. We use the peak and trough positions to establish whether or not the two traps are resolvable according to the Rayleigh criterion, that is whether or not they satisfy the equation [23]

$$I_{\text{trough}} \leqslant \frac{8}{\pi^2} I_{\text{peak,min}} \approx 0.81 I_{\text{peak,min}}, \tag{6}$$

as the intensity cross-section in figure 4(d) clearly does. We also calculate p, the actual fraction of power in the area of interest, A, using equation (5); for the intensity distribution shown in figure 4(c), this fraction is $p \approx 0.26\%$.

Figure 5 shows the separations d between resolvable traps calculated by the DS algorithm, plotted as a function of the logarithm of the actual power in the area of interest, p. Note that the graph contains a relatively small number of data points due to our choice of plotting only those data points (out of many more we calculate) that correspond to a ratio of $I_{\text{trough}}/I_{\text{peak,min}}$ between 0.8 and $8/\pi^2 \approx 0.811$.

The trap separations in figure 5 are plotted as a fraction of the separation achieved by switching off superresolution, that is by running the algorithm with a target fraction of beam power in the area of interest of $P = 1$; this separation is $d_{\max} = 4.0$ pixels. Note that this value is already well below the four-corner-pixel benchmark value of 4.6 pixels (section 6). The other data points are even lower: the minimum separation we achieved—$d = 3.0$ pixels for a power ratio $p = 0.036$—is about 65% of the four-corner-pixel benchmark value. Note however that this is not comparing like with like: the trough intensity separating the peaks in the four-corner-pixel benchmark drops to zero, while the superresolved intensity distributions corresponding to the data points shown in figure 5 are selected such that the trough intensity only drops to $\approx 81\%$. For this reason, we do not plot the separations as a fraction of the four-corner-pixel benchmark value, but as a fraction of the separation achieved without superresolution, $d_{\max}$. Our minimum separation of $d = 3.0$ pixels is 75% of $d_{\max}$, which is a significantly reduced peak separation.

The straight line fitted to the data points in figure 5 gives a rough idea of the cost at which a reduction of the peak separation comes: it has a slope of ≈ 0.18, which means that a reduction of the power in A from 100 to 10%—i.e., directing 90% of the power out of the area of interest— reduces the separation between resolvable traps by merely $\approx 18\%$ of the original separation.
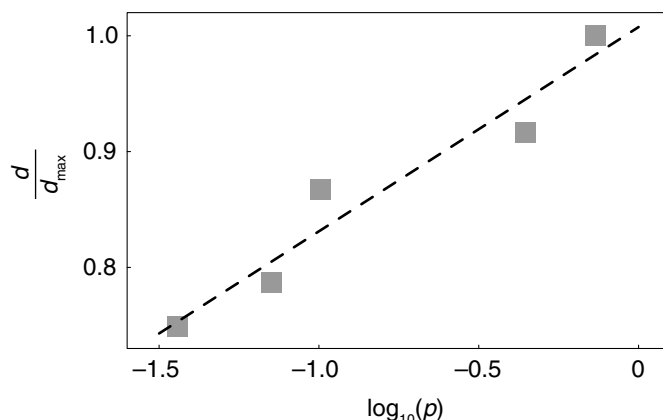
Figure 5. Resolvable peak separation as a function of the logarithm (base 10) of the actual fraction of the power ratio in the area of interest, p. The data points shown here represent light distributions only just resolvable according to the Rayleigh criterion: the ratio of the trough intensity to the intensity of the darker peak is between 0.8 and $8/\pi^2$. The resolvable peak separation is shown relative to that closest to $p = 1$, $d_{max} = 4.0$ pixels, and is therefore an indication for the factor by which the separation between two traps can be reduced through superresolution. The leftmost data point corresponds to a separation $d = 3.0$ pixels and a power ratio $p = 0.036$. The dashed line is a linear fit to the DS data. The data are calculated for illumination of the SLM with a wide ($w = 4s$) Gaussian.

8. Relative trap stiffness

In this section, we investigate the increase in the relative trap stiffness by using our superresolution algorithms. We chose a target intensity distribution consisting of one bright element surrounded by elements of zero brightness. (As an alternative, in the DS algorithm we could have chosen to optimize the relative trap stiffness directly.)

In order to evaluate the relative stiffness corresponding to a peak in the Fourier plane, we fit parabolas to the x and y intensity cross-sections through the peak. The stiffnesses in the x- and y-directions, s_x and s_y, are proportional to the second derivative at the vertex of the parabolas fitted to the x and y intensity cross-sections, respectively. We define the relative trap stiffnesses, S_x and S_y, as the corresponding stiffness divided by the intensity at the vertex. In the intensity distributions generated by our algorithms, these two stiffnesses are generally very similar, which is why we study only their average, $S_{xy} = (S_x + S_y)/2$.

Figure 6 shows plots of S_{xy} as a function of p, the fraction of the beam power in A, for both algorithms and two illuminations. This clearly shows that superresolution holography can significantly increase the relative transverse stiffness of an optical trap. The increase is greatest in the DS algorithm for illumination of the SLM with a small Gaussian; sacrificing 90% of power in A (i.e. $\log_{10}(p) = -1$), for example, improves S_{xy} by a factor ≈ 1.7 compared to the value achieved with the DS algorithm run with $p = 1$ (no superresolution), and by a factor ≈ 1.9 compared to the uniform-phase benchmark value.
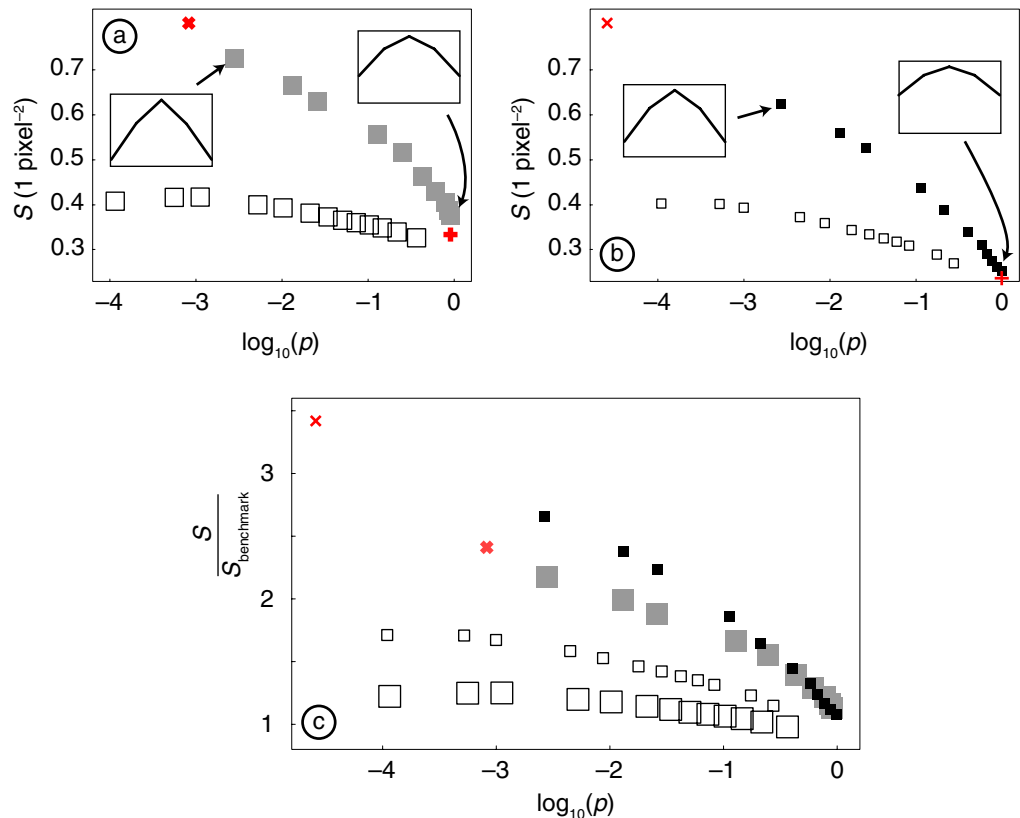
Figure 6. Relative transverse stiffness, S_{xy}, as a function of the actual fraction of the power in the area of interest, p. S_{xy} is plotted for illumination of the SLM with a wide ($w = 4s$) Gaussian (a) and a narrow ($w = 0.4s$) Gaussian (b). In each case, it was calculated using two superresolution algorithms, the DS algorithm (solid squares) and Gerchberg algorithm (empty squares). Insets indicate examples of the resulting peak shape by plotting a normalized intensity cross-section through the peak; the intensity range between 0 and the peak intensity is shown. (c) Shows the ratio of S_{xy} corresponding to different fractions of the power in A to the value achieved for uniform-phase illumination. Again, this is an indication for the ratio by which the relative stiffness can be increased using superresolution. As always, large squares indicate wide-Gaussian ($w = 4s$) illumination, small squares indicate narrow-Gaussian ($w = 0.4s$) illumination.

9. *z*-trapping properties

We have so far been concerned with shaping the light in one transverse plane ($z = 0$) only. The result was one or more small foci in a disc of darkness, which in turn is surrounded by brightness. Here we discuss briefly the trapping properties of such light fields in the direction perpendicular to the trapping plane, z.

When moving out of the $z = 0$ plane, the small foci will diffract very quickly, which suggests that the z-trapping properties of such fields are very good. On the other hand, the brightness surrounding the central dark disc might diffract into the center very quickly, spoiling
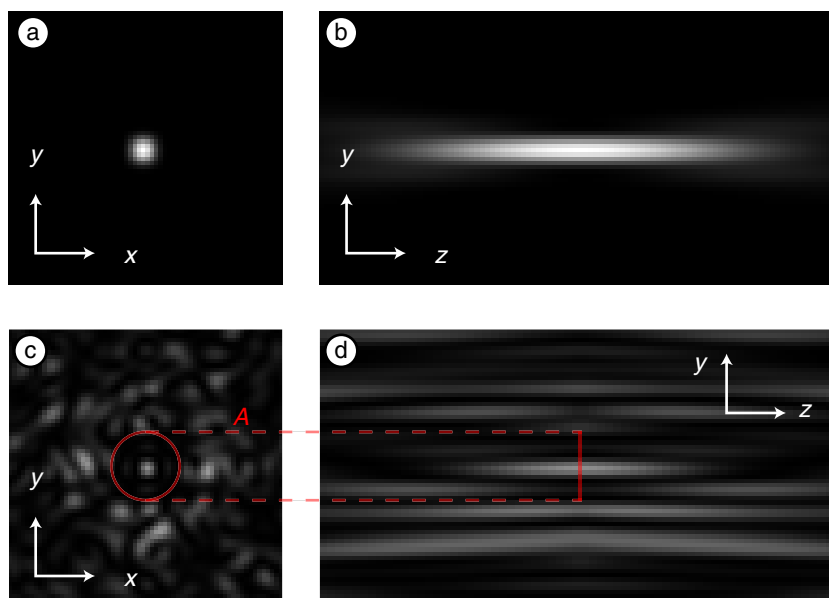
Figure 7. Intensity cross-sections in the xy- (left; i.e. (a) and (c)) and yz-planes (right; (b) and (d)) of a focus created without (top; (a) and (b)) and with (bottom; (c) and (d)) superresolution. xy cross-sections represent $1\,\text{mm} \times 1\,\text{mm}$ area, yz cross-sections show an area of $1\,\text{mm} \times 10\,\text{mm}$ area (compressed in the z-direction). The non-superresolution beam is the flat-phase benchmark beam (section 6). The superresolved beam was calculated with the DS algorithm with $p = 0.029$ (the target power fraction was $P = 0.01$). Both were calculated for narrow-Gaussian ($w = 0.4s$) illumination. Propagation of both beams was calculated for $\lambda = 633\,\text{nm}$ and $L = 1\text{mm}$. The normalized stiffnesses corresponding to the four frames are $S_{xy} = 0.24$ (a); $S_{xy} = 0.52$ (c), a 2.2-fold increase with respect to (a); $S_z = 0.000\,048$ (b); and $S_z = 0.000\,69$ (d), a 14-fold increase compared to (b).

the z-trapping properties. It is not obvious which of the two mechanisms dominates, and therefore what the z-trapping properties of our shaped light fields are.

We restrict ourselves here to demonstrating that it is possible to create situations in which the former mechanism dominates, corresponding to good z-trapping properties. This is also backed up by a recent investigation into the lifetime of superoscillations in quantum-mechanical wave functions [31]. We calculate the z-dependence of a specific light field. The light field we chose to investigate was the single peak calculated with the DS algorithm for illumination with a narrow-Gaussian-beam ($w = 0.4s$) and a power ratio in the area of interest of $p = 0.029$. To calculate the z evolution of the field we had to make a choice of λ and L; we chose $\lambda = 633\,\text{nm}$ and $L = 1\,\text{mm}$. Figure 7 shows the resulting intensity both in the xy (transverse) plane (figure 7(c)) and in the yz (longitudinal) plane (figure 7(d)). It can clearly be seen that the central intensity peak is surrounded by darkness not just in the xy-plane, but also in the yz-plane.

Figure 7 also shows the corresponding intensity cross-sections of the narrow-Gaussian, flat-phase, benchmark beam (figures 7(a) and (b)). Comparison confirms that the focus of the superresolved beam is narrower in the xy-plane, and that it is significantly narrower in the
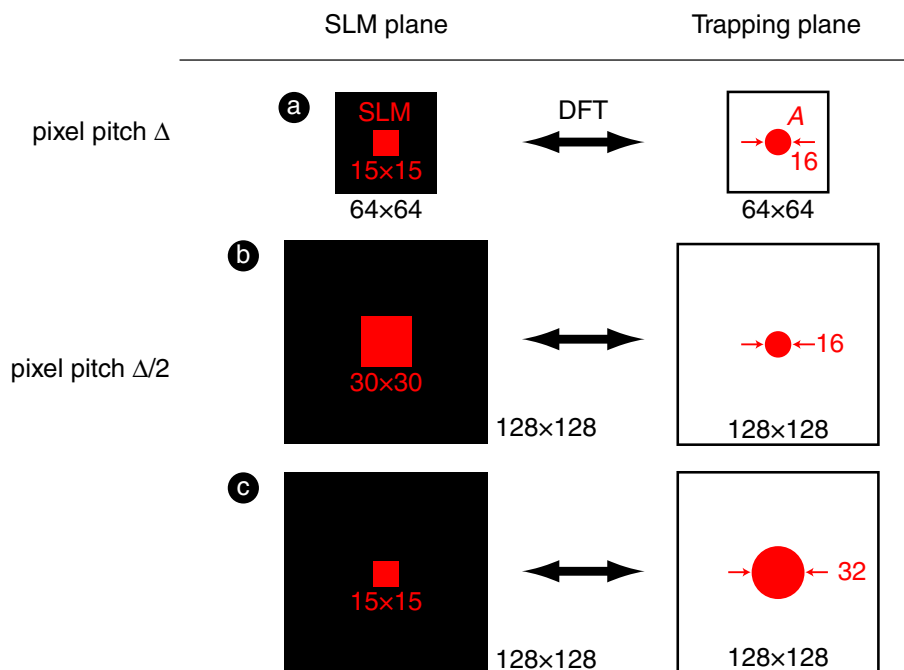
Figure 8. Modifications to the model system to test for dependence on model resolution and number of SLM pixels. The figure shows the relative size of the representation of the SLM plane (black square; left column) and the trapping plane (white square; right column) in the array of complex numbers that represents the light beam's complex field in three cases: (a) standard case; (b) double model resolution in SLM plane; and (c) double model resolution in trapping plane. In the model, the fields in the SLM plane and the trapping plane are related through a DFT. The red square in the center of the represented SLM-plane area indicates the SLM; the red circle in the center of the trapping-plane area indicates the area of interest. The dimensions of the 2D array representing the entire area is written in black (64×64 in (a); 128×128 in (b) and (c)), the size of the part of the array representing the SLM and area of interest, A, is written in red.

yz-plane. This is reflected in the 2.2-fold and dramatic 14-fold increases in the relative trap stiffnesses in the transverse (xy) and longitudinal (z) directions, respectively.

10. Are results from the model system indicative of the performance of systems with more SLM pixels?

As described in section 4, our model system uses a very small number of SLM pixels. To test the sensitivity of our results to the number of SLM pixels—and therefore to establish whether or not the results from our model system are indicative of the performance of high-resolution-SLM systems—we re-calculate some of the results with the number of SLM pixels doubled in each direction. This corresponds to a doubling of the model resolution in the SLM plane. We also re-calculated some of the results with the model resolution in the trapping plane doubled. Figure 8 summarizes the standard model system and the two modifications.
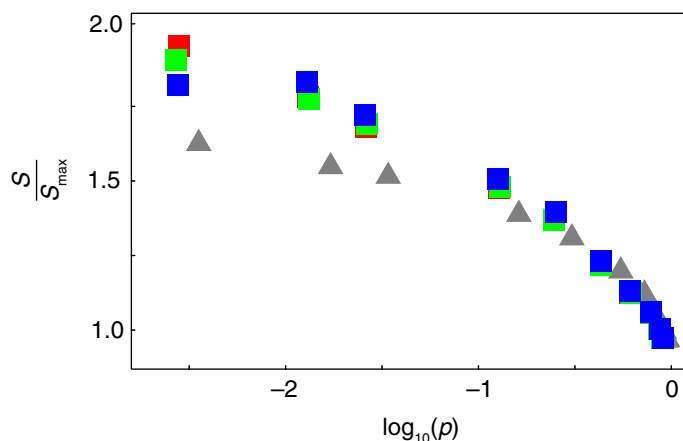
Figure 9. Improvements in the relative stiffness relative to the non-superresolution relative stiffness for a number of model resolutions, plotted as a function of p. Large squares correspond to the same physical area of interest represented in different resolutions: standard model resolution (red); double model resolution in the SLM plane (green); double model resolution in trapping plane (blue). Triangles correspond to a larger area of interest (represented using double model resolution in the SLM plane and standard model resolution in the trapping plane). All curves were calculated for wide-Gaussian ($w = 4s$) illumination.

Figure 9 shows the improvement factors of the relative stiffness versus $\log_{10}(p)$, calculated for the three cases outlined above. The graphs demonstrate that over the range of model resolutions we tested, the improvement factors relative to the non-superresolution values are independent of model resolution, but are dependent on other parameters, specifically SLM illumination and size of area of interest.

11. Are superresolution algorithms sufficiently fast for interactive use?

On our computer (Intel Pentium M processor, 2 GHz), the DS algorithm typically converges within 4 min, having completed of order 10^5 iterations, and the Gerchberg algorithm converges typically within 2 min, having completed of order 10^4 iterations. In a real optical-tweezers set-up, the number of SLM pixels would be significantly higher, as SLMs typically have of the order of 1000×1000 pixels, and like in our simulations this would have to be zero-padded so that the algorithm represents the trapping plane at a higher model resolution so as to be able to simulate any superresolution effects. This would increase the convergence time further. We therefore suspect that neither of our algorithms is suitable for interactive use, not even when implemented in a faster programming language.

There are a number of ways of using iterative hologram-calculation algorithms in optical tweezers. Perhaps the most obvious way is to run the algorithm for the latest target arrangement of trap positions/characteristics until it has converged, and then display the resulting hologram on the SLM. Both algorithms are clearly too slow to be used interactively in this way.

There are other ways to use these algorithms in optical tweezers; these might be suitable for interactive use in optical tweezers, but we do not investigate them here. One example is

IOP Institute of Physics Φ DEUTSCHE PHYSIKALISCHE GESELLSCHAFT

displaying on the SLM the hologram resulting from each iteration (or each ith iteration, where i is chosen to achieve a suitable update rate), which works well for the (non-superresolution) Gerchberg–Saxton algorithm and variations thereof [32]. Another example is using a pre-calculated hologram for a single, on-axis, high-relative-stiffness trap surrounded by a dark disc or sphere (which could be calculated with a straightforward extension of our DS algorithm, whereby the area/volume of interest includes the trap and the surrounding dark disc/sphere) as the kernel in the Curtis–Koss–Grier algorithm [20] to create multiple traps with high relative stiffness. Note that the latter would work only if each trap lies in the shifted dark discs/spheres around every other trap. If it does work, it might well turn out to be a fast general superresolution algorithm.

12. Conclusions and future work

Here, we have started to investigate theoretically the application of superresolution holography to optical tweezers. We have found that in our scaled-down model system, this can improve optical tweezers in a number of ways: it can improve the relative stiffness of individual traps, and it can make it possible to bring neighboring traps closer together while maintaining the identities of the individual traps. Unfortunately, these improvements come at the cost of a significant loss of trap depth.

We tried to make our results as useful as possible by formulating them in terms of improvement factors over benchmark values. However, these improvement factors are not universal: they do, for example, depend significantly on the intensity profile of the beam illuminating the SLM (see figure 6(c)) and on the size of the area of interest (section 10). The dependence on illumination and area of interest needs to be investigated in more detail, as does the lack of dependence on model resolution and number of SLM pixels, tested here only for a few examples over a small range.

Our algorithms in their simplest form are not fast enough to be used interactively in optical tweezers with moderate computing power, but they might form part of other algorithms, which we are planning to explore.

We are considering testing our predictions experimentally. We will need to overcome a number of difficulties, including aberrations in the optics, most likely the SLM, which will decrease the trap resolution, and light from the bright, and un-interesting, part of the beam being scattered into the interesting, but dark, part. We hope that it will be possible to overcome the latter problem, certainly in samples that inherently scatter little light. In order to overcome the former problem, we will be able to draw on expertise in our research group for correcting for such aberrations [33].

Acknowledgments

References

[1] Berry M 1994 *Faster than Fourier Fundamental Problems in Quantum Theory* ed J A Anandan and J Safko (Singapore: World Scientific)

[2] Karman G P, Beijersbergen M W, van Duijl A, Bouwmeester D and Woerdman J P 1998 Airy pattern reorganization and subwavelength structure in a focus *J. Opt. Soc. Am.* A **15** 884–99

[3] Gundu P N, Hack E and Rastogi P 2005 Superspeckles: a new application of optical superresolution *Opt. Express* **13** 6468–75

[4] Zambrini R, Thomson L C, Barnett S M and Padgett M 2005 Momentum paradox in a vortex core *J. Mod. Opt.* **52** 1135–44

[5] Gerchberg R W 1974 Super-resolution through error energy reduction *Opt. Acta* **21** 709–20

[6] Noel E, Khan R R and Dhadwal H S 1993 Optical implementation of a regularized Gerchberg iterative algorithm for superresolution *Opt. Eng.* **32** 2866–71

[7] Solomon J, Zalevsky Z and Mendlovic D 2003 Superresolution by use of code division multiplexing *Appl. Opt.* **42** 1451–62

[8] Jia J, Zhou C, Sun X and Liu L 2004 Superresolution laser beam shaping *Appl. Opt.* **43** 2112–7

[9] Zalevsky Z, García-Martínez P and García J 2006 Superresolution using gray level coding *Opt. Express* **14** 5178–82

[10] Kempf A and Ferreira P J S G 2003 Unusual properties of superoscillating particles *Preprint* quant-ph/0305148

[11] Blanca C M and Hell S 2002 Axial superresolution with ultrahigh aperture lenses *Opt. Express* **10** 893–8

[12] Wei P, Tan O F, Zhu Y and Duan G H 2007 Axial superresolution of two-photon microfabrication *Appl. Opt.* **46** 3694–9

[13] Helseth L 2006 Smallest focal hole *Opt. Commun.* **257** 1–8

[14] Molloy J E and Padgett M J 2002 Lights, action: optical tweezers *Contemp. Phys.* **43** 241–58

[15] Lu Z, Murakowski J A, Schuetz C A, Shi S, Schneider G J, Samluk J P and Prather D W 2006 Perfect lens makes a perfect trap *Opt. Express* **14** 2228–35

[16] Pendry J B 2000 Negative refraction makes a perfect lens *Phys. Rev. Lett.* **85** 3966–9

[17] de Juana D M, Oti J E, Canales V F and Cagigal M P 2003 Design of superresolving continuous phase filters *Opt. Lett.* **28** 607–9

[18] Cagigal M P, Oti J E, Canales V F and Valle P J 2004 Analytical design of superresolving phase filters *Opt. Commun.* **241** 249–53

[19] Canales V F, Oti J E and Cagigal M P 2005 Three-dimensional control of the focal light intensity distribution by analytically designed phase masks *Opt. Commun.* **247** 11–8

[20] Curtis J E, Koss B A and Grier D G 2002 Dynamic holographic optical tweezers *Opt. Commun.* **207** 169–75

[21] Ashkin A, Dziedzic J M, Bjorkholm J E and Chu S 1986 Observation of a single-beam gradient force optical trap for dielectric particles *Opt. Lett.* **11** 288–90

[22] Ashkin A, Dziedzic J M and Yamane T 1987 Optical trapping and manipulation of single cells using infrared-laser beams *Nature* **330** 769–71

[23] Born M and Wolf E 1993 *Principles of Optics* (Oxford: Pergamon)

[24] Rohrbach A 2005 Stiffness of optical traps: quantitative agreement between experiment and electromagnetic theory *Phys. Rev. Lett.* **95** 168102

[25] Schmitz C H J, Spatz J P and Curtis J E 2005 High-precision steering of multiple holographic optical traps *Opt. Express* **13** 8678–85

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

[27] National Instruments, 'LabVIEW'

[28] Ota T, Kawata S, Sugiura T, Booth M J, Neil M A A, Juškaitis R and Wilson T 2003 Dynamic axial-position control of a laser-trapped particle by wave-front modification *Opt. Lett.* **28** 465–7

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

[30] Soifer V, Kotlyar V and Doskolovich L 1997 *Iterative Methods for Diffractive Optical Elements Computation* (London: Taylor and Francis)

[31] Berry M V and Popescu S 2006 Evolution of quantum superoscillations and optical superresolution without evanescent waves *J. Phys. A: Math. Gen.* **39** 6965–77

[32] Sinclair G, Leach J, Jordan P, Gibson G, Yao E, Laczik Z J, Padgett M J and Courtial J 2004 Interactive application in holographic optical tweezers of a multi-plane Gerchberg–Saxton algorithm for three-dimensional light shaping *Opt. Express* **12** 1665–70

[33] Wulff K D, Cole D G, Clark R L, Di Leonardo R, Leach J, Cooper J, Gibson G and Padgett M J 2006 Aberration correction in holographic optical tweezers *Opt. Express* **14** 4170–5