

Phase identification despite amplitude variation in coherent beam combination using deep learning

JAMES A. GRANT-JACOB,^{1,\$} BEN MILLS,^{1,\$,*} AND MICHALIS N. ZERVAS¹

¹*Optoelectronics Research Centre, University of Southampton, SO17 1BJ, UK*

^{\$}*Authors contributed equally*

^{*}*bm602@orc.soton.ac.uk*

Abstract: Coherent beam combination offers the potential for surpassing the power limit of a single fibre laser, as well as achieving agile far-field beam-shaping. However, the spatial beam profile of the combined beam is significantly dependent on the phase of each fibre. Recent results have shown that deep learning can be used to extract phase information from a far-field intensity profile, hence unlocking the potential for real-time control. However, the far-field intensity profile is also dependent on the amplitude of each fibre, and therefore phase identification may also need to occur whilst the fibre amplitudes are not equal. Here, it is shown that a neural network trained to identify phase when all fibres have equal amplitudes can also identify phase values when the amplitudes are not equal, without requiring additional training data.

1. Introduction

Fibre lasers have transformed manufacturing, as they offer a highly efficient method of directed energy transfer [1-4]. However, the maximum power that can be generated by a single fibre laser is limited, due to ever more critical nonlinear effects as the fibre laser power is increased, most notably stimulated Raman scattering and transverse mode instability [5]. A well-known solution is through the coherent combination of multiple fibre lasers. Whilst there are many methods for combination [6, 7], here we focus on the tiled array approach, where a hexagonal close packed array of fibre lasers is combined at a single plane, and the resultant laser light is focussed onto a work piece. Whilst the coherent combination approach bypasses the single fibre power limit, this technique uncovers a different challenge, namely the requirement for precise control of the phase of each fibre, since the combined spatial intensity profile is hugely dependent on the phase of each fibre. Whilst the optimal case in such an arrangement is generally where all fibres have the same phase value, as this generally results in constructive interference with a single strong interference peak, there are also many applications for shaped beams [8-10]. Regardless of the final application, or the preferred spatial intensity profile, a fundamental challenge in this field is the identification of the phase value of each fibre. Further to this, this phase identification must be achieved in real-time, as the phases of each fibre can change continuously during practical operation.

Deep learning is a machine learning technique that uses artificial neural networks with multiple layers to automatically learn and extract hierarchical representations of data [11-15]. These learned representations enable the network to make predictions and decisions on data that was not used during the training process. Deep learning models are capable of automatically learning complex features from data, without being explicitly programmed, making them highly versatile and applicable to a wide range of tasks [16-18]. Since the first practical demonstration of the convolutional neural network [19], neural networks have dominated the field of object identification. The network architecture of a convolutional neural network is designed to identify objects in images through incremental levels of abstraction of spatial information as the image data progresses through the network layers. The result is a

powerful tool that can be applied in a huge variety of identification tasks, for example the identification of the tens of thousands of categories of objects in the ImageNet database [20]. More recently, the concept of adversarial training of convolutional neural networks was introduced, where two networks work in parallel to enhance the image generation capability [21]. This network was subsequently modified to enable image-to-image transformation. Known as a conditional generative adversarial network (cGAN) [22], this network has an extraordinary capability to transform one style of image into another style, and has seen many applications within the field of photonics, including in holography and lensless imaging [23-25], microscope resolution enhancement [26-28], and modelling of femtosecond laser machining [29-31].

Deep learning has also been applied to solving some of the challenges associated with coherent beam combination. Key breakthroughs in the field include experimental control of 107 beams [32], and a 7kW combination [33] via stochastic gradient descent [34, 35] by maximising the amount of laser light transmitted through an aperture. Whilst this approach can achieve a high intensity in a central interference peak, random noise fluctuations are generally needed (i.e., stochastic) for finding an optimal solution. In addition, this approach is also limited to the chosen aperture size and shape, and hence is less appropriate for bespoke beam shaping applications. To achieve this, a camera-based observation of the combined spatial intensity is generally needed, along with a direct technique for identification (and then optimisation) of the phase of each fibre. The challenge then becomes the well-known phase-retrieval problem [36], where phase information is extracted from an intensity pattern. In the field of coherent beam combination, there have been many suggested designs to allow additional information to be captured to eliminate the non-uniqueness of the problem [37], such as interference with a reference beam [38], and use of diffractive elements [39] to assist in the phase identification and enable beam shaping [40]. Reinforcement learning [41], which is a technique for enabling a neural network to learn through exploration whilst being given rewards, has also been applied to this field [42-46].

Recently, the authors demonstrated the identification of phase information from a far-field intensity pattern using deep learning, directly from the focal intensity pattern [47]. This result presented the potential for using a neural network for the real-time identification of phases for bespoke beam shaping when using a 19-fibre hexagonal tiled array. The neural network was trained to transform a simulated far-field intensity pattern into an image corresponding to the phases of each fibre. In other words, the network could identify the phases of all fibres directly from the intensity pattern. However, in this previous work, the amplitudes of the fibres were equal, and hence the intensity patterns always corresponded to fibres with equal amplitudes. In general, such an intensity pattern is dependent on both the phases and the amplitudes of the fibres, and hence an assumption might be that the neural network would need additional training data to identify the phase values when the amplitudes values are not equal. However, as shown here, an important result is that a neural network trained to identify phase values from intensity patterns corresponding to random phases and equal amplitudes can indeed identify phase values from intensity patterns where the amplitudes are also randomly chosen.

This observation could lead to several practical benefits for the application of neural networks for control of an experimental coherent beam combination system. Firstly, this result reduces the complexity of the task of training the neural network. Rather than having to train a neural network to identify fibre phase values for varying fibre amplitude values, the network need only be trained on varying phase values with constant amplitude values. In this case, reduction in the complexity of the task for the network to learn could allow for easier implementation on microcomputers [48, 49], which could further reduce the cost of industrial implementation. Secondly, even though the network has not encountered variation in the fibre amplitudes during training, on application to a physical system, if one of the fibre sources stops emitting light, or becomes misaligned, then the neural network could still identify the phase values of all other fibres. Thirdly, as the amplitude of each fibre could realistically be measured

in real-time with a power meter and beam pickoff, then the addition of the trained network could potentially allow both the phase and amplitude of each fibre to be monitored in real-time. This combination of measurement techniques could therefore be used to support bespoke beam shaping where the amplitudes of each fibre are also deliberately modified to unlock additional beam shaping capabilities, beyond that which could be achieved with control of phase only. In such a case, despite the amplitudes varying, the network would still be able to identify the phase values of each fibre.

In general, the outcome presented in this work is particularly important for an industrial perspective, as a considerable challenge in deep learning is the requirement for additional training data as the complexity of a challenge increases, and hence there is much interest in finding efficient applications, such as the approach presented here.

2. The interrelationship of phase and amplitude in interference patterns

The far-field intensity pattern from an array of light sources is generally dependent on the phase and the amplitude of each light source. In this work, this interrelationship was explored using a simulation of seven light sources in a hexagonal close packed arrangement, representing fibre lasers, which were then propagated to a focal plane, where the intensity was calculated. Each simulated fibre had both a phase value between $-\pi$ and $+\pi$, and an amplitude value between 0 and 1, where this amplitude value represents the maximum value (i.e., the centre) of the simulated Gaussian amplitude profile for each fibre. The simulated fibres were organised in a hexagonal close packed arrangement with a fibre radius of $500\text{ }\mu\text{m}$, a wavelength of $1\text{ }\mu\text{m}$, and where the generated light was focussed through a simulated lens and propagated towards the focal plane of 25 cm . The Gaussian amplitude profile had a radius of 0.8 of the fibre radius and had zero amplitude outside the spatial extent of each simulated fibre. Figure 1 shows three examples of phase values (shown on the left as a single column) and three examples of amplitude values (shown on the top as a single row). It is important to note that, for ease of viewing, the amplitude plots show the maximum value of the Gaussian amplitude profile used in the simulations, rather than the Gaussian spatial distribution. The figure shows the calculated far-field intensity profiles for each of the nine combinations of the phase values and the amplitude values. Of interest here, is that the intensity pattern has a clear dependence on both the phase and amplitude. In the case of a set of equal amplitude values, the phase strongly changes the pattern. However, those same phase values result in a less striking change in the intensity pattern for non-equal amplitude values.

Investigating further, figure 2a) shows firstly the combination of a set of phase values with a set of equal amplitude values, and secondly the combination of the same set of phase values with 100 randomly chosen sets of amplitudes. As shown in the figure, the average of all 100 intensity patterns is comparable to the intensity pattern from a set of equal amplitude values, with metrics for the intersection over union and image subtraction shown in parts b) and c) respectively. To calculate the intersection over union, the two images were thresholded to produce binary maps, and hence the value is provided for all possible threshold levels. This figure visually demonstrates that whilst both the phase and the amplitude affect the intensity pattern, their effects are distinct.

Importantly, changes in the phase and the amplitude of an electric field affect its far-field diffraction pattern in different ways. In general, the phase profile determines the position of the interference peaks, and the amplitude modifies the fringe visibility. This effect can be demonstrated through a simulation of a pair of slits, as shown in Figure 3. The figure shows the simulated far-field diffraction pattern, for different values of amplitude and phase for the right-hand side slit. As shown in the figure, a change in the relative phase of the two slits changes the position of the interference peaks, and a change in the relative amplitude changes the fringe visibility (but it does not change the peak position). Therefore, in the case of this simulation, the spatial frequencies corresponding to the relative phase of the two slits exists in the diffraction pattern even when the amplitude is changed. The signal-to-noise ratio (i.e., fringe

visibility) of this spatial frequency information corresponds directly to the relative amplitudes of the two slits, and this ratio will drop to zero for cases where the amplitude of either slit is zero. Whilst the complexity of the diffraction pattern increases when the number of light sources is increased to seven (i.e., in the case for the simulated fibres in this work), the physical principle remains. It could therefore be argued that the neural network can identify the phase values of the fibres when the amplitudes have changed (despite never encountering this during training) as the position of the interference peaks does not change. Similarly, it could be argued that the phase prediction accuracy decreases when the amplitude of a fibre decreases, as the fringe visibility also decreases.

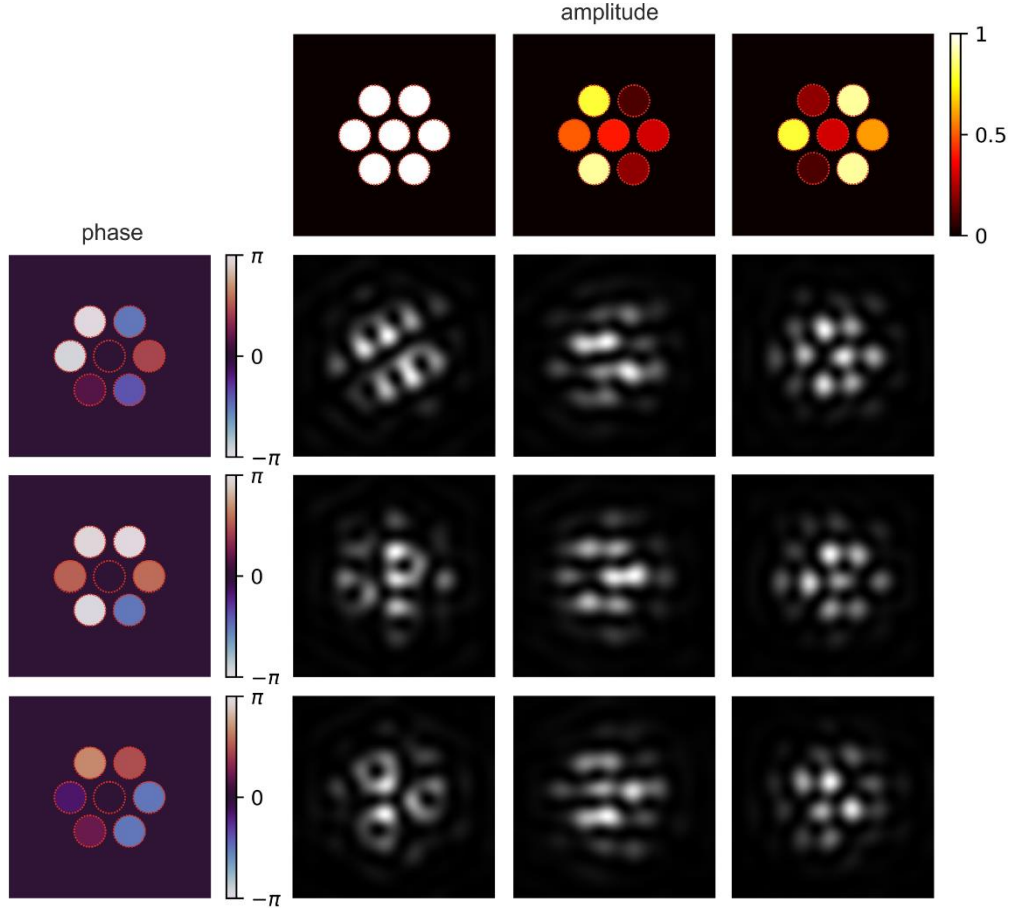


Fig. 1. Simulated far-field intensity patterns corresponding to different combinations of phase values and amplitude values.

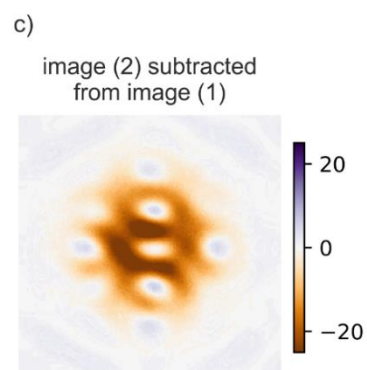
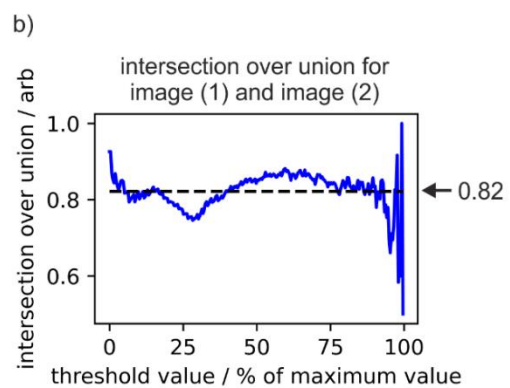
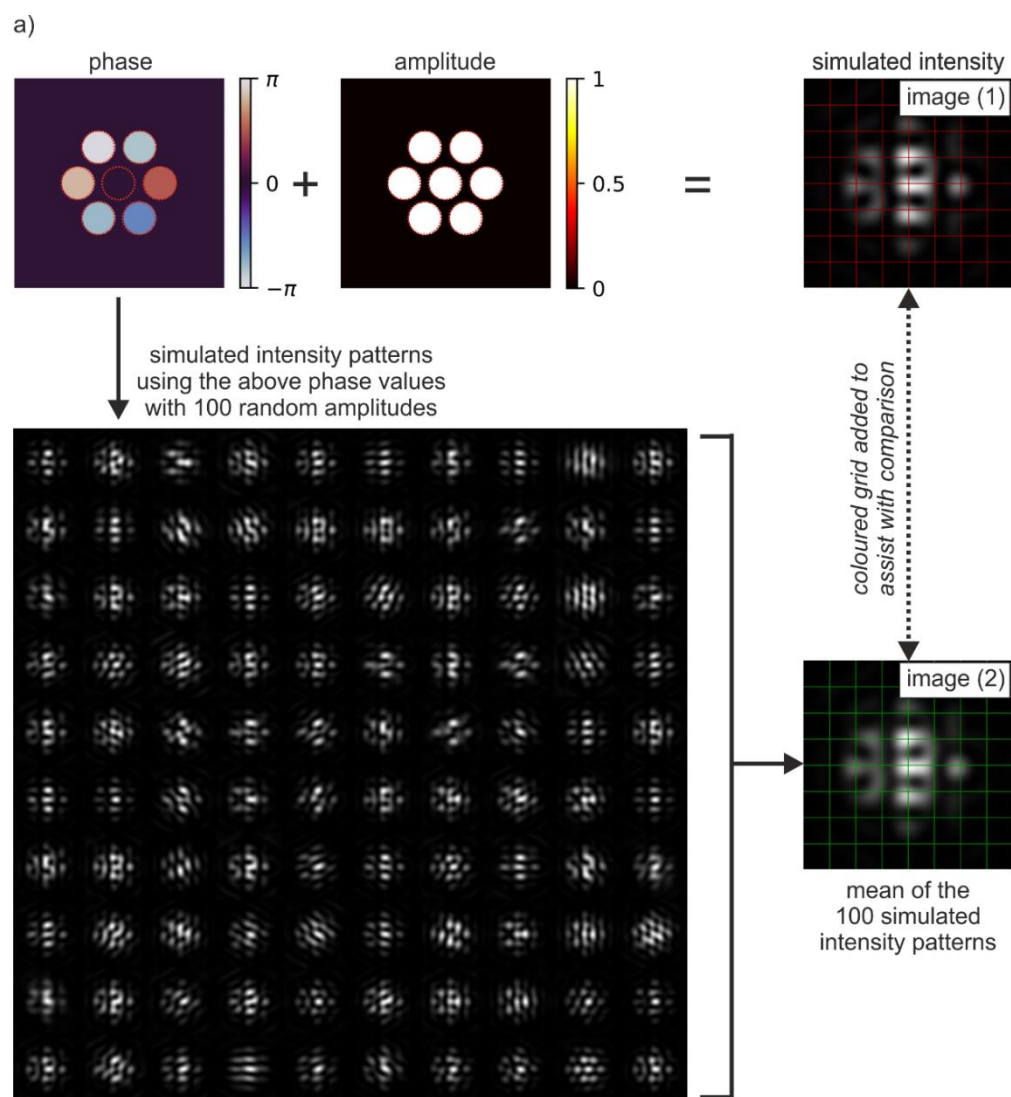


Fig. 2. Showing a) that the intensity pattern for equal amplitude values is comparable to the mean of the intensity patterns for non-equal amplitude values, with quantitative metrics of b) intersection over union and c) image subtraction.

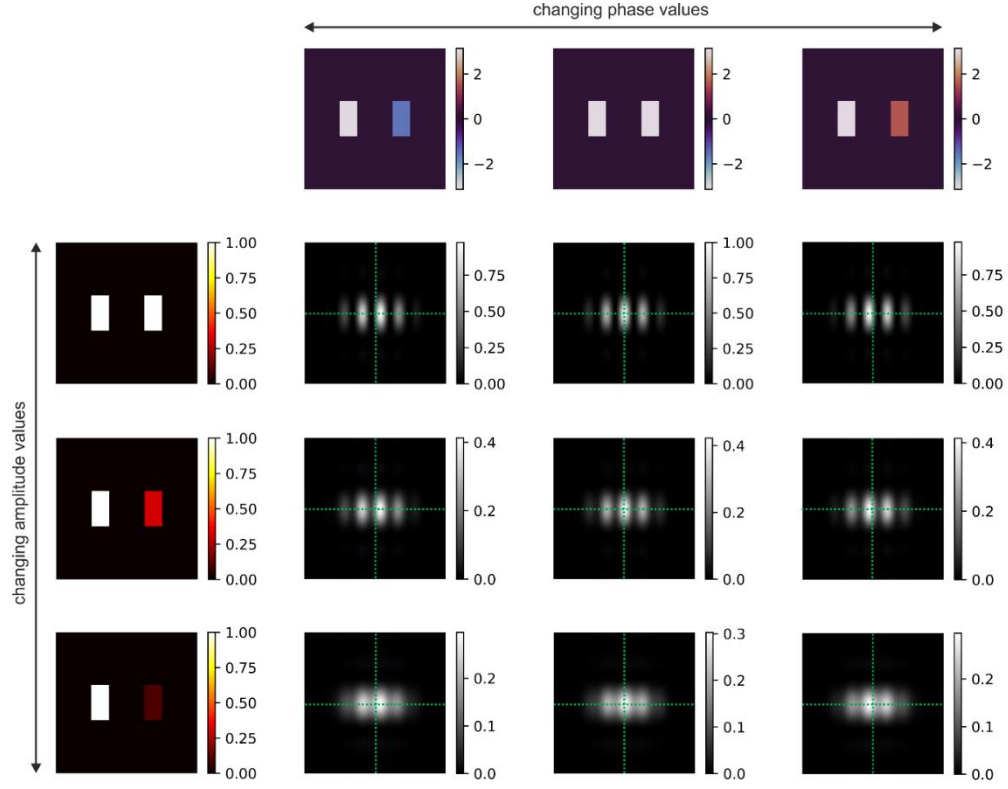


Fig. 3. Simulated far-field diffraction patterns from two slits, for changing amplitude and phase values. The dotted green lines are included for assisting with visual comparison.

3. The neural network

Here, a cGAN was used to transform an image corresponding to the simulated intensity profile, into an image corresponding to the phase profile of the simulated fibres. The neural network had a generator and discriminator adversarial architecture, with the generator based on an encoder-decoder U-Net model with skip connections (as shown in Figure 4), and the discriminator downscaling the generated 256×256 resolution images down to a 32×32 images before judging them.

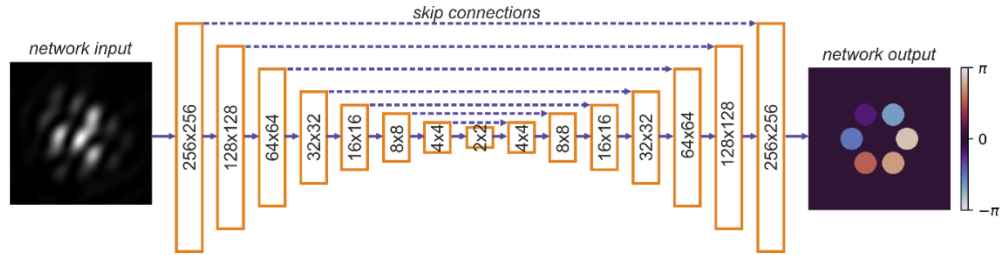


Fig. 4. Schematic of the image-to-image network used for transforming simulated diffraction patterns into predictions of phase.

The network was trained with a batch size of 2, learning rate of 0.0002 (Adam, beta=0.5), and with a L1-to-GANloss ratio of 100:1. The beam propagation simulation, and the neural network and its training parameters, were identical to those presented previously [47], except for the use of seven fibres in this work. As shown in figure 5a), the neural network was trained to identify the phase values directly from the simulated intensity profiles. The training data was formed of

100,000 pairs of images, corresponding to seven fibres with randomly generated phase profiles and equal amplitudes, and their associated intensity profiles. The neural network was therefore trained on data where all fibres had an amplitude value of one, and then tested on data where all fibres had a random amplitude value between zero and one. This ensured that the testing data was not encountered during training (except for the single case where all amplitude values were one).

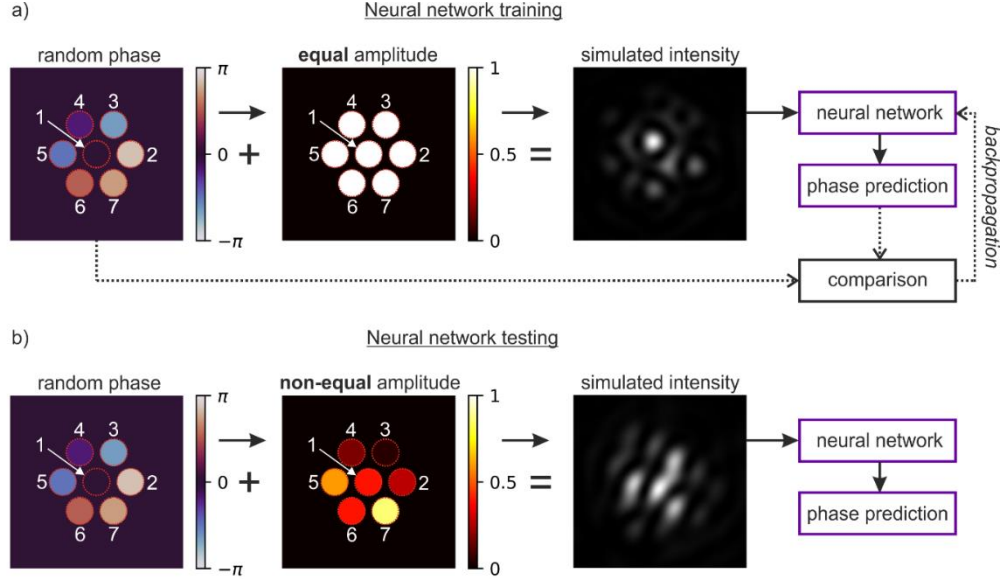


Fig. 5. Concept for a) training and b) testing the neural network. The neural network is trained to identify phase values from simulated far-field intensity patterns. The training data contains intensity patterns corresponding to random phase values and equal amplitude values (the amplitudes are always set to a value of one). The testing data contains intensity patterns corresponding to random phase values and random amplitude values (the amplitudes are chosen from a uniform distribution from zero to one).

As the diffraction patterns would, in practice, be recorded as a camera image for an experimental implementation, an image-to-image neural network architecture was chosen. This meant that the phase information would likewise need to be converted into an RGB image. Due to the cyclic nature of phase, a single channel of the RGB image was not sufficient for the neural network to learn the process of phase identification, and instead the red and blue image channels were used to represent the sine and cosine of the phase for each fibre. A flow chart describing this approach in detail is shown in figure 1 in [47]. The phase and amplitude information were encoded into the RGB channels of a 256×256 -pixel resolution image with 24-bit depth, with the phase encoded using a $\sin(\theta)$ and $\cos(\theta)$ approach to ensure continuity across the $-\pi$ to $+\pi$ boundary. The phase of the central fibre was always set to zero, to avoid the infinite number of equivalent solutions associated with a global change in phase value. The amplitude was encoded in the green channel of the image and was scaled accordingly between the values of 0 and 255. The simulated intensity pattern was encoded through the scaling of intensity values between the values of 0 and 255, with all RGB taking these values, hence giving a grayscale image. The input for the neural network was therefore a grayscale image corresponding to the intensity pattern, and the output (i.e., the neural network prediction) was an RGB image that encoded the phase and amplitude information. As indicated by the figure, all training data pairs corresponded to a set of equal amplitude values. As shown in figure 5b), after the neural network was trained, it was tasked with the identification of phase values from simulated intensity patterns corresponding to non-equal (i.e., randomly chosen)

amplitude values. In other words, the neural network was trained on equal amplitude values, and was tested on non-equal amplitude values.

4. Results and discussions

One of the primary challenges associated with the application of neural networks is quantifying the encoding that the network has learnt during training. Typically, an image-based neural network will have of order millions or billions of programmable neurons, and hence a direct investigation of neurons is rarely effective. Despite the challenges, a wide range of approaches have been investigated, such as Grad-Cam [50] and Saliency maps [51]. Alternatively, the network can be applied to a range of tasks, and the network outputs can be analysed. This is, in essence, the approach taken in this work, where the neural network was trained on task A (i.e., phase identification when the fibre amplitudes are equal) and then tested on task B (i.e., phase identification when the fibre amplitudes are non-equal). The behaviour of the network on task B can then provide some level of qualitative understanding of the capability of the network. In this section, three such task B examples are presented. Firstly, when the fibre amplitudes are either 0 or 1 (figure 6). Secondly, when the central fibre has an amplitude of 0, and when the spatial positions of the fibres are changed (figure 7). Thirdly, when the fibre amplitudes can take any value between 0 and 1 (figure 8).

Figure 6 shows the capability of the neural network for identification of phases values where fibres have an amplitude of 0 or 1 (i.e., ‘on’ or ‘off’). In a), the columns correspond to the phase, amplitude, and simulated intensity, along with the predicted phase, where all four columns are shown for cases of 2 to 7 fibres with an amplitude of 1. In all cases, the central fibre has an amplitude of 1. As shown in the figure (by the green ticks), the neural network accurately predicts the phase for all fibres that have a non-zero amplitude. Whilst the phase value for zero amplitude is of course meaningless, the neural network still predicts a phase value for each fibre. It is notable that the phase values are predicted correctly for all cases, as the intensity patterns for 2 to 6 fibres ‘on’ were not experienced by the network during training, as only intensity patterns for 7 fibres ‘on’ were used. Part b) shows that the prediction accuracy increases slightly for higher numbers of ‘on’ fibres. This could be since the larger number of ‘on’ fibres corresponds more similarly to the training data, or that the interference from multiple fibres is required, and hence each ‘off’ fibre removes some of the information used to identify the phases.

Figure 7a) shows the capability of phase prediction when the central fibre is removed. As noted earlier, in the training data, the phase of the central fibre was always set to zero, to remove the infinite number of solutions associated with the property that a global phase value change would produce an identical intensity pattern. Therefore, when the central fibre is removed, the phase predictions are significantly less accurate, with analysis showing that 93% of the prediction accuracy comes from interference with the central fibre and 7% from interference with other outer ring fibres. This provides strong evidence that the neural network uses the central fibre as a fixed reference when predicting the phase values of the surrounding 6 fibres. Figure 7b) shows the effect of translating the position of the fibres. Here, the hexagonal close packed arrangement is transformed into a ‘plus’ pattern. The ‘plus’ pattern was specifically chosen as it shares the position of three fibres with the hexagonal arrangement, but also has two positions that are not shared. Of course, the neural network only experienced the 7-fibre hexagonal arrangement during training, and hence the intensity pattern from the ‘plus’ arrangement would not have been experienced during training. Here, the neural network correctly predicts the phase values for the fibres that match the position of the hexagonal arrangement and fails to predict the phase values for the fibres for the un-matched positions. This is clear evidence that the neural network is only able to identify the phase of fibres that are in a position that was experienced during training. Interestingly, the neural network still predicts a value for the phase of the fibres in the un-matched position, and understandably they

are incorrect, as there is no mathematical combination of phases that could allow the hexagonal arrangement to produce the intensity patterns that was associated with the ‘plus’ arrangement.

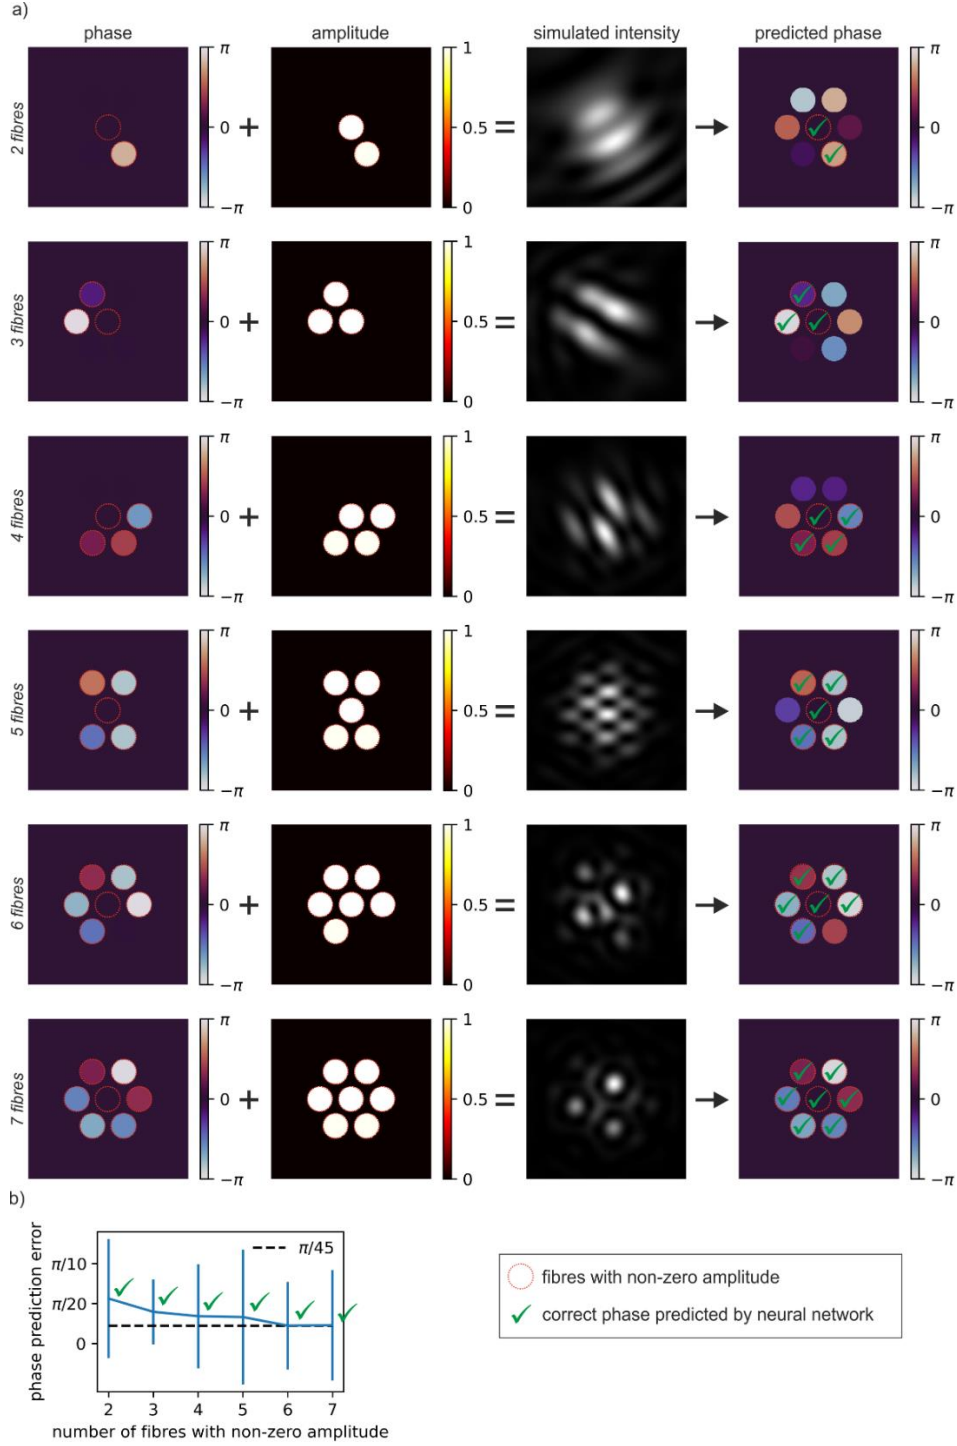


Fig. 6. Showing a) phase, amplitude, and simulated intensity patterns, along with the phase values predicted by the neural network directly from the simulated intensity patterns, and b) the phase prediction error for different numbers of fibres with non-zero amplitude values.

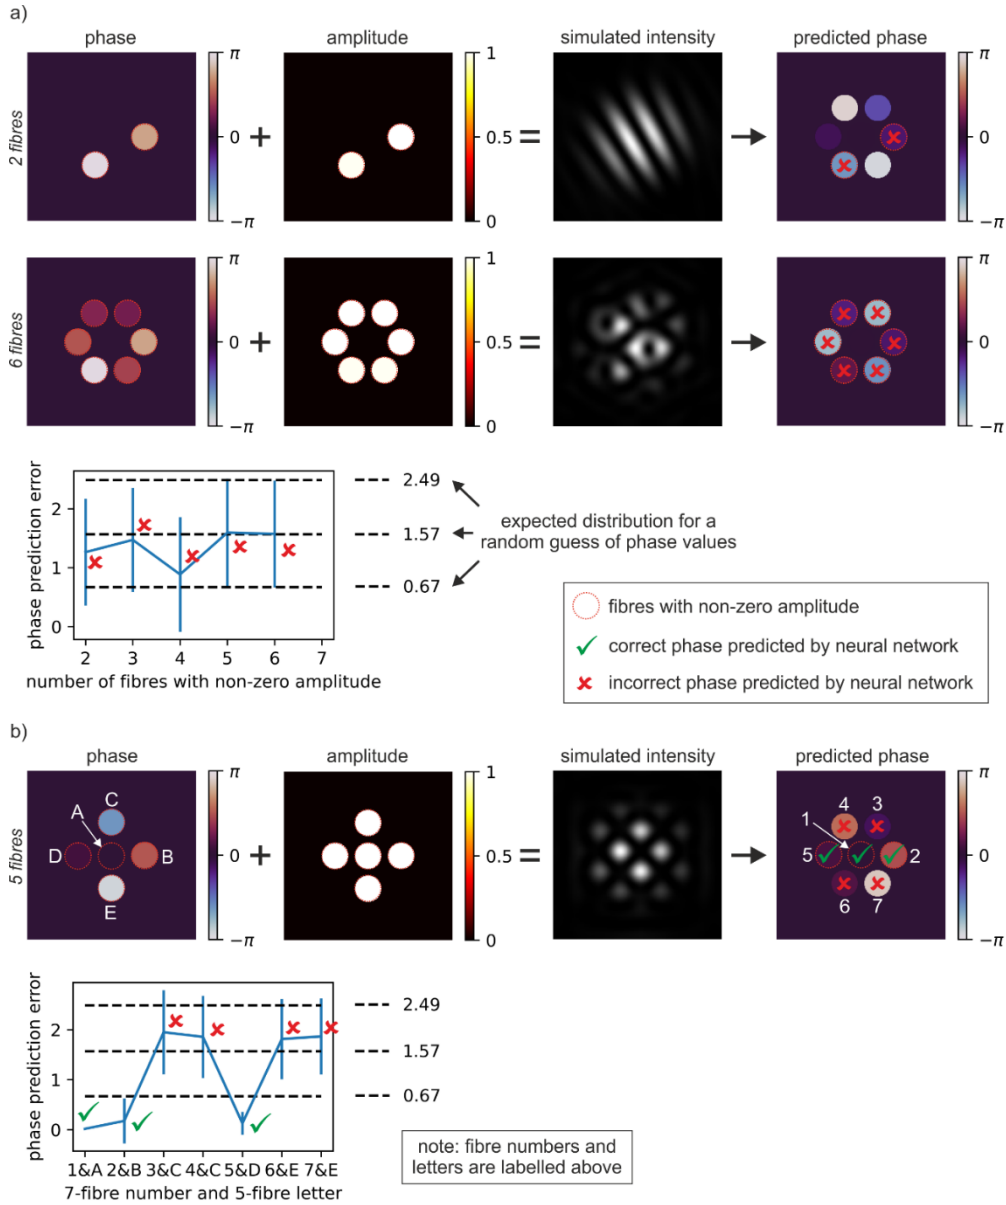


Fig. 7. Showing phase, amplitude, and simulated intensity patterns, along with the phase values predicted by the neural network for a) cases where the central fibre amplitude is set to zero (i.e., removed), and b) where the positions of the fibres are changed from a 7-fibre hexagonal arrangement to a 5-fibre 'plus' arrangement.

Figure 8 shows the capability of the neural network for predicting phase values when the amplitude values for each fibre can take any value between 0 and 1. For each of the examples shown in a-c), the phase values are constant, and the amplitude values are randomly chosen. Firstly, 30 examples of intensity patterns (i.e., fixed phase, random amplitude) are presented. Secondly, the neural network phase predictions for each of these intensity patterns are graphed (solid green circles), along with the correct phase values (red plusses). The graphs clearly show that the neural network can identify the phase values in the presence of randomly chosen amplitude values. The intensity of the solid green circles in the figure corresponds to the multiplication of the amplitude of the specified fibre and the amplitude of the central fibre. As

this multiplied value decreases below 25%, the prediction error sharply increases. This effect is likely due to the decreased visibility of the interference fringes when the amplitudes are reduced. Figure 9 shows the ability of a neural network to predict phase values of each fibre when a linear gradient phase profile is added to the simulated phase profile. This capability is also observed when the amplitudes of each fibre are randomly chosen.

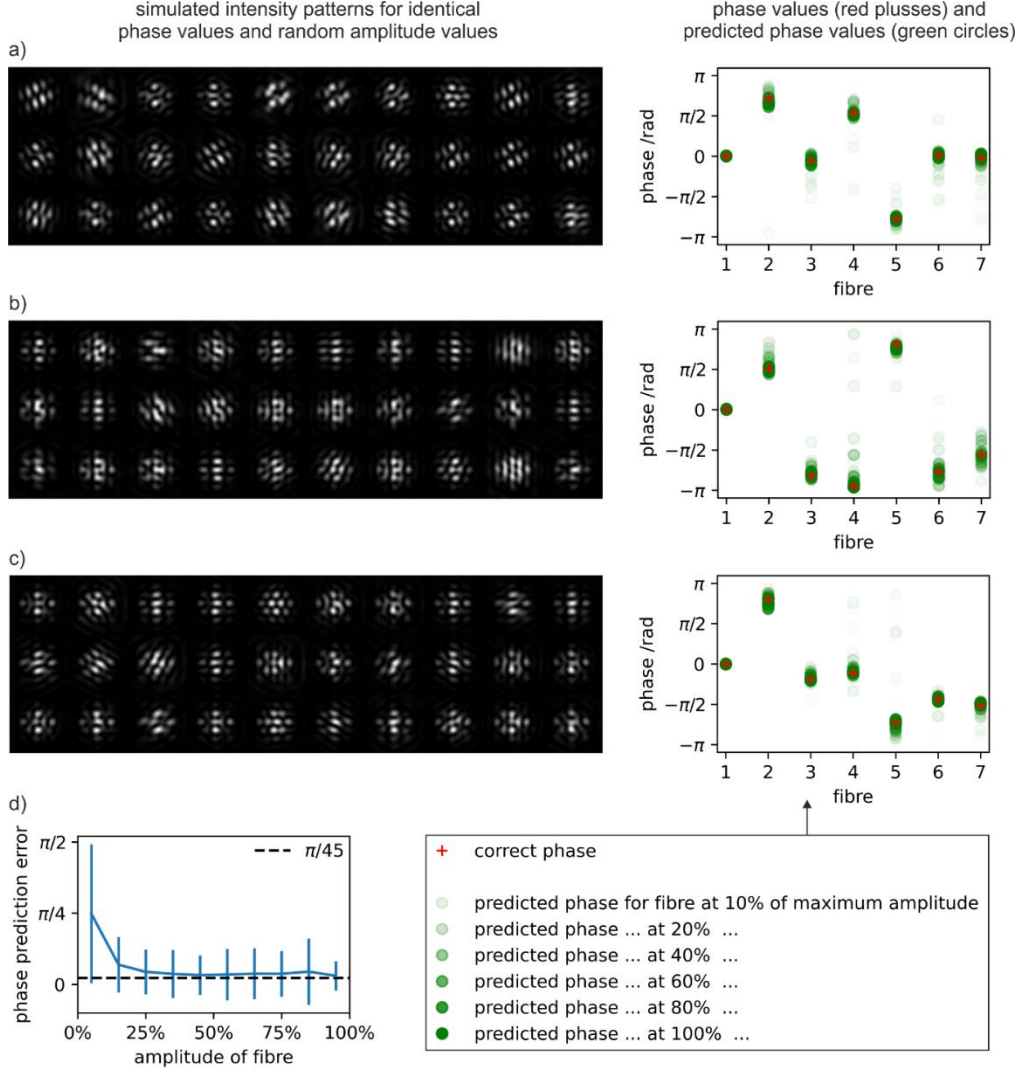


Fig. 8. The capability of the neural network for predicting the phase directly from intensity patterns that correspond to amplitude values with a uniform distribution between zero and one. Showing for each of a-c), thirty examples of simulated intensity patterns corresponding to a fixed phase and random amplitude values, along with the graph showing the phase values predicted by the neural network for each of these thirty intensity patterns. Showing d), the phase prediction error for different amplitudes.

This work has presented analysis of a neural network trained on simulated data, and hence it is important to consider the approach required to apply this technique to a real-world experimental setup. The key challenge will likely be the collection of suitable labelled training data. As the phase cannot be directly identified (after all, this is the motivation for the work), the experimental setup will require an additional interferometric component on each single beam so that the phase of each fibre can be identified through direct measurement of the

produced interference fringes. Such an approach would be costly and challenging to implement experimentally and would therefore not be suitable for use on a commercial coherent beam combination system. The expectation is that an augmented setup would be able to provide labelled experimental training data that would then enable other equivalent coherent beam combination systems to be controlled without the need for this additional interferometric measurement device.

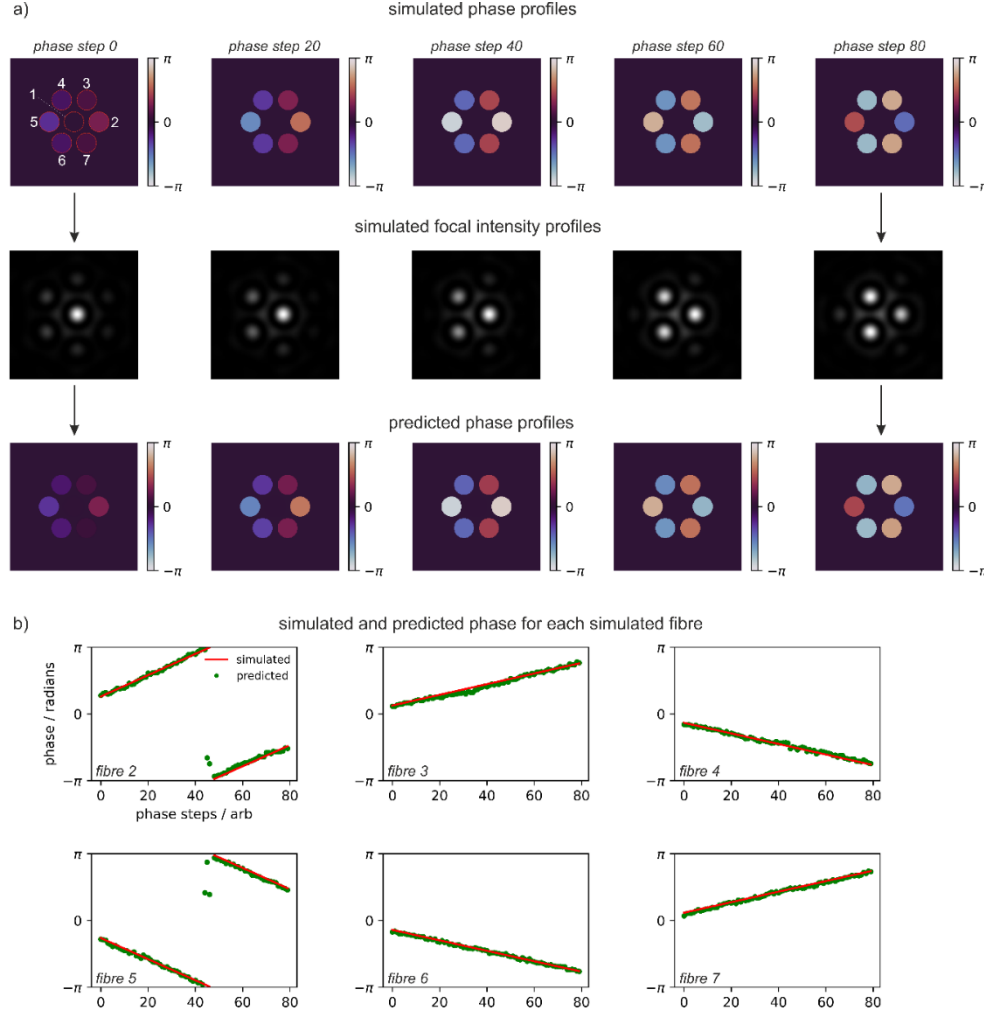


Fig. 9. The capability of the neural network in identification of the phase values when a phase gradient is present in the simulation. Showing a) the simulated phase profile, the associated simulated focal intensity profile, and the prediction phase profile, for a range of phase gradients, and b) comparison of phase values for each fibre.

5. Conclusion

In conclusion, a neural network trained to identify the phase information from intensity patterns corresponding to simulated fibres with equal amplitude profiles was shown capable of accurately completing this task for intensity patterns from fibres with non-equal amplitude profiles. This result has an important practical implication for coherent beam combination, namely that the identification of phase via deep learning may not require additional training data if the fibre amplitudes are intentionally, or unintentionally, varying to some degree.

Funding. Engineering & Physical Sciences Research Council (EP/W028786/1, EP/T026197/1, EP/P027644/1). For the purpose of open access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

Disclosures. The authors declare no conflicts of interest.

Data availability. Data underlying the results presented in this paper are available in Ref. [52].

References

1. D. J. Richardson, J. Nilsson, and W. A. Clarkson, "High power fiber lasers: current status and future perspectives," *JOSA B* **27**, B63-B92 (2010).
2. M. N. Zervas and C. A. Codemard, "High power fiber lasers: a review," *IEEE Journal of selected topics in Quantum Electronics* **20**, 219-241 (2014).
3. C. Jauregui, J. Limpert, and A. Tünnermann, "High-power fibre lasers," *Nature photonics* **7**, 861-867 (2013).
4. J. Canning, "Fibre lasers and related technologies," *Optics and lasers in engineering* **44**, 647-676 (2006).
5. M. N. Zervas, "Transverse mode instability, thermal lensing and power scaling in Yb 3+-doped high-power fiber amplifiers," *Optics express* **27**, 19019-19041 (2019).
6. Z. Liu, X. Jin, R. Su, P. Ma, and P. Zhou, "Development status of high power fiber lasers and their coherent beam combination," *Science China Information Sciences* **62**, 1-32 (2019).
7. G. D. Goodno, C.-C. Shih, and J. E. Rothenberg, "Perturbative analysis of coherent combining efficiency with mismatched lasers," *Optics express* **18**, 25403-25414 (2010).
8. R. Shi, S. A. Khairallah, T. T. Roehling, T. W. Heo, J. T. McKeown, and M. J. Matthews, "Microstructural control in metal laser powder bed fusion additive manufacturing using laser beam shaping strategy," *Acta Materialia* **184**, 284-305 (2020).
9. T. T. Roehling, R. Shi, S. A. Khairallah, J. D. Roehling, G. M. Guss, J. T. McKeown, and M. J. Matthews, "Controlling grain nucleation and morphology by laser beam shaping in metal additive manufacturing," *Materials & Design* **195**, 109071 (2020).
10. A. Möhl, S. Kaldun, C. Kunz, F. A. Müller, U. Fuchs, and S. Gräf, "Tailored focal beam shaping and its application in laser material processing," *Journal of Laser Applications* **31**, 042019 (2019).
11. Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature* **521**, 436-444 (2015).
12. G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. Van Der Laak, B. Van Ginneken, and C. I. Sánchez, "A survey on deep learning in medical image analysis," *Medical image analysis* **42**, 60-88 (2017).
13. J. Gu, Z. Wang, J. Kuen, L. Ma, A. Shahroudy, B. Shuai, T. Liu, X. Wang, G. Wang, and J. Cai, "Recent advances in convolutional neural networks," *Pattern recognition* **77**, 354-377 (2018).
14. Z. Li, F. Liu, W. Yang, S. Peng, and J. Zhou, "A survey of convolutional neural networks: analysis, applications, and prospects," *IEEE transactions on neural networks and learning systems* (2021).
15. A. Ajit, K. Acharya, and A. Samanta, "A review of convolutional neural networks," in *2020 international conference on emerging trends in information technology and engineering (ic-ETITE)*, (2020), 1-5.
16. A. Barucci, C. Cucci, M. Franci, M. Loschiavo, and F. Argenti, "A deep learning approach to ancient egyptian hieroglyphs classification," *Ieee Access* **9**, 123438-123447 (2021).
17. T. Guidi, L. Python, M. Forasassi, C. Cucci, M. Franci, F. Argenti, and A. Barucci, "Egyptian Hieroglyphs Segmentation with Convolutional Neural Networks," *Algorithms* **16**, 79 (2023).
18. G. Carleo, Y. Nomura, and M. Imada, "Constructing exact representations of quantum many-body systems with deep neural networks," *Nature communications* **9**, 5322 (2018).
19. A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Communications of the ACM* **60**, 84-90 (2017).
20. J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*, (2009), 248-255.
21. I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," *Communications of the ACM* **63**, 139-144 (2020).
22. P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, (2017), 1125-1134.
23. Y. Rivenson, Y. Wu, and A. Ozcan, "Deep learning in holography and coherent imaging," *Light: Science & Applications* **8**, 1-8 (2019).
24. Y. Rivenson, Y. Zhang, H. Günaydin, D. Teng, and A. Ozcan, "Phase recovery and holographic image reconstruction using deep learning in neural networks," *Light: Science & Applications* **7**, 17141-17141 (2018).
25. J. A. Grant-Jacob, M. Praeger, M. Loxham, R. W. Eason, and B. Mills, "Lensless imaging of pollen grains at three-wavelengths using deep learning," *Environmental Research Communications* **2**, 075005 (2020).
26. J. A. Grant-Jacob, B. S. Mackay, J. A. Baker, Y. Xie, D. J. Heath, M. Loxham, R. W. Eason, and B. Mills, "A neural lens for super-resolution biological imaging," *Journal of Physics Communications* **3**, 065004 (2019).
27. H. Wang, Y. Rivenson, Y. Jin, Z. Wei, R. Gao, H. Günaydin, L. A. Bentolila, C. Kural, and A. Ozcan, "Deep learning enables cross-modality super-resolution in fluorescence microscopy," *Nature methods* **16**, 103-110 (2019).

28. Y. Rivenson, H. Ceylan Koydemir, H. Wang, Z. Wei, Z. Ren, H. Günaydin, Y. Zhang, Z. Gorocs, K. Liang, and D. Tseng, "Deep learning enhanced mobile-phone microscopy," *Acs Photonics* **5**, 2354-2364 (2018).
29. B. Mills, D. J. Heath, J. A. Grant-Jacob, and R. W. Eason, "Predictive capabilities for laser machining via a neural network," *Optics express* **26**, 17245-17253 (2018).
30. D. J. Heath, J. A. Grant-Jacob, Y. Xie, B. S. Mackay, J. A. Baker, R. W. Eason, and B. Mills, "Machine learning for 3D simulated visualization of laser machining," *Optics Express* **26**, 21574-21584 (2018).
31. B. Mills and J. A. Grant-Jacob, "Lasers that learn: The interface of laser machining and machine learning," *IET Optoelectronics* **15**, 207-224 (2021).
32. H. Chang, Q. Chang, J. Xi, T. Hou, R. Su, P. Ma, J. Wu, C. Li, M. Jiang, and Y. Ma, "First experimental demonstration of coherent beam combining of more than 100 beams," *Photonics Research* **8**, 1943-1948 (2020).
33. P. Ma, H. Chang, Y. Ma, R. Su, Y. Qi, J. Wu, C. Li, J. Long, W. Lai, and Q. Chang, "7.1 kW coherent beam combining system based on a seven-channel fiber amplifier array," *Optics & Laser Technology* **140**, 107016 (2021).
34. J. Song, Y. Li, D. Che, J. Guo, and T. Wang, "Coherent beam combining based on the SPGD algorithm with a momentum term," *Optik* **202**, 163650 (2020).
35. P. Zhou, Z. Liu, X. Wang, Y. Ma, H. Ma, X. Xu, and S. Guo, "Coherent beam combining of fiber amplifiers using stochastic parallel gradient descent algorithm and its application," *IEEE journal of selected topics in quantum electronics* **15**, 248-256 (2009).
36. J. R. Fienup, "Phase retrieval algorithms: a comparison," *Applied optics* **21**, 2758-2769 (1982).
37. T. Hou, Y. An, Q. Chang, P. Ma, J. Li, D. Zhi, L. Huang, R. Su, J. Wu, and Y. Ma, "Deep-learning-based phase control method for tiled aperture coherent beam combining systems," *High Power Laser Science and Engineering* **7**(2019).
38. Q. Chang, T. Hou, H. Chang, P. Ma, R. Su, Y. Ma, and P. Zhou, "Iteration-free, simultaneous correction of piston and tilt distortions in large-scale coherent beam combination systems," *Optics Express* **29**, 34863-34879 (2021).
39. D. Wang, Q. Du, T. Zhou, D. Li, and R. Wilcox, "Stabilization of the 81-channel coherent beam combination using machine learning," *Optics express* **29**, 5694-5709 (2021).
40. T. Hou, Y. An, Q. Chang, P. Ma, J. Li, L. Huang, D. Zhi, J. Wu, R. Su, and Y. Ma, "Deep-learning-assisted, two-stage phase control method for high-power mode-programmable orbital angular momentum beam generation," *Photonics Research* **8**, 715-722 (2020).
41. V. François-Lavet, P. Henderson, R. Islam, M. G. Bellemare, and J. Pineau, "An introduction to deep reinforcement learning," *Foundations and Trends in Machine Learning* **11**, 219-354 (2018).
42. H. Tünnermann and A. Shirakawa, "Deep reinforcement learning for coherent beam combining applications," *Optics express* **27**, 24223-24230 (2019).
43. H. Tünnermann and A. Shirakawa, "Deep reinforcement learning for tiled aperture beam combining in a simulated environment," *Journal of Physics: Photonics* **3**, 015004 (2021).
44. X. Zhang, P. Li, Y. Zhu, C. Li, C. Yao, L. Wang, X. Dong, and S. Li, "Coherent beam combination based on Q-learning algorithm," *Optics Communications* **490**, 126930 (2021).
45. M. Shpakovych, G. Maulion, V. Kermene, A. Boju, P. Armand, A. Desfarges-Berthelemy, and A. Barthelemy, "Experimental phase control of a 100 laser beam array with quasi-reinforcement learning of a neural network in an error reduction loop," *Optics Express* **29**, 12307-12318 (2021).
46. M. Shpakovych, G. Maulion, A. Boju, P. Armand, A. Barthélémy, A. Desfarges-Berthelemy, and V. Kermene, "On-Demand Phase Control of a 7-Fiber Amplifiers Array with Neural Network and Quasi-Reinforcement Learning," in *Photonics*, (MDPI, 2022), 243.
47. B. Mills, J. A. Grant-Jacob, M. Praeger, R. W. Eason, J. Nilsson, and M. N. Zervas, "Single step phase optimisation for coherent beam combination using deep learning," *Scientific Reports* **12**, 1-12 (2022).
48. J. A. Grant-Jacob, Y. Xie, B. S. Mackay, M. Praeger, M. D. McDonnell, D. J. Heath, M. Loxham, R. W. Eason, and B. Mills, "Particle and salinity sensing for the marine environment via deep learning using a Raspberry Pi," *Environmental Research Communications* **1**, 035001 (2019).
49. A. K. Jain, "Working model of self-driving car using convolutional neural network, Raspberry Pi and Arduino," in *2018 Second International Conference on Electronics, Communication and Aerospace Technology (ICECA)*, (IEEE, 2018), 1630-1635.
50. R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-cam: Visual explanations from deep networks via gradient-based localization," in *Proceedings of the IEEE international conference on computer vision*, (2017), 618-626.
51. K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," *arXiv preprint arXiv:1312.6034* (2013).
52. J. A. Grant-Jacob, B. Mills, and M. N. Zervas, "Phase identification despite amplitude variation in coherent beam combination using deep learning Dataset," (2023), <https://doi.org/10.5258/SOTON/D2344>.