

Conjugate gradient minimisation approach to generating holographic traps for ultracold atoms

Tiffany Harte,^{1,2} Graham D. Bruce,^{1,3} Jonathan Keeling,¹
and Donatella Cassettari^{1,*}

¹*SUPA School of Physics and Astronomy, University of St Andrews,
St Andrews, Fife KY16 9SS, UK*

²*Clarendon Laboratory, University of Oxford, Parks Road, Oxford, OX1 3PU, UK*

³*SUPA Dept. of Physics, University of Strathclyde, 107 Rottenrow, Glasgow, G4 0NG, UK*

[*dc43@st-andrews.ac.uk](mailto:dc43@st-andrews.ac.uk)

Abstract: Direct minimisation of a cost function can in principle provide a versatile and highly controllable route to computational hologram generation. Here we show that the careful design of cost functions, combined with numerically efficient conjugate gradient minimisation, establishes a practical method for the generation of holograms for a wide range of target light distributions. This results in a guided optimisation process, with a crucial advantage illustrated by the ability to circumvent optical vortex formation during hologram calculation. We demonstrate the implementation of the conjugate gradient method for both discrete and continuous intensity distributions and discuss its applicability to optical trapping of ultracold atoms.

© 2014 Optical Society of America

OCIS codes: (020.0020) Atomic and molecular physics; (020.1475) Bose-Einstein condensates; (020.7010) Laser trapping; (090.1760) Computer holography; (090.1995) Digital holography; (230.6120) Spatial light modulators.

References and links

1. I. Bloch, J. Dalibard, and S. Nascimbene, “Quantum simulations with ultracold quantum gases,” *Nat. Phys.* **8**, 267–276 (2012).
2. A. D. Cronin, J. Schmiedmayer, and D. E. Pritchard, “Optics and interferometry with atoms and molecules,” *Rev. Mod. Phys.* **81**, 1051–1129 (2009).
3. K. Henderson, C. Ryu, C. MacCormick, and M. G. Boshier, “Experimental demonstration of painting arbitrary and dynamic potentials for Bose-Einstein condensates,” *New J. Phys.* **11**, 043030 (2009).
4. C. Muldoon, L. Brandt, J. Dong, D. Stuart, E. Brainis, M. Himsforth, and A. Kuhn, “Control and manipulation of cold atoms in optical tweezers,” *New J. Phys.* **14**, 073051 (2012).
5. S. Bergamini, B. Darquié, M. Jones, L. Jacubowicz, A. Browaeys, and P. Grangier, “Holographic generation of microtrap arrays for single atoms by use of a programmable phase modulator,” *J. Opt. Soc. Am. B* **21**, 1889–1894 (2004).
6. V. Boyer, R. M. Godun, G. Smirne, D. Cassettari, C. M. Chandrashekar, A. B. Deb, Z. J. Laczik, and C. J. Foot, “Dynamic manipulation of Bose-Einstein condensates with a spatial light modulator,” *Phys. Rev. A* **73**, 031402 (2006).
7. B. Zimmermann, T. Müller, J. Meineke, T. Esslinger, and H. Moritz, “High-resolution imaging of ultracold fermions in microscopically tailored optical potentials,” *New J. Phys.* **13**, 043007 (2011).
8. D. Trypogeorgos, T. Harte, A. Bonnin, and C. Foot, “Precise shaping of laser light by an acousto-optic deflector,” *Opt. Express* **21**, 24837–24846 (2013).
9. G. D. Bruce, S. L. Bromley, G. Smirne, L. Torralbo-Campo, and D. Cassettari, “Holographic power-law traps for the efficient production of Bose-Einstein condensates,” *Phys. Rev. A* **84**, 053410 (2011).
10. G. D. Bruce, J. Mayoh, G. Smirne, L. Torralbo-Campo, and D. Cassettari, “A smooth, holographically generated ring trap for the investigation of superfluidity in ultracold atoms,” *Phys. Scr.* **T143**, 014008 (2011).

11. N. Houston, E. Riis, and A. S. Arnold, "Reproducible dynamic dark ring lattices for ultracold atoms," *J. Phys. B* **41**, 211001 (2008).
12. D. McGloin, G. Spalding, H. Melville, W. Sibbett, and K. Dholakia, "Applications of spatial light modulators in atom optics," *Opt. Express* **11**, 158–166 (2003).
13. M. Pasienski and B. DeMarco, "A high-accuracy algorithm for designing arbitrary holographic atom traps," *Opt. Express* **16**, 2176–2190 (2008).
14. A. L. Gaunt and Z. Hadzibabic, "Robust digital holography for ultracold atom trapping," *Sci. Rep.* **2**, 721 (2012).
15. F. Nogrette, H. Labuhn, S. Ravets, D. Barredo, L. Béguin, A. Vernier, T. Lahaye, and A. Browaeys, "Single-atom trapping in holographic 2D arrays of microtraps with arbitrary geometries," *Phys. Rev. X* **4**, 021034 (2014).
16. J. G. Lee and W. T. Hill III, "Spatial shaping for generating arbitrary optical dipoles traps for ultracold degenerate gases" (2014), <http://arxiv.org/abs/1406.4084>.
17. G. D. Bruce, M. Y. H. Johnson, E. Cormack, D. Richards, J. Mayoh, and D. Cassetari, "Feedback-enhanced algorithm for aberration correction of holographic atom traps" (2014), <http://arxiv.org/abs/1409.3151>.
18. J. R. Hui, X. Wu, and C. Warde, "Addressing large arrays of electrostatic actuators for adaptive optics applications," *Proc. SPIE* **5553**, 17–27 (2004).
19. J. Fortágh, H. Ott, S. Kraft, A. Günther, and C. Zimmermann, "Surface effects in magnetic microtraps," *Phys. Rev. A* **66**, 041604 (2002).
20. M. Mitchell, *An Introduction to Genetic Algorithms* (MIT, 1998).
21. M. Clark and R. Smith, "A direct-search method for the computer design of holograms," *Opt. Commun.* **124**, 150–164 (1996).
22. V. Boyer, C. M. Chandrashekar, C. J. Foot, and Z. J. Laczik, "Dynamic optical trap generation using FLC SLMs for the manipulation of cold atoms," *J. Mod. Opt.* **51**, 2235–2240 (2004).
23. M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos, "Iterative minimization techniques for ab initio total-energy calculations: molecular dynamics and conjugate gradients," *Rev. Mod. Phys.* **64**, 1045–1097 (1992).
24. R. Liu, B.-Z. Dong, G.-Z. Yang, and B.-Y. Gu, "Generation of pseudo-nondiffracting beams with use of diffractive phase elements designed by the conjugate-gradient method," *J. Opt. Soc. Am. A* **15**, 144–151 (1998).
25. R. Liu, B.-Y. Gu, B.-Z. Dong, and G.-Z. Yang, "Design of diffractive phase elements that realize axial-intensity modulation based on the conjugate-gradient method," *J. Opt. Soc. Am. A* **15**, 689–694 (1998).
26. R. Liu, B.-Y. Gu, B.-Z. Dong, and G.-Z. Yang, "Diffractive phase elements that synthesize color pseudo-nondiffracting beams," *Opt. Lett.* **23**, 633–635 (1998).
27. G. Zhou, X. Yuan, P. Dowd, Y.-L. Lam, and Y.-C. Chan, "Efficient method for evaluation of the diffraction efficiency upper bound of diffractive phase elements," *Opt. Lett.* **25**, 1288–1290 (2000).
28. G. Zhou, X. Yuan, P. Dowd, Y.-L. Lam, and Y.-C. Chan, "Design of diffractive phase elements for beam shaping: hybrid approach," *J. Opt. Soc. Am. A* **18**, 791–800 (2001).
29. W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes: The Art of Scientific Computing* (Cambridge University, 1987).
30. J. R. Shewchuk, "An introduction to the conjugate gradient method without the agonizing pain," (Carnegie Mellon University, School of Computer Science, Pittsburgh, Pa., 1994).
31. M. Johansson and J. Bengtsson, "Robust design method for highly efficient beam-shaping diffractive optical elements using an iterative-fourier-transform algorithm with soft operations," *J. Mod. Opt.* **47**, 1385–1398 (2000).
32. A. Bartok-Partay, S. Cereda, G. Csanyi, J. Kermodé, I. Solt, W. Szlachta, C. Varnai, and S. Winfield, <http://www.libatoms.org>.
33. P. Senthilkumar, F. Wyrowski, and H. Schemmel, "Vortex stagnation problem in iterative Fourier transform algorithms," *Opt. Laser Eng.* **43**, 43–56 (2005).
34. D. G. Grier, "A revolution in optical manipulation," *Nature* **424**, 810–816 (2003).
35. C. Maurer, A. Jesacher, S. Bernet, and M. Ritsch-Marte, "What spatial light modulators can do for optical microscopy," *Laser Photon. Rev.* **5**, 81–101 (2011).
36. M. Reichert, M. Wu, and L. Seifert, "Using graphics boards to compute holograms," *Comput. Sci. Eng.* **8**, 8–13 (2006).
37. N. Masuda, T. Ito, T. Tanaka, A. Shiraki, and T. Sugie, "Computer generated holography using a graphics processing unit," *Opt. Express* **14**, 603–608 (2006).

1. Introduction

In recent years there has been extraordinary progress in cold atom physics and its applications in fields such as quantum computation and simulation of condensed-matter systems, precision measurements, and matter-wave interferometry [1, 2]. In this context, arbitrary time-dependent optical trapping potentials are particularly appealing, with a variety of geometries including

toroids and ring lattices already realised by acousto-optic or holographic means [3, 4]. Experiments have been performed with discrete arrays of optical dipole traps, loaded with either thermal atoms [4, 5] or quantum degenerate atomic gases [3, 6, 7], in which individual trapping sites can be moved, addressed and manipulated. Important too are continuous trapping geometries: the primary subject of the present work are extended (as opposed to diffraction-limited) power-law potentials, proposed both as a static supplement to a trapping potential to cancel unwanted external potentials [8], and in a dynamic sequence as a tool for the efficient production of Bose-Einstein condensates [9]. Other interesting continuous potentials include engineered waveguides with dynamic bright regions, shown to be suitable for studies of BEC superfluidity [10].

Technologies employed so far in the realisation of these arbitrary optical trapping patterns include acousto-optic deflection of a laser beam to produce either a composite static intensity distribution [8] or a rapidly-scanned profile [3, 7, 11], digital micro-mirror devices (DMDs) [4], and computer-generated holograms implemented with phase-only spatial light modulators (SLMs) [5, 6, 9, 10, 12–17]. The high phase-resolution available in phase-only SLMs offers significant advantages for versatility of the accessible trapping patterns, though at the cost of lower switching speed between frames if compared to acousto-optic modulators and digital mirror devices. However, with new technologies currently being developed for grey-scale phase-only SLMs with kHz refresh rates [18], this versatility may become accessible at sufficiently high update rates for high-speed dynamic manipulation of trapped atoms. The primary challenge of phase-only SLMs is the computational complexity inherent in reproducing the target intensity distribution on the trapping plane. This paper demonstrates an alternative reliable and efficient method to address this problem.

Our investigation concerns an SLM consisting of 256×256 programmable pixels; each pixel is able to impose a phase retardation between 0 and 2π in steps of $2\pi/256$ on an incident laser beam. The resulting digital hologram is calculated to reconstruct a given target intensity pattern in the far field, or equivalently in the focal plane of a lens, in which atoms will be trapped. The calculated phase mask ϕ_{pq} and the incident laser field $A_0 S_{pq}$, with indices p and q denoting pixel position, determine the SLM-plane electric field:

$$E_{in} = A_0 S_{pq} \exp(i\phi_{pq}). \quad (1)$$

We express this electric field as an array of N pixels; propagation through focussing optics can be calculated by a fast Fourier transform. The electric field in the output plane is therefore given by

$$E_{out} = \frac{A_0}{N} \sum_{pq} S_{pq} \exp(i\phi_{pq}) \exp\left(-\frac{2\pi i}{N}(pn + qm)\right), \quad (2)$$

with output-plane coordinates denoted by n and m . As only the modulus of E_{out} is relevant for optical trapping, we have output-plane phase freedom: consequently the phase ϕ_{pq} required to recreate a given target intensity is not unique, and solutions are found numerically. The general problem of phase retrieval, which includes both the above case of laser beam shaping where the modulus of the field is known in both planes, and the related problem of image reconstruction where the field modulus is known only in the output plane, can be solved using a variety of methods. These fall broadly into two categories: iterative Fourier transform algorithms (IFTAs), and algorithms based on the minimisation of a cost function. The IFTA calculation encourages convergence from an initial phase guess to one yielding the target intensity using successive Fourier transforms between SLM and output planes, imposing the known or desired electric field amplitude at each step. For the purpose of atom trapping in arbitrary geometries, where smoothness is a primary consideration (as corrugations of the trapping potential cause fragmentation of cold atom clouds [19]), the best results so far have been achieved by the mixed-region

amplitude freedom (MRAF) variant of the IFTA [9, 10, 13, 14, 17]. In MRAF, the output plane is divided into two regions: a signal region in which the intensity is restricted to match the target intensity pattern, and a noise region in which the intensity is unconstrained. This separation allows for increased accuracy and smoothness in the signal region, leading to computed intensity patterns with residual root-mean-square (RMS) errors of less than a few percent. A secondary consideration in designing optical traps is the light-usage efficiency of the computer-generated hologram. One motivation for the use of phase-only spatial light modulators rather than amplitude modulators is that the former do not deliberately remove light from the incident beam. However, the MRAF algorithm gains accuracy by deliberately lowering this efficiency. Furthermore, lacking a minimisation principle, IFTA approaches provide no guarantee of converging, and their final state can be highly dependent on the initial phase pattern used.

In contrast, cost function minimisation algorithms are inherently more directional than IFTAs: the cost function encodes all constraints and desired properties of the intensity pattern, and can be designed with terms accounting for specific output plane features (such as high light-usage efficiency) in addition to adherence to the target intensity profile. Within this category, established beam shaping methods include genetic algorithms and direct search algorithms, both of which are less computationally efficient than IFTAs [13]. Genetic algorithms [20] seek the global minimum of the cost function, and as such are computationally demanding but accurate. Direct search algorithms [21, 22], in which SLM pixel values are sequentially altered with only changes reducing the cost function being retained, are limited to just a few phase levels due to computational intensity, and as such work well for simple targets but struggle to reproduce more intricate patterns.

In this paper, we consider an alternative approach to the beam shaping problem, in which the cost function is minimised by a conjugate gradient local search algorithm. Conjugate gradient minimisation, a well-established method for minimising high-dimensional smooth functions, is widely used in contexts such as electronic structure [23]. Here we find that this approach successfully combines computational efficiency and algorithm versatility, allowing the accurate reproduction of a variety of target intensity profiles relevant for optical trapping of atoms. The simplest cost function we study, a least-squares difference from the desired pattern, leaves localised defects which have low cost but present significant problems for atom trapping. However, these defects can be removed by systematically modifying the cost function, and we discuss the forms of cost functions required to eliminate them. This flexibility in cost function definition, a useful feature common to all minimisation algorithms, also allows us to go beyond the simple definitions of signal and noise regions, and to fine-tune our algorithm for different experimental requirements simply by adding cost function terms and applying different weightings across the output plane.

Laser beam shaping via an algorithm that relies entirely on gradient-based local search has been relatively unexplored so far. Examples of this are found in [24–26], where conjugate gradient minimisation is used to generate pseudo-non diffractive beams in which the target is a given axial intensity distribution. Gradient-based local search techniques are also used as part of more complex algorithms, such as hybrid algorithms for beam shaping in which they are combined with genetic algorithms to increase the reliability in locating the global minimum for the generation of flat top laser beams [27, 28]. The present work shows that gradient-based local search can be applied to a much wider range of beam shaping problems, and that it can be applied on its own: finding a local minimum is sufficient for a good reproduction of a variety of targets.

2. The conjugate gradient calculation method

At the heart of this method is quantifying the error between target and predicted intensity by defining a cost function C . Minimisation of C is performed over all SLM pixel phase values, while computational efficiency is ensured by incorporating gradient information. Figure 1 illustrates the application of conjugate gradient minimisation to hologram calculation.

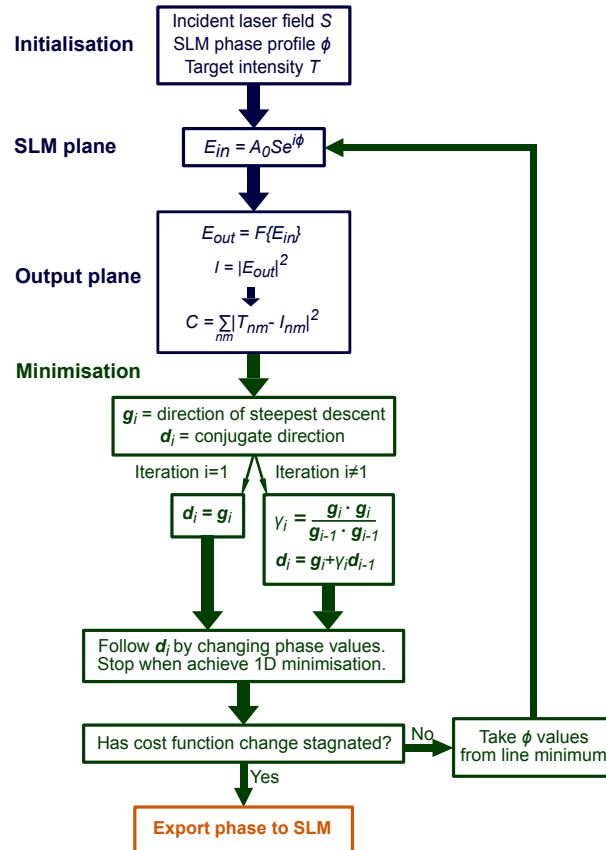


Fig. 1. Block diagram illustrating the conjugate gradient minimisation approach to hologram solution.

The process is initialised by defining a target intensity distribution, an incident laser field amplitude, and an initial phase guess. This may be the summation of analytically-expressed phase patterns to form an educated guess, e.g. quadratic and linear phase gradients giving expansion and position-offset of the intensity in the output plane [13], or it may be a random array taking values between 0 and 2π , with each element corresponding to an SLM pixel.

Weighting arrays are also defined during this initialisation stage. This weighting process is significantly more flexible than the simple definition of signal and noise regions characteristic of the MRAF IFTA. Output plane regions can be arbitrarily weighted according to their importance: we can allocate a pixel-dependent prefactor to individual cost function terms according to their relative importance in different output plane regions. For the remainder of this paper, the high-intensity output region forming the trapping pattern is referred to as the trapping region, with the signal and noise regions retaining the same meanings as in MRAF: the signal region is the trapping region plus some border which will remain devoid of light, while the noise re-

gion is the remainder of the plane where light may be deposited without adversely affecting the trapping potential. For example, the cost function can be allocated a larger weighting, so greater importance, in the trapping region than the remainder of the output plane, while the highest intensity parts of this region preferentially seen by the atoms can be given yet more prominence than those of lower intensity. With our cost function approach, we can provide a smooth transition between signal and noise region weightings which reduces noise accumulation at the signal region border. By comparison, a typical light pattern calculated using MRAF places much of the noise region light at the boundary of the signal and noise regions.

Upon each iteration, the output electric field corresponding to the current phase profile is calculated, with phase array components generated by the minimisation routine. This multidimensional minimisation is composed of one-dimensional steps, each seeking to minimise the cost function by changing phase values. The initial step minimises the cost function in the local gradient direction; subsequent consecutive minimisation directions are conjugate and independent, to avoid repetition of minimisation directions [29, 30]. Conjugate directions \mathbf{d} are those satisfying [23, 30]:

$$\mathbf{d}_i \cdot H_C \cdot \mathbf{d}_j = 0, \quad (3)$$

where H_C is the Hessian matrix of the cost function. However, the power of the conjugate gradient descent approach is that the Hessian does not need to be calculated explicitly, but is rather built up by application of the following procedure. The i th direction is calculated using [23]:

$$\mathbf{d}_i = \mathbf{g}_i + \gamma_i \mathbf{d}_{i-1} \quad (4)$$

where \mathbf{g}_i is the direction of steepest descent at the termination point of step $i - 1$ and γ_i is the scalar

$$\gamma_i = \frac{\mathbf{g}_i \cdot \mathbf{g}_i}{\mathbf{g}_{i-1} \cdot \mathbf{g}_{i-1}}. \quad (5)$$

The fast Fourier transforms used in the calculations map a $N \times N$ array in the SLM plane onto a $N \times N$ array in the output plane. If N is chosen to be the number of pixels in the SLM, then the size of each output plane pixel is exactly the diffraction limit of the system. Aliasing in the output plane is avoided by selecting a value for N , in accordance with the Nyquist criterion, of twice the number of pixels in the SLM [31]. Therefore the phase array is surrounded by zeroes to double its size and optimise output plane sampling at each iteration. Correspondingly enlarging the target array, the resolution of the cost calculation is optimised. Iteration continues until the difference between consecutive cost values stagnates: a minimum of the chosen cost function has been located. Our calculations make use of the libatoms library [32].

3. Versatility via cost function definition

The cost function should be such that its minimum corresponds to the desired pattern. There are however a number of other features required for the algorithm to operate efficiently: it should be efficient to evaluate derivatives of the cost function; the function should not have local minima which give poor trapping profiles. Indeed, as we will see below, a naive choice of cost function leads to local minima containing optical vortices. Furthermore, the conjugate gradient approach assumes an approximately quadratic function [29]. A simple cost function may purely concern target reproduction accuracy, expressed as a sum over output plane pixels (n, m) of differences between target T_{nm} and calculated output intensity:

$$C = \sum_{nm} \left(T_{nm} - |E_{out, nm}|^2 \right)^2 = \sum_{nm} \left(T_{nm} - |A_0 \tilde{\psi}_{nm}|^2 \right)^2 \quad (6)$$

using

$$\tilde{\psi}_{nm} = \frac{1}{N} \sum_{pq} \psi_{pq} \exp\left(-\frac{2\pi i}{N}(pn + qm)\right) \quad (7)$$

with $\psi_{pq} = S_{pq} \exp(i\phi_{pq})$, where indices p and q indicate general input plane pixels. A_0 is a free scale parameter allowing for the fact that the overall scale of the target potential is not necessarily matched to the laser power and S_{pq} allows for the possibility of an input field with a slowly varying intensity, e.g. a Gaussian input beam. The cost function gradient required by the minimisation algorithm must be calculated with respect to ϕ_{rs} , the phase value at a specific SLM pixel. For this cost function, the corresponding gradient is

$$\frac{\partial C}{\partial \phi_{rs}} = 4A_0^2 \text{Re}(i\tilde{\psi}_{rs}^* X_{rs}) \quad (8)$$

where

$$X_{rs} = \frac{1}{N} \sum_{nm} \left[\left(T_{nm} - |A_0 \tilde{\psi}_{nm}|^2 \right) \tilde{\psi}_{nm}^* \exp\left(\frac{2\pi i}{N}(rn + sm)\right) \right]. \quad (9)$$

Additional terms in the cost function can incorporate experimentally relevant output plane features. Such features could, for instance, include noise suppression at the signal region boundary to aid trap loading, or in a dynamic sequence for real-time manipulation of trapped atoms, a cost function term could be introduced to reduce intensity fluctuations between consecutive frames in the sequence. The application of additional cost function terms is illustrated here by the suppression of optical vortex formation. We find that the cost function in Eq. (6) is effective for lattice distributions, but inadequate for large continuous patterns due to the emergence of optical vortices within the trapping region during calculation. These vortices are characterised by a sudden drop in intensity coinciding with a local phase winding by a multiple of 2π , and they arise because they can be initially beneficial to cost function reduction. However, their prevention is imperative to all hologram calculation schemes. From following the evolution of our conjugate gradient minimisation, it appears the local vorticity cannot change, and so these vortices can only be removed by annihilation of oppositely charged vortex pairs, or by moving vortices to regions of low intensity [33].

Figure 2 illustrates the vortices formed within the output plane of a second-order power-law intensity distribution calculated using the cost function defined in Eq. (6) with no regional weightings applied, starting from an educated guess. Since the algorithm has identified a local minimum of the cost function, the observation of these vortices suggests that the cost function does not sufficiently penalise them. Indeed, since the vortex cores are small, they only introduce a very localised deviation from the pattern. Moreover, there is only a small change in cost function as vortices move through the pattern, and hence minimisation of this cost function does not effectively eliminate vortices once formed. In minimising this cost function, vortex removal is principally achieved by gradually shifting them towards lower-intensity regions where their cost is reduced, but this is obstructed in regions of high vortex density where phase contours can become tangled [33]. Vortex elimination is therefore only realistically achievable if their early formation is suppressed such that their numbers remain manageable. Given that the cost function can be chosen at will (within the constraints given above), our approach to eliminating vortices becomes a question of choosing a *better* cost function such that vortices do not remain frozen in the final pattern.

A cost function that penalises large localised deviations more than the simple cost function in Eq. (6) is

$$C_t = \sum_{nm} \left(T_{nm} - |A_0 \tilde{\psi}_{nm}|^2 \right)^t. \quad (10)$$




Fig. 2. Second-order power-law trapping potential calculated using the cost function in Eq. (6). (left) Two-dimensional profiles and vertical line profiles taken through the centre of the pattern for target (upper inset and cyan line) and calculated output (lower inset and red points). The colorbar applies to the insets and to all subsequent figures. The predicted output is distorted by deep optical vortices. These vortices inhibit further improvements in accuracy, resulting in a poor fit to the target intensity. The root-mean square (RMS) fractional error between target and predicted output is 26%. (right) Vortex locations on the output plane, with red pixels indicating 2π phase windings and blue pixels -2π . The black circle indicates the trapping region. While the vortex density is reduced within this region in comparison to the remainder of the output plane, 232 are established within the trapping region with an insignificant fraction removed with further iterations due to the tangled phase contours.

The higher the value of t , the higher the cost of large discrepancies relative to small, increasing the cost contribution of trapping-region vortices. Fewer vortices persist, but at the expense of trap smoothness. In practice we find that powers higher than four produce too rough an intensity distribution with insufficient vortex improvements to justify this sacrifice. The C_t gradient is

$$\frac{\partial C_t}{\partial \phi_{rs}} = 2tA_0^2 \text{Re}(i\psi_{rs}^* X_{rs}). \quad (11)$$

Alternatively, we can also specify cost functions that perform active smoothing by associating a cost with intensity variations between neighbouring pixels. For example, to apply active smoothing over the four nearest-neighbour pixels, we use the cost function

$$C_s = \sum_{nm} \left[\left(|\tilde{\psi}_{nm}|^2 - |\tilde{\psi}_{n(m-1)}|^2 \right)^2 + \left(|\tilde{\psi}_{nm}|^2 - |\tilde{\psi}_{n(m+1)}|^2 \right)^2 \right. \\ \left. + \left(|\tilde{\psi}_{nm}|^2 - |\tilde{\psi}_{(n-1)m}|^2 \right)^2 + \left(|\tilde{\psi}_{nm}|^2 - |\tilde{\psi}_{(n+1)m}|^2 \right)^2 \right]. \quad (12)$$

The gradient is calculated in the same way for all four terms in C_s . For instance, the gradient for the first term, $C_s^{(1)} = \sum_{nm} \left(|\tilde{\psi}_{nm}|^2 - |\tilde{\psi}_{n(m-1)}|^2 \right)^2$, is

$$\frac{\partial C_s^{(1)}}{\partial \phi_{rs}} = \frac{4}{N} \text{Re}(i\psi_{rs} (X_{1rs}^* - X_{2rs}^*)) \quad (13)$$

where

$$X_{1rs} = \frac{1}{N} \sum_{nm} \left[\left(|\tilde{\psi}_{nm}|^2 - |\tilde{\psi}_{n(m-1)}|^2 \right) \tilde{\psi}_{nm} \exp\left(\frac{2\pi i}{N}(rn + sm)\right) \right] \quad (14)$$

$$X_{2rs} = \frac{1}{N} \sum_{nm} \left[\left(|\tilde{\psi}_{nm}|^2 - |\tilde{\psi}_{n(m-1)}|^2 \right) \tilde{\psi}_{n(m-1)} \exp \left(\frac{2\pi i}{N} (rn + sm) \right) \right]. \quad (15)$$

Sequential combination of $C_{t=4}$ and $C_{t=2}$ establishes a vortex-free trap region with subsequent smoothing; as an alternative approach, simultaneous combination of $C_{t=2}$ and C_s terms is also successfully implemented to demand both accuracy and smoothness. Calculated outputs corresponding to these examples are illustrated in Fig. 3 for the same second-order power-law pattern shown in Fig. 2.

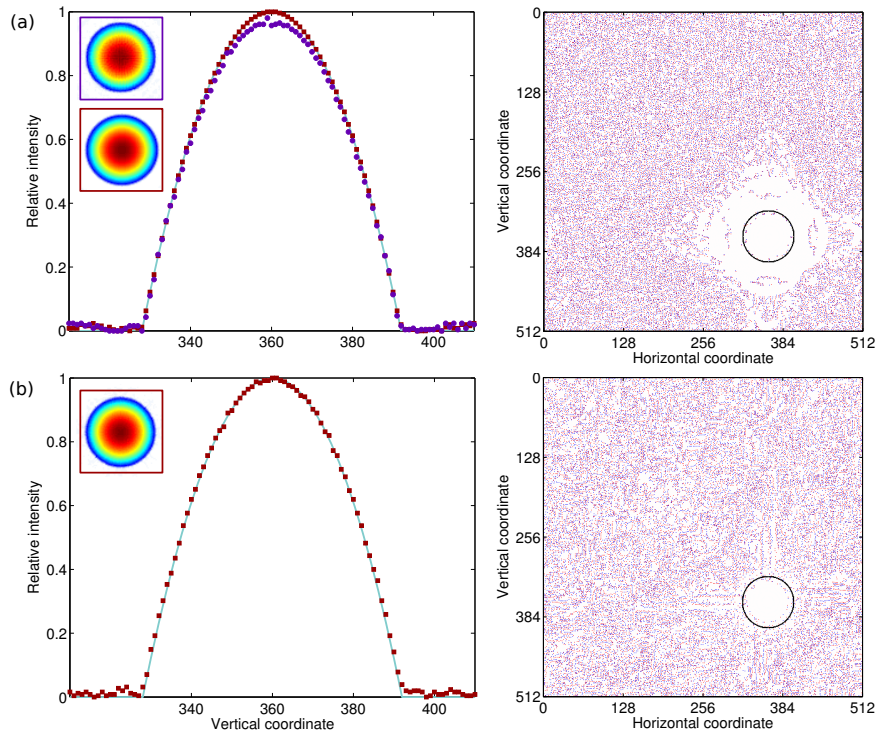


Fig. 3. Second-order power-law trapping potentials calculated with more sophisticated cost functions, starting with an educated guess phase. (left) Line profiles through pattern centre and two-dimensional profiles as insets. Calculated intensity patterns are denoted by points while the target profile is indicated with a cyan line. (right) output plane vortex map with the trapping region indicated by the black circle. (a) $C_{t=4}$ (purple, upper inset) and subsequent $C_{t=2}$ application (red, lower inset). Applying $C_{t=4}$ generates an approximate fit to the target while suppressing the initial vortex number. Subsequent smoothing using $C_{t=2}$ improves the accuracy and smoothness, and clears residual vortices from the trapping region. (b) A combination of $C_{t=2}$ and C_s achieves both direct smoothing and high reproduction accuracy.

For the sequential C_t application shown in Fig. 3(a), we initially use $C_{t=4}$ to apply coarse corrections to the calculated intensity pattern, then follow this with more refined corrections using $C_{t=2}$. The cost function on signal region pixels is given a weighting of 10 times that of noise region pixels during the $C_{t=4}$ stage; for $C_{t=2}$ application the signal region is weighted by a factor of $10^{12}(1+T)$ relative to the noise region with T the target value of a given pixel, while a linear slope over 8 pixels smooths the weightings between these two regions and discourages noise accumulation near the target intensity distribution. The calculated fractional RMS error is 1.4% after 3000 iterations of $C_{t=4}$ application and 0.07% following an additional 10000 it-

erations of $C_{t=2}$ smoothing, with 70 vortices remaining in the signal region. However, as they are confined to low-intensity regions these vortices do not degrade the trapping pattern. The efficiency of the algorithm in placing light within the trapping region is 47% prior to smoothing and 45% afterwards. However, efficiency varies widely according to weighting choice: one method of improving efficiency is to demand an accuracy across the entire output plane comparable to that of the signal region, requiring more iterations to achieve the desired trapping region accuracy.

In the combined $C_{t=2}$ and C_s approach, the signal and noise regions are not given relative weightings for the $C_{t=2}$ term, but the C_s term is weighted according to the target value for each pixel within the signal region and set to zero in the noise region. Competition between the accuracy and smoothing terms reduce accuracy as compared to the pure discrepancy power method, with a fractional RMS error of 0.43% after 30000 iterations in the example shown in Fig. 3(b). However, appropriate balancing of terms results in effective vortex suppression and sufficient prediction accuracy. 130 vortices remain in the low-intensity region of the trapping pattern, which again is less relevant for atom trapping. Furthermore, regional weightings increase the efficiency to 64%. We also find that this active smoothing method is particularly resilient to initialisation conditions, increasing the chance of success of a given iteration run.




Fig. 4. Target (i) and calculated output intensity (ii) for the square lattice (a) and stirring ring pattern (b). Both are generated from a random initial phase and have RMS errors of 0.58% and 3.0% respectively.

With MRAF, the output quality depends critically on the initial phase guess and on the initialisation parameters, which have to be carefully chosen to suppress the formation of optical vortices during the calculation process [13]. In contrast, with conjugate gradient minimisation, optical vortex suppression is achieved by a judicious cost function choice. Having determined these cost functions, the output quality is then largely insensitive to the initial phase guess, to the point that high accuracy can be achieved even with a random phase guess. For this reason, while conjugate gradient minimisation converges in more iterations than are required in MRAF, the two methods end up with a comparable computational efficiency, because with conjugate gradient minimisation it is not necessary to run the code for many different choices of initial phase patterns. Initialisation resilience would be of further benefit in dynamical sequences as it increases the chance of success of all frames from a single initialisation step.

Displayed in Fig. 4 are examples illustrating the general applicability of the method to both continuous distributions and discrete arrays, showcasing the successful elimination of high-intensity borders next to the signal region. Both patterns are generated from a random initial guess; remarkably, for the ring pattern, we find that an initially random phase pattern results in quicker convergence to the final form than an apparently educated guess. The lattice pattern is calculated using solely a $C_{t=2}$ term: smoothing is found to cause blurring of the pattern edges, and the $C_{t=2}$ term is sufficient to remove vortices from these smaller intensity features before they become established. The small size of the spots allows the vortices to escape more easily and they are therefore not frozen into the final pattern. The signal region is weighted by a factor

of 10^4 relative to the noise region, with a linearly sloped border of 4 pixels connecting these two regions sufficient to discourage noise accumulation near the signal region boundary. After 4000 iterations, the RMS signal region error is 0.58%. The stirring ring pattern, so called because it can be used to induce superfluid rotation [10], has an RMS signal region error of 3.0% after 2500 iterations. This example is calculated using a combination of $C_{t=2}$ and C_s terms, with the signal region given an overall weighting of 10 relative to the noise region with a border of 8 pixels connecting these to prevent disruptive noise accumulation, and the smoothing term given a weighting of $\frac{1}{2}(T_{max} - T_{nm})$, with T_{max} the maximum target value and T_{nm} the target values on individual pixels, in the signal region only.

4. Conclusion

Conjugate gradient minimisation of an appropriate cost function has been verified as a viable alternative to the established methods of hologram generation. By applying well-developed conjugate gradient approaches and optimised numerical libraries, the cost function minimisation approach allows careful guiding of the calculation process by choice of a sensible cost function with an analytical gradient. In particular, we show that tailoring the cost function beyond its simplest form is important and that the flexibility inherent in the cost function definition should be exploited to guide output plane features of interest. We illustrate this by directly suppressing optical vortices in the calculated intensity profiles by associating a cost either with large deviations from the target intensity or with intensity fluctuations between neighbouring pixels. Both methods successfully suppress vortices to optimise trapping potential accuracy, though active smoothing is more appropriate in potentials without sharp features. This precision guiding and the ability to tailor the weightings assigned to each pixel has also allowed us to avoid the formation of the high-intensity signal region border characteristic of the MRAF approach, while retaining the flexibility characteristic of regional definitions and the ability to accurately reproduce a wide range of intensity patterns suitable for trapping ultracold atoms. The method has also proven to be resilient to initialisation conditions, which may be of benefit in designing dynamic sequences of intensity patterns from a single initialisation step.

The method of conjugate gradient optimisation may also find applications in the fields of biological, chemical and soft condensed matter physics, where the use of phase-only SLMs is already well established [34, 35]. Further improvements could be achieved by incorporating a measured laser beam profile into the calculation process to correct for beam imperfections, and considering Helmholtz propagation of light within the model [14], while overall computation times could be reduced by performing the calculations on Graphics Processing Units [36, 37].

Acknowledgments

The authors acknowledge funding from the Leverhulme Trust Research Project Grant RPG-2013-074 and from the EPSRC grant EP/G004714/2.