



RESEARCH ARTICLE

Hyperspectral compressive wavefront sensing

Sunny Howard^{1,2}, Jannik Esslinger², Robin H. W. Wang¹, Peter Norreys^{1,3}, and Andreas Döpp^{1,2}

¹Department of Physics, Clarendon Laboratory, University of Oxford, Oxford, UK

²Centre for Advanced Laser Applications, Ludwig-Maximilians-Universität München, Garching, Germany

³John Adams Institute for Accelerator Science, Oxford, UK

(Received 27 August 2022; revised 11 October 2022; accepted 2 November 2022)

Abstract

Presented is a novel way to combine snapshot compressive imaging and lateral shearing interferometry in order to capture the spatio-spectral phase of an ultrashort laser pulse in a single shot. A deep unrolling algorithm is utilized for snapshot compressive imaging reconstruction due to its parameter efficiency and superior speed relative to other methods, potentially allowing for online reconstruction. The algorithm's regularization term is represented using a neural network with 3D convolutional layers to exploit the spatio-spectral correlations that exist in laser wavefronts. Compressed sensing is not typically applied to modulated signals, but we demonstrate its success here. Furthermore, we train a neural network to predict the wavefronts from a lateral shearing interferogram in terms of Zernike polynomials, which again increases the speed of our technique without sacrificing fidelity. This method is supported with simulation-based results. While applied to the example of lateral shearing interferometry, the methods presented here are generally applicable to a wide range of signals, including Shack–Hartmann-type sensors. The results may be of interest beyond the context of laser wavefront characterization, including within quantitative phase imaging.

Keywords: artificial neural networks; compressed sensing; high-power laser characterization; wavefront measurement

1. Introduction

Ultrashort laser pulses possess a necessarily broad spectral bandwidth^[1]. The chromatic properties of the optical elements that are used for the generation or application of such pulses can then create relations between the spatial and temporal profiles, called spatio-temporal couplings (STCs)^[2]. These phenomena can lead to a variety of effects including, for example, the broadening of a focused laser pulse either spatially or temporally, thereby reducing its peak intensity^[3]. Deliberately introduced STCs can lead to exotic light pulses that behave very differently from 'normal' pulses. An example of this is the so-called flying focus^[4] with its potential application in laser-driven wakefield accelerators^[5] and orbital angular momentum beams^[6]. Universally, the expansion in the applications of ultrafast laser pulses has exacerbated the need for a robust way to measure their properties.

To resolve STCs, one must gain wavefront information over the 3D hypercube (x, y, t) or equivalently its spatio-spectral analogue (x, y, ω) . Due to the limitation that array sensors (such as complementary metal–oxide–semiconductor (CMOS) cameras) capture information in a maximum of two dimensions, the majority of current techniques resort to scanning over one or two dimensions, whether it is a spatial^[7,8], spectral^[9,10] or temporal^[11] scan. Such techniques are time consuming and are blind to shot-to-shot variations and drift of the laser. While there exist some methods that are single shot^[12] – that is, those that capture the hypercube in one shot – these currently lack resolution and spectral range and are cumbersome to implement.

Inspired by recent progress in machine-learning-based laser science^[13], here we present the concept for a single-shot method that utilizes compressed sensing (CS) to resolve the wavefront in both the spectral and spatial domains. The paper is structured as follows. In Section 2 we will discuss the wavefront sensor and in Section 3 we introduce snapshot compressive imaging (SCI) as a way to expand the wavefront sensor to measuring multiple colours at once. Our implementation is based on deep unrolling, which yields high

Correspondence to: Andreas Döpp, Centre for Advanced Laser Applications, Ludwig-Maximilians-Universität München, Am Coulombwall 1, 85748 Garching, Germany. Email: a.doepf@lmu.de

performance in both reconstruction fidelity and speed, as required for use as a real-time diagnostic. Section 4 provides a thorough description of all neural network architectures used, and Section 5 contains a description of how training data was generated, before Section 6 displays the results of the proposed method.

2. Wavefront sensing

The wavefront sensor that was simulated in this example was a quadriwave lateral shearing interferometer (QWLSI), which is known for its high resolution and reconstruction fidelity. Nonetheless, our method can in general be applied to any kind of wavefront retrieval technique, including the popular Shack–Hartmann sensor or multi-plane techniques, such as Gerchberg–Saxton phase retrieval.

A lateral shearing interferometer (LSI) measures the spatially varying phase of a light beam, and was first applied to the measurement of ultrashort laser pulses in the late 1990s^[14]. The LSI works by creating multiple copies of the laser pulse and shearing them laterally relative to each other before their interference pattern is captured on a sensor. Due to the shear, information about the spatial gradient of the wavefront is encoded in the interferogram. This can then be extracted using Fourier filtering^[15] and stitched together to form the wavefront via methods such as modal reconstruction^[16] or Fourier integration^[17]. The most popular implementation is the aforementioned QWLSI. By generating and shearing four (2×2) copies of the pulse under investigation, this setup also enables the extraction of two separate pairs of orthogonal gradients, meaning two distinct estimates for the wavefront can be found, providing error estimation. This property is highly desirable for a sensor based on CS, because inevitable noise in the measurement can corrupt the wavefront with reconstruction artifacts. This

would, for instance, be the case in a two-plane Gerchberg–Saxton algorithm. In contrast, redundancy of phase information in the QWLSI provides direct validation and, thus, makes the wavefront retrieval much more resilient to noise. A sketch illustrating the concept of the QWLSI is shown in the red box of Figure 1.

2.1. QWLSI simulation

A physical implementation of the QWLSI usually consists of a phase grating with ‘pixels’ of alternating phase arranged in a checkerboard pattern^[18], which leads to dominant diffraction in 2×2 copies of the beam. Instead of simulating this process, we consider an idealized setup where we analytically generate the four copies. We begin with creating the pulse of interest by defining a spatial-spectral intensity, I_0 , and phase, ϕ_0 :

$$E_0(x, y, z = 0, \omega) = \sqrt{I_0(x, y, z = 0, \omega)} e^{i\phi_0(x, y, z = 0, \omega)}. \quad (1)$$

The pulse is copied four times, and each copy’s field is propagated to the detector plane according to the following rules. Note that for brevity, when the z index is not stated, $z = 0$ and the ω index will be suppressed from the electric field.

2.1.1. Propagation

Considering the j th copy has travelled a distance Δz , its poloidal angle is θ_j and its azimuthal angle is ζ , then one finds its displacement in the x and y directions to be as follows:

$$\begin{aligned} \Delta x_j(\omega) &= \Delta z \sin(\theta_j) \sin(\zeta(\omega)), \\ \Delta y_j(\omega) &= \Delta z \cos(\theta_j) \sin(\zeta(\omega)). \end{aligned}$$




Figure 1. Schematic of the experimental setup that was simulated. The pulse first travels through a quadriwave lateral shearing interferometer, yielding a hypercube of interferograms, a slice of which is shown in the green box. The hypercube is then passed through a CASSI setup. This consists of a random mask and a relay system encompassing a prism, before the coded shot is captured with the camera. This diagram is not to scale.

In a QWLSI, the poloidal angles are $\theta_j \in \{0, \frac{\pi}{2}, \pi, \frac{3\pi}{2}\}$. In Figure 1, one identifies the azimuthal angle, ζ , as that between the white copy lines and the central yellow line. This is related to both the pitch of the grating Λ and the wavelength λ by the following:

$$\zeta(\omega) = \arcsin\left(2\pi \frac{\lambda}{\Lambda}\right).$$

The resulting electric field of the copy is as follows:

$$E_j(x, y, \Delta z) = \frac{1}{4} \sqrt{I_0(x - \Delta x_j, y - \Delta y_j)} e^{i\phi_0(x - \Delta x_j, y - \Delta y_j)}. \tag{2}$$

2.1.2. Tilt

As diffraction occurs at an angle ζ , the grating imparts a tilt onto the copy. This translates to an additional phase shift dependent on both the spectral and spatial domains:

$$\Delta\phi_j(x, y, \omega) = k(\omega) (x \cos(\theta_j) + y \sin(\theta_j)) \sin(\zeta),$$

where $k = \frac{2\pi}{\lambda}$ is the wavevector of the pulse. This tilt is crucial in reconstruction as it provides a high-frequency modulation that separates the gradients in Fourier space.

Combining these two effects and summing over copies, we obtain the final changes to the field:

$$E(x, y, \Delta z) = \frac{1}{4} \sum_{j=1}^4 \sqrt{I_0(x - \Delta x_j, y - \Delta y_j)} \cdot e^{i(\phi(x - \Delta x_j, y - \Delta y_j) + \Delta\phi_j(x, y))}. \tag{3}$$

At the Talbot self-imaging plane, $\Delta z = 2\Lambda^2/\lambda$, one has a hypercube of interferograms. An example of a one frequency channel slice is shown in the green box of Figure 1.

In other applications one would collapse the cube onto a sensor at this point; however, this would eliminate the chance of retrieving the spectrally resolved phase. Instead, as discussed in Section 3, we use SCI to aid in the capturing of the cube.

2.2. Wavefront reconstruction

Once the interferogram is captured, one must extract the wavefront. As previously mentioned, current reconstruction methods usually involve multiple steps, that is, extracting the gradients, integrating and stitching them together. This can be a time-consuming process, especially in a hyperspectral setting where the reconstruction has to be done for every channel. To address this problem, we present a deep learning approach to wavefront reconstruction for the LSI. While similar work has been done in the context of Shack–Hartmann sensors^[19,20], this is the first application of deep learning to LSI reconstruction, to the best of our knowledge. The network that was used will be discussed in Section 4.

3. Snapshot compressive imaging

CS describes the highly efficient acquisition of a sparse signal from fewer samples than would classically be required according to the Nyquist theorem by utilizing optimization methods to solve underdetermined equations. SCI is an example of CS, capturing 3D data on a 2D sensor in a single shot.

Fundamentally, there are two requirements to be fulfilled for CS to work. Firstly, the signal must be sparse on some basis and, secondly, the signal must be sampled on the basis that it is incoherent with respect to the sparse basis^[21]. The first condition was hypothesized to be satisfied given the fact that laser wavefronts are known to be well-expressed with a few coefficients of the Zernike basis. When one does not have prior knowledge about which basis the signal is sparse on, the second condition is often solved by performing random sampling. Whilst being trivial for 2D data, in the context of SCI it is challenging, as the 3D hypercube must be randomly sampled onto a 2D sensor. To do so, nearly all research in this area uses hardware based on the coded aperture snapshot compressive imaging (CASSI) system^[22,23].

3.1. CASSI

The hypercube is first imaged onto a coded aperture. This is a binary random mask with each pixel transmitting either 100% or 0% of the light. The cube is then also passed through some dispersive media, for example, a prism or grating, before being captured by a sensor, resulting in what is known as the coded shot. The effect of this optical system is that when the hypercube reaches the detector plane, each spectral channel is encoded with a different coded aperture, thereby approximating random sampling across the whole cube. It is then possible for a reconstruction algorithm to retrieve the cube. A diagram of a CASSI system is shown in the yellow box in Figure 1, with an example of a coded shot for an interferogram hypercube shown on the far left of Figure 3. The setup can easily be simulated by multiplying the cube by the mask, then shifting the channels of the cube according to the amount of (angular) dispersion imparted onto them, and finally summing over the spectral axis.

Mathematically, the CASSI system discussed above is summarized into a matrix Φ , which operates on \mathbf{m} , a vectorized representation of the hypercube, to give \mathbf{n} , a vectorized coded shot:

$$\mathbf{n} = \Phi \mathbf{m}. \tag{4}$$

In order to reconstruct \mathbf{m} , one can solve the following:

$$\tilde{\mathbf{m}} = \underset{\mathbf{m}}{\operatorname{argmin}} \left[\underbrace{\|\mathbf{n} - \Phi \mathbf{m}\|^2}_{\text{data term}} + \underbrace{\eta \mathcal{R}(\mathbf{m}, \psi)}_{\text{regularizer}} \right]. \tag{5}$$

The first term on the right-hand side is labelled the data term, and enforces that the hypercube must match the coded shot when captured. This alone would be an underdetermined system, so a regularization term, parameterized by ψ , is added that restricts the solution space and selects the correct hypercube.

Most methods that have been developed to solve this non-convex equation can be sorted into two classes: iterative algorithms or end-to-end neural networks. The former offers good generalization but lacks abstraction capability and is slow, whilst deep nets are fast and have been shown to learn almost any function, but can be prone to overfitting^[24]. A middle ground that offers state-of-the-art performance is deep unrolling.

3.2. Deep unrolling

While an end-to-end neural net would attempt to solve Equation (5) directly, if it were possible to split the equation, the data term can actually be solved analytically. This is desirable as it alleviates the abstraction needed to be done by the network, resulting in greater generalization and better parameter efficiency^[25]. To perform such a separation, half quadratic splitting is employed. Firstly, an auxiliary variable \mathbf{p} is substituted into the regularization term, with Equation (6) being equivalent to Equation (5). Then, the constraint is relaxed and replaced by a quadratic loss term:

$$\hat{\mathbf{m}}, \hat{\mathbf{p}} = \operatorname{argmin}_{\mathbf{m}, \mathbf{p}} [|\mathbf{n} - \Phi \mathbf{m}|^2 + \eta R(\mathbf{p})] \text{ s.t. } \mathbf{m} = \mathbf{p}, \quad (6)$$

$$\approx \operatorname{argmin}_{\mathbf{m}, \mathbf{p}} [|\mathbf{n} - \Phi \mathbf{m}|^2 + \eta R(\mathbf{p}) + \beta |\mathbf{m} - \mathbf{p}|^2]. \quad (7)$$

Here, β is a variable that controls the strength of the constraint. High values of β will strongly enforce $\mathbf{m} = \mathbf{p}$ and approximate the subject-to statement.

The benefit of this problem formulation is that it is then possible to split Equation (7) into two minimization sub-problems in \mathbf{m} and \mathbf{p} , and effectively separate the data term from the regularization term. When minimized iteratively, the following sub-problems can approximate Equation (7):

$$\hat{\mathbf{p}}^{k+1} = \operatorname{argmin}_{\mathbf{p}} [\beta |\mathbf{p} - \mathbf{m}^k|^2 + \eta R(\mathbf{p})] \sim \mathcal{S}(\mathbf{m}^k), \quad (8)$$

$$\hat{\mathbf{m}}^{k+1} = \operatorname{argmin}_{\mathbf{m}} [|\mathbf{n} - \Phi \mathbf{m}|^2 + \beta |\mathbf{p}^{k+1} - \mathbf{m}|^2]. \quad (9)$$

Equation (9) is a convex equation and can be solved via a conjugate gradient algorithm, which provides better stability than solving analytically. On the right-hand side of Equation (8), \mathcal{S} represents that a neural network will be used to solve the equation.

The deep unrolling process is shown in Figure 2(b) panel (i). Firstly $\mathbf{m}^{(0)}$ is initialized: $\mathbf{m}^{(0)} = \Phi^T \mathbf{n}$. Then the two equations are solved for a fixed number of iterations, with

the same architecture neural net being used to represent Equation (8) in each iteration. However, the network has its own set of weights for each iteration, hence the unrolling of the algorithm. The architecture of the network will be discussed in Section 4.

4. Network architecture

This section contains the architectures of the neural networks that were used. They will be discussed in the order they are used in the reconstruction process, which is displayed in the flow chart of Figure 2(a). Firstly, the deep unrolling algorithm performs reconstruction of the interferogram hypercube from the coded shot, and secondly another network, Xception-LSI, reconstructs the spatial-spectral wavefront from the hypercube.

4.1. Deep unrolling regularizer

As previously discussed, the neural network, \mathcal{S} , represents a regularization term. This means one can exploit prior knowledge about the data to choose a suitable architecture. As will be discussed in the following section, STCs can be described by a correlation between Zernike polynomial coefficients and the wavelength. Accordingly, there will likely be a strong similarity in spot positions for neighbouring spectral channels. Due to this, an architecture with 3D convolutions was developed, which can exploit these relations. Inspired by recent work in video SCI^[26], a simplified ResUNet architecture was chosen^[27], with the standard 2D convolutions replaced with 3D ones. We used 10 iterations for our model, as it has been found that adding more than this produces negligible performance gains^[28]. A diagram of the network is displayed in Figure 2(b) panel (i).

4.2. Xception-LSI

A wavefront retrieval network was developed that takes a single spectral channel QWLSI interferogram and predicts the spatial wavefront in terms of Zernike coefficients. The network is based on the Xception network^[29], but as the original 71-layer network is designed for classification, some changes were made to adapt Xception to our application. Firstly, the final two layers were removed. A max pool layer and a convolutional layer were added to shrink the output in the spatial and spectral dimensions, respectively. Dropout was applied before using three dense layers with 1000, 500 and 100 nodes using the ReLU activation function^[30]. The output layer consists of 15 nodes with linear activation, corresponding to the number of Zernike coefficients to predict. We name the network Xception-LSI, and it can be seen in Figure 2(c) panel (i).

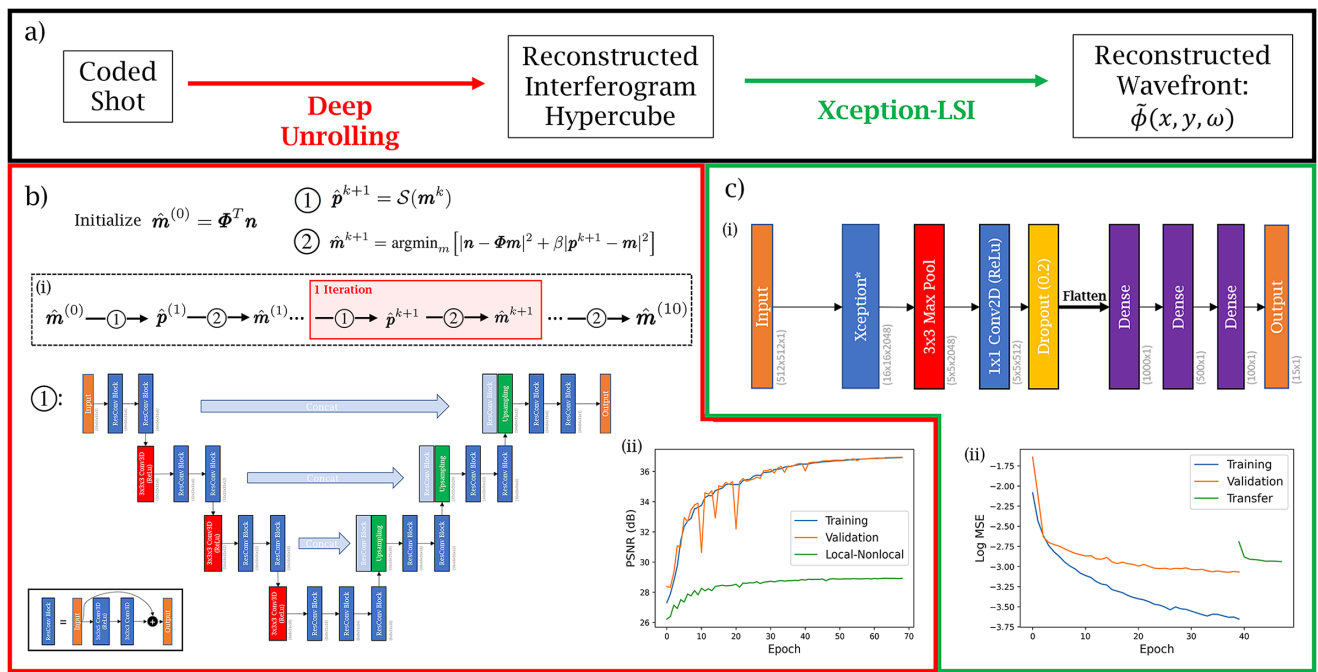


Figure 2. A diagram showing the full reconstruction process of the wavefront from the coded shot. (a) A flow chart of the reconstruction process. (b) (i) The deep unrolling process, where sub-problems ① and ② are solved recursively for 10 iterations. Also shown is the neural network structure used to represent $\mathcal{S}(m^k)$. (ii) The training curve for the deep unrolling algorithm. Plotted is the training and validation PSNR for the 3D ResUNet prior that was used, as well as the validation score for a local–nonlocal prior. Here is demonstrated the superior power of 3D convolutions in this setting. (c) (i) The network design for the Xception-LSI network. The Xception* block represents that the last two layers were stripped from the conventional Xception network. (c) (ii) The training curve for Xception-LSI for training and validation sets, with the loss shown in log mean squared error. Also plotted is the validation loss when further training the model on the deep unrolling reconstruction of the data (transfer).

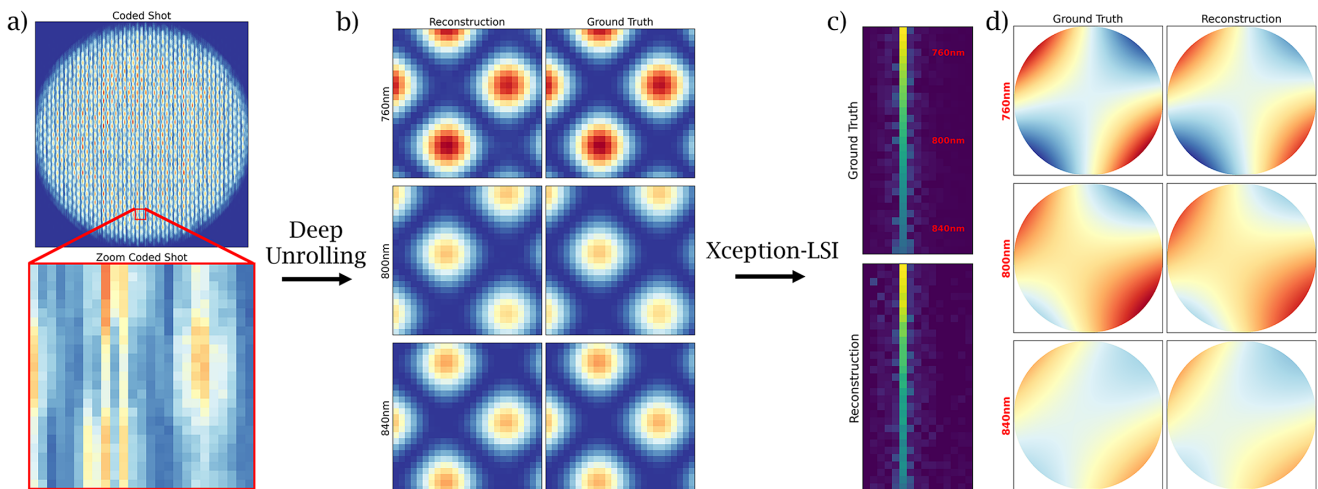


Figure 3. Example results of the reconstruction process. (a) An example of the coded shot, along with a zoomed section. (b) Deep unrolling reconstruction of the interferogram hypercube in the same zoomed section at different wavelength slices. (c) The Xception-LSI reconstruction of the spatio-spectral wavefront displayed in terms of Zernike coefficients, where the x -axis of each plot is the Zernike function, the y -axis is the wavelength and the colour represents the value of the coefficient. (d) The spatial wavefront resulting from a Zernike basis expansion of the coefficients in (c) at the labelled spectral channels.

5. Training data generation

To represent the initial pulse, a total of 300 cubes were generated with dimensions $(n_x \times n_y \times n_\omega) = (512 \times 512 \times 31)$. The data was randomly split at a ratio of 4 : 1 : 1 into training, validation and test sets, respectively. The wavelength range considered was 750–850 nm, representing a broadband

Ti:sapphire laser, giving $\Delta\lambda \approx 3.23$ nm. For each cube, the wavefront for each channel was first initialized to a randomly weighted sum of 15 Zernike basis functions. Then, to simulate an STC, one Zernike function was chosen and was made to vary either linearly or quadratically with frequency. Indeed, common STCs, such as pulse front tilt and pulse front curvature, can be represented in this way^[1]. The mean

amplitude of this coefficient was also made to be higher. This choice of Zernike coefficients is arbitrary, but allows for a demonstration that the method can identify all Zernike basis functions. The intensity of slices of the cube was set to an image taken of a real laser intensity.

Each cube was then processed according to Figure 1. Firstly, it was passed through the QWLSI simulation (see Section 2.1), yielding a hypercube of interferograms – these are the training labels for the deep unrolling algorithm. This hypercube was then passed through the SCI simulation, yielding a coded shot – the training data. The wavefront was reconstructed via the process in Figure 2(a). The interferogram hypercube was reconstructed via deep unrolling, before being passed into the Xception-LSI network to predict the spectral Zernike coefficients.

The pitch of the LSI was set to $\Lambda = 80 \mu\text{m}$, and the dispersion of the prism, measured at the camera plane, was set to 1 pixel per channel (each channel having a width of 3.23 nm).

Before being passed through the deep unrolling network, the cubes and coded shots were split spatially into 64×64 oblique parallelepiped patches, allowing for a one-to-one reconstruction between the input and output^[28]. The initial learning rate was set to 0.01 and decayed by 10% every five epochs. The total number of epochs was 70, and the batch size was 8.

The Xception-LSI network was fed individual channels of the ground truth interferogram hypercubes and predicted Zernike coefficients. Normal random noise ($\mathcal{N}(\mu = 0, \sigma = 0.1)$) was applied to the input, to make the model robust to noise produced by the SCI reconstruction. The initial learning rate was set to 10^{-5} and decayed by 10% every five epochs. The total number of epochs was 40, and the batch size was 16. Once trained on the ground truth hypercubes, the model was trained on interferogram hypercubes that had been reconstructed by deep unrolling, for the further eight epochs. The aim of this transfer learning was to allow the network to account for any systematic noise in the SCI reconstruction, resulting in a more accurate wavefront reconstruction.

6. Results and discussion

6.1. Snapshot compressive imaging

Crucial to this method's success is the SCI reconstruction of the hypercube of interferograms. As can be seen from the green box of Figure 1, the image slices are modulated and appear as spot patterns. As a result, the images do not exhibit the same sparsity in, for example, the wavelet domain, as most natural images used in SCI research do. Because of this, there was uncertainty in whether it would be possible to recover the cube.

Here it is demonstrated that it is indeed possible to reconstruct such modulated signals with SCI. The training curve can be seen in Figure 2(b) panel (ii). Also plotted is the validation loss when a local–nonlocal prior^[31], which is state-of-the-art for natural images, was used. One sees that when both architectures were used with 10 iterations of unrolling, the 3D convolutional model achieved a far superior peak signal-to-noise ratio (PSNR) of 36 compared to 29. Furthermore, it contains approximately 45% fewer parameters.

6.2. QWLSI

In order to reconstruct the wavefront for a full hypercube, each spectral channel is fed through the network sequentially. After training, the final mean squared error on the ground truth test set was 6.80×10^{-4} . Figure 2(c) panel (ii) displays the training curve with the training, validation and transfer loss curves. The additional transfer learning proves to be extremely effective in reducing the error of the wavefront predictions when working with reconstructed interferogram hypercubes. The final mean squared error on the reconstructed test set was 9.18×10^{-4} .

6.3. Hyperspectral compressive wavefront sensing

An example of the full reconstruction process, from coded shot to spatial-spectral wavefront, is displayed in Figure 3. It is apparent that the deep unrolling network was able to accurately reconstruct the interferogram hypercube, and the Xception-LSI network was able to reconstruct the wavefront.

7. Summary and outlook

In this report we have demonstrated the possibility of combining a wavefront sensor with SCI in order to achieve a single-shot measurement of the spatial-spectral phase. Crucially, it has been shown that SCI has the ability to reconstruct modulated signals, such as those produced by a QWLSI.

A natural progression to this study is to realize the results in an experimental setting, where challenges arise from the more complicated dispersion, transfer functions and noise. Other further work could include extending the deep learning LSI analysis to the hyperspectral setting. By passing the network a hypercube of interferograms, rather than individual slices, it may be possible to exploit spectral correlations in order to improve accuracy and detect STCs more easily. Also, work can be done on testing the model with a more varied set of Zernike polynomials. Finally, there has been recent interest in the possibility of spreading phase contrast imaging to a hyperspectral setting. However, current methods take many seconds to capture a hypercube of phase^[32]. The proposed method would be able to collect information

with higher spectral resolution in a single shot, allowing for dynamic events to be recorded hyperspectrally.

Acknowledgements

We would like to acknowledge the useful discussions with Dr. Ramy Aboushelbaya and the rest of Professor Peter Norreys' group. This work was supported by the Independent Junior Research Group 'Characterization and control of high-intensity laser pulses for particle acceleration', DFG Project No. 453619281. We would also like to acknowledge UKRI-STFC grant ST/V001655/1.

References

1. S. Jolly, O. Gobert, and F. Quere, *J. Opt.* **22**, 103501 (2020).
2. A. Jeandet, S. W. Jolly, A. Borot, B. Bussière, P. Dumont, J. Gautier, O. Gobert, J.-P. Goddet, A. Gonsalves, A. Irman, W. P. Leemans, R. Lopez-Martens, G. Mennerat, K. Nakamura, M. Ouillé, G. Pariente, M. Pittman, T. Püschel, F. Sanson, F. Sylla, C. Thauray, K. Zeil, and F. Quéré, *Opt. Express* **30**, 3262 (2022).
3. C. Bourassin-Bouchet, M. Stephens, S. de Rossi, F. Delmotte, and P. Chavel, *Opt. Express* **19**, 17357 (2011).
4. D. H. Froula, J. P. Palastro, D. Turnbull, A. Davies, L. Nguyen, A. Howard, D. Ramsey, P. Franke, S.-W. Bahk, I. A. Begishev, R. Boni, J. Bromage, S. Bucht, R. K. Follett, D. Haberberger, G. W. Jenkins, J. Katz, T. J. Kessler, J. L. Shaw, and J. Vieira, *Phys. Plasmas* **26**, 032109 (2019).
5. C. Caizergues, S. Smartsev, V. Malka, and C. Thauray, *Nat. Photonics* **14**, 475 (2020).
6. R. Aboushelbaya, "Orbital angular momentum in high-intensity laser interactions", PhD. Thesis (University of Oxford, 2021).
7. P. Bowlan, P. Gabolde, A. Shreenath, K. McGresham, R. Trebino, and S. Akturk, *Opt. Express* **14**, 11892 (2006).
8. M. López-Ripa, I. J. Sola, and B. Alonso, *Photon. Res.* **10**, 922 (2022).
9. S. L. Cousin, J. M. Bueno, N. Forget, D. R. Austin, and J. Biegert, *Opt. Lett.* **37**, 3291 (2012).
10. N. Weiße, J. Esslinger, S. Howard, F. M. Foerster, F. Haberstroh, L. Doyle, P. Norreys, J. Schreiber, S. Karsch, and A. Döpp, [arXiv:2303.01360](https://arxiv.org/abs/2303.01360) (2023).
11. G. Pariente, V. Gallet, A. Borot, O. Gobert, and F. Quéré, *Nat. Photonics* **10**, 547 (2016).
12. P. Gabolde and R. Trebino, *Opt. Express* **14**, 11460 (2006).
13. A. Döpp, C. Eberle, S. Howard, F. Irshad, J. Lin, and M. Streeter, [arXiv:2212.00026](https://arxiv.org/abs/2212.00026) (2022).
14. J.-C. Chanteloup, F. Druon, M. Nantel, A. Maksimchuk, and G. Mourou, *Opt. Lett.* **23**, 621 (1998).
15. M. Takeda, H. Ina, and S. Kobayashi, *J. Opt. Soc. Am.* **72**, 156 (1982).
16. F. Dai, F. Tang, X. Wang, O. Sasaki, and P. Feng, *Appl. Opt.* **51**, 5028 (2012).
17. S. Velghe, J. Primot, N. Guérineau, M. Cohen, and B. Wattellier, *Opt. Lett.* **30**, 245 (2005).
18. S. Velghe, J. Primot, N. Guérineau, R. Haïdar, S. Demoustier, M. Cohen, and B. Wattellier, *Proc. SPIE* **6292**, 62920E (2006).
19. L. Hu, S. Hu, W. Gong, and K. Si, *Opt. Express* **27**, 33504 (2019).
20. L. Hu, S. Hu, W. Gong, and K. Si, *Opt. Lett.* **45**, 3741 (2020).
21. E. J. Candes, *Comptes Rendus Math.* **346**, 589 (2008).
22. M. E. Gehm, R. John, D. J. Brady, R. M. Willett, and T. J. Schulz, *Opt. Express* **15**, 14013 (2007).
23. A. Wagadarikar, R. John, R. Willett, and D. Brady, *Appl. Opt.* **47**, B44 (2008).
24. X. Yuan, D. J. Brady, and A. K. Katsaggelos, *IEEE Signal Process. Mag.* **38**, 65 (2021).
25. V. Monga, Y. Li, and Y. C. Eldar, *IEEE Signal Process. Mag.* **38**, 18 (2021).
26. Z. Wu, J. Zhang, and C. Mou, [arXiv:2109.06548](https://arxiv.org/abs/2109.06548) (2021).
27. Z. Zhang, Q. Liu, and Y. Wang, *IEEE Geosci. Remote. Sens. Lett.* **15**, 749 (2018).
28. L. Wang, C. Sun, Y. Fu, M. H. Kim, and H. Huang, in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2019), p. 8024.
29. F. Chollet in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017), p. 1800.
30. A. F. Agarap, [arXiv:1803.08375](https://arxiv.org/abs/1803.08375) (2018).
31. L. Wang, C. Sun, M. Zhang, Y. Fu, and H. Huang, in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2020), p. 1658.
32. C. Ba, J.-M. Tsang, and J. Mertz, *Opt. Lett.* **43**, 2058 (2018).