

Stabilisation of transverse mode purity in a radially polarised Ho:YAG laser using machine learning

Thomas L. Jefferson-Brain¹, Matthew J. Barber¹, Azaria D. Coupe², W. Andrew Clarkson^{1*}, Peter C. Shardlow¹

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

² Physics & Astronomy, University of Southampton, Southampton, SO17 1BJ, UK

* wac@orc.soton.ac.uk

Received: date / Revised version: date

Abstract Radially polarised solid-state lasers offer attractive improvements in materials processing applications, but selection and stabilisation of the appropriate radially polarised mode is much more challenging than for the fundamental mode. Here, we demonstrate automated stabilisation of a radially polarised Ho:YAG laser by utilising laser mode analysis computed from a convolutional neural network. The neural network predicts the transverse modal content from single plane intensity images with high accuracy on timescales of a few milliseconds, permitting real-time self-adjustment of the laser cavity. Radially polarised emission has been maintained across a 30 W range of pump power, with the stabilisation of other arbitrary laser modes using the same neural network also demonstrated.

1 Introduction

Radially polarised transverse laser modes have a polarisation which continuously varies across the beam profile, with the electric field directed radially outwards from the centre of the mode. This leads to a polarisation discontinuity at the centre of the mode and hence the intensity falls to zero at the centre, resulting in a ‘donut-like’ appearance. The lowest order radially polarised mode shares the same intensity profile as the LG_{01} (Laguerre-Gaussian) transverse mode, and has shown particular promise in applications such as high-resolution imaging [1], particle acceleration [2] and laser materials processing [3]. To this end, the use of radially polarised modes in melt shearing laser cutting applications may enable cutting efficiencies which are 1.5–2 times higher than linearly polarised and circularly polarised beams [4].

Solid-state laser sources are an attractive platform for generating radial polarisation as thermally-induced stress birefringence and the associated bifocussing in the gain medium can be exploited to permit relatively easy

selection of a radially polarised mode [5]. However, due to the thermal dependence of this process, it is generally difficult to control the output power via variation of the pump power without adversely affecting the laser’s radial polarisation purity.

Furthermore, extra-cavity control of the laser power at a specific pump power cannot be accomplished with a half-wave plate and polariser in the same way that it can be for a linearly polarised laser source, so an alternative strategy is required whereby the laser resonator length is continuously re-optimised as the pump power is varied to maintain mode purity. While the beam degradation produced by a change in the pump power can generally be compensated by adjusting the laser cavity, the need for frequent optimisation of the system restricts the real-world usage of radially polarised sources.

If the modal purity optimisation could be performed in real-time by an automated feedback protocol, the implementation of radially polarised sources in a variety of applications would be much more attractive. To this end, a range of simple, fast and accurate techniques for the measurement of modal purity have recently been demonstrated using deep learning convolutional neural networks (CNNs) [6–8]. Such use of a CNN has been shown to allow high-accuracy predictions of laser modal content using single plane intensity distributions on a sub-millisecond timescale [7,8].

In this paper, we present the application of a bespoke CNN as a pseudo-quantitative modal purity metric to enable real-time feedback control of a radially polarised solid-state laser. By predicting the modal content of the laser output from two cameras (with and without a linear polariser), a control algorithm can be employed to automatically adjust the laser cavity in order to improve the quality of the desired transverse laser mode. As such, we demonstrate self-stabilisation of a radially polarised Ho:YAG laser across 30 W of pump power variation.

Recently, there have been major advancements in pattern recognition algorithms facilitated by the use of

deep learning CNNs, which can extract characteristics or features of an image rather than identifying more simplistic linear relationships between patterns of data and a series of labels and parameters. To process an image with a CNN, the input is subject to a multitude of convolutional filter transformations, arranged in a ‘deep’ series of layers. By connecting the layer outputs with optimally weighted nodes in the network (known as ‘neurons’), the convolutional filters are able to assign image characteristics such as edges, arrangements of edges and collections of arrangements to an output which documents the features recognised by the network [9].

The relationship between the data characteristics and the labels/parameters is learned automatically in an iterative process, where the CNN is provided with a set of training data containing known features. Initial neuron weightings are iteratively adjusted in order to minimise the error between the output feature vector from the network and the known feature vector from the training data. The main computational load of implementing the CNN occurs in the training of the network, which may take many hours, such that the final optimised algorithm is very fast at performing data feature recognition.

The speed of CNNs and their ability to predict parameters from complex non-linear problems lends them to applications in optical physics where dynamics can be difficult to model and takes place on short timescales. Recent applications have included laser mode-locking [10], imaging [11, 12], microscopy [13] and coherent beam combination [14]. Furthermore, CNNs have been used to rapidly predict the relative modal compositions of solid-state and fibre laser beam profiles [7, 8].

In contrast, traditional (non-CNN) modal content prediction algorithms have relied on experimentally [15–19] or computationally [20–22] intensive methods, requiring computation times on the order of hundredths of a second or more [23]. As such, CNN modal prediction techniques can enable substantially faster and computationally less intensive optimisation of the transverse mode in a solid-state laser when compared to the more traditional algorithms that are used for the computation of modal content.

2 Machine learning mode identification

Here, we have developed a bespoke CNN for the analysis of single plane intensity images, typical of those obtained from low-cost CCD and microbolometer-based beam imaging cameras. Figure 1 provides a visualisation of the structure of the final CNN, where a 120×120 8-bit greyscale image is input into the network, with the image size stipulated by the camera sensor resolution (FLIR Lepton 3.0). The input image is then convolved with 8 filters of 5×5 dimensions, after which a rectified linear unit (ReLU) layer is used to set any resulting negative values to zero. The rectified convolution results




Fig. 1 Visualisation of the CNN structure.

are then down-sampled using a pooling layer with size 2 and stride 2, where the input is divided into rectangular regions and the maximum of each region is output as the down-sampled result. These steps are repeated twice more, with the original image becoming increasingly down-sampled while the convolution processes increase the contrast of key features.

A fully connected layer with 4,096 elements then multiplies the outputs of the previous layers by a weighting vector and adds a bias vector, where the values of the weight and bias are established during the training procedure. A dropout layer is then included which discards 20% of the weighted data to help prevent over-fitting in the next layer. Finally, another fully connected layer with N elements (where N is the number of modes) is combined with a regression layer in order to estimate the modal composition of the input image.

To train the CNN, 23,000 120×120 8-bit greyscale images of transverse laser modes were generated from the equations for the intensity distributions of Hermite-Gaussian (HG) and Laguerre-Gaussian (LG) modes. In particular, nine transverse laser modes were selected, based on those that were observed during construction of the Ho:YAG laser: HG₀₀, HG₀₁, HG₁₁, LG₀₁, LG₀₂, LG₀₃, LG₀₄, LG₀₅ and LG₀₆.

The generated beam profiles each had randomised orientations, positions, radii, and intensities in order to emulate realistic experimental conditions that would be encountered during operation. Incoherent combinations of neighbouring mode orders were also included in the data, with the intensity-normalised percentage modal composition of each training image recorded in a label vector. Images comprised of filtered Gaussian noise were used to train the CNN to recognise the absence of a laser beam on the camera.

The network was trained for 2,000 epochs on a Nvidia RTX 2060 GPU, taking approximately 80 minutes to complete. The training iteratively increased the scale of the parameters until the improvement in the theoretical modal content prediction accuracy was negligible. A second set of 7,300 randomised beam images was also generated theoretically, to be used as the validation data for the training process.




Fig. 2 Optical schematic of the 2090 nm radially polarised Ho:YAG rod laser.

3 Radially polarised Ho:YAG laser

To demonstrate machine learning stabilisation of a radially polarised transverse mode, an end-pumped Ho:YAG laser was constructed, as shown in Figure 2. The laser cavity was formed between a highly reflective (HR, $R = 99.9\%$ at $2.1\text{ }\mu\text{m}$) mirror and a partially reflective output coupler (OC, $R = 50\%$ at $2.1\text{ }\mu\text{m}$) mirror, both being plane-parallel and anti-reflection (AR)-coated at $1.9\text{ }\mu\text{m}$. The laser had a 250 mm total cavity length and contained a Ho:YAG rod crystal that was 40 mm long, 5 mm in diameter and 0.5 at.% Ho-doped, with AR coatings for 1.9 and $2.1\text{ }\mu\text{m}$ wavelengths on each end. The rod was mounted in a water-cooled copper heatsink maintained at 18°C and positioned close to the HR mirror.

An $f = 40$ mm intra-cavity lens was positioned close to the output coupler mirror. This lens was mounted onto a computer-controlled translation stage driven by a stepper motor, allowing the separation between the lens and the output coupler to be varied. The lens position was used to control the size of the cavity modes within the Ho:YAG crystal and therefore modify the overlap with the population inversion, allowing the preferred laser modes to be selected.

The radially polarised mode of interest shares the same donut-shaped intensity profile with the azimuthally polarised and randomly polarised LG_{01} modes. As a result, these modes are essentially degenerate (in terms of the spatial overlap with the inversion distribution), having nominally the same threshold pump power. However, this degeneracy can be lifted by the thermally-induced stress birefringence and bifocussing in the Ho:YAG crystal by virtue of different thermal lens strengths, which can be exploited to favour preferred lasing on the radially polarised mode via careful optimisation of the resonator length and/or position of the intra-cavity lens [24].

The Ho:YAG laser was pumped by a home-built 52 W thulium fibre laser, locked to 1907 nm with a pair of fibre Bragg gratings. The fibre laser itself was pumped by three 793 nm fibre-coupled diode lasers (DILAS 35 W in a $105/125\text{ }\mu\text{m}$ 0.22 NA fibre) combined by a 3×1 pump combiner (Gooch and Housego). The 1907 nm laser light was emitted from a $10\text{ }\mu\text{m}$ 0.15 NA single-mode passive fibre, with a maximum measured beam propagation factor (M^2) of 1.03 at 52 W. This was spliced onto a $105/125\text{ }\mu\text{m}$ 0.22 NA multimode fibre, which itself was spliced onto a low-loss adiabatically-tapered $100/200\text{ }\mu\text{m}$ capillary fibre with an outer fluorine-doped 0.22 NA ring

(tapered and spliced by Gooch and Housego), with a theoretical M^2 of around 36. The capillary fibre conditioned the pump beam into a ring of light which was relay-imaged into the Ho:YAG crystal via a $6 \times$ telescope, formed from an $f = 25$ mm and an $f = 150$ mm plano-convex lens pair. As a result, the pump telescope produced a 1.2 mm diameter pump spot waist within the Ho:YAG crystal. The ring-shaped pump spot profile was utilised as it can allow maximum extraction of the population inversion when the laser is operating on the donut-shaped LG_{01} cavity mode.

A verification of the laser's ability to operate on the radially polarised transverse mode was undertaken for 28 W of launched 1907 nm pump power and 6 W of 2090 nm output power. When the intra-cavity lens position was optimised for radially polarised emission, the M^2 was measured to be 2.0 and 2.1 in the x and y axes respectively (Figure 3), which closely matches the theoretical value of 2.0 for a pure LG_{01} transverse mode. The polarisation purity of the radially polarised laser was also measured in order to evaluate the desired laser output quality. To do so, a sample of the laser beam was passed through a linear polariser and onto an imaging sensor (FLIR Lepton 3.0), generating a two-lobe intensity pattern that is aligned with the polariser axis for an incident radially polarised beam (Figure 4). Variation of the beam intensity was mapped along a circular path




Fig. 3 Beam propagation factor (M^2) of the radially polarised Ho:YAG laser at 28 W of incident pump power, indicating a high quality LG_{01} mode. Inset: near-field 1907 nm pump beam profile.




Fig. 4 (a) Radially polarised beam profile at 28 W of incident pump power; (b)–(e) two-lobe intensity patterns for the radially polarised beam transmitted through a linear polariser. The polariser axis is orientated as shown by the arrows.

which intersects both maxima of the two-lobe pattern. The intensity map, I , was fitted to the function

$$I = a + b \sin^2(\theta + c) + d \sin(\theta + f), \quad (1)$$

where θ is the azimuthal angle and a, b, c, d and f are the fitting coefficients. From Equation 1, the radial polarisation purity can be defined as

$$P = \frac{b}{a+b}, \quad (2)$$

where P can take values between 0 and 1, such that 0 represents no radial polarisation in the beam and 1 indicates a perfect radially polarised transverse mode. With 28 W of 1907 nm pump power and 6 W of 2090 nm Ho:YAG output, the linear polariser was rotated in four 45° increments (Figure 4b–e). By maintaining a two-lobe pattern at diagonal orientations of the polariser (Figure 4d and 4e), the emission is confirmed as being radially polarised and not an incoherent superposition of degenerate HG₀₁ and HG₁₀ transverse modes. For the four polariser images of Figure 4, the lowest polarisation purity value of the Ho:YAG source was measured to be $P = 0.96$, indicating a high-purity radially polarised output.

4 Gradient descent locking

Figure 5 shows a conceptual schematic of the process used to lock the laser onto the desired transverse mode. The output beam from the Ho:YAG laser is imaged onto two cameras via a 50/50 beamsplitter. The second camera is placed behind a linear polarisation filter in a rotation mount, allowing simultaneous imaging of the intensity profile and polarisation state of the laser. The camera images are then fed into the CNN, which outputs a vector defining the predicted modal composition.

If the laser emission has some presence of higher-order transverse modes with more complicated polarisation states, the vectorial output from the CNN will differ between the polarised and unpolarised camera views. For




Fig. 5 Feedback loop schematic for the stabilisation of mode purity with a convolutional neural network (CNN). The polarised and unpolarised beam profiles are analysed for modal composition, with the Ho:YAG intra-cavity lens adjusted to optimise the cavity for the chosen transverse mode.

instance in the case of the radially polarised LG₀₁ mode, the CNN will identify an LG₀₁ mode for the unpolarised camera and an HG₀₁ mode for the polarised camera (like those shown in Figure 4). Based on the output of the CNN, a direction of travel is sent to the stepper motor that moves the intra-cavity lens. From multiple passes through the control loop, the lens position is optimised for emission of the desired transverse mode.

To lock onto a particular mode, a one-dimensional gradient descent algorithm has been implemented. The CNN provides a vector output of the predicted modal content for the beam, \vec{F} , whilst the desired mode is defined by a separate vector, \vec{D} , which take the form:

$$\vec{F} = [F_{\text{HG}_{00}}, F_{\text{HG}_{11}}, F_{\text{LG}_{01}}, \dots, F_{\text{LG}_{nm}}, F_{\text{HG}_{nm}}] \quad (3)$$

$$\vec{D} = [D_{\text{HG}_{00}}, D_{\text{HG}_{11}}, D_{\text{LG}_{01}}, \dots, D_{\text{LG}_{nm}}, D_{\text{HG}_{nm}}]. \quad (4)$$

From this, we can define an error function as the length between the two vectors given as

$$E = \frac{\|\vec{D} - \vec{F}\|}{\sqrt{N}}, \quad (5)$$

such that

$$E = \frac{\sqrt{(F_{\text{HG}_{00}} - D_{\text{HG}_{00}})^2 + \dots + (F_{\text{LG}_{nm}} - D_{\text{LG}_{nm}})^2}}{\sqrt{N}}, \quad (6)$$

where N is the number of components in the vector. If there is a good match between the desired and obtained vectors, the error function (E) will tend to 0, which increases up to 1 as the mode quality deviates from the desired mode. When locking with the gradient descent algorithm, the error function is calculated for a series of neighbouring positions of the intra-cavity lens. The gradient of the error at these positions is then calculated and a step of the central position is undertaken

in the direction specified by the negative gradient. This process is then repeated until the error reaches an acceptable threshold, or can be left running if changes in the underlying parameter space are expected.

In this implementation, we are specifically interested in locking onto the radially polarised LG_{01} transverse mode. A single error vector from just one camera image would not be sufficient as the intensity distribution and the polarisation distribution must be fed into the locking algorithm to distinguish the radial, azimuthal and random LG_{01} polarisations. As such, an error signal from the polarised camera (E_{pol}) and the unpolarised camera (E_{unpol}) are combined as

$$E_{\text{tot}} = E_{\text{pol}} \cdot E_{\text{unpol}}. \quad (7)$$

Here, the intra-cavity lens position was dithered to take 5 equally spaced readings of the error function either side of the current central position, after which the gradient is calculated and a positional step made in the direction of reduced error. The process is then repeated on the new central position. The spacing of the dithered positions was set to be small compared to the distances over which the laser modal content would change, such that the dithering should not greatly affect the modal quality.

5 Experimental results

5.1 Without mode stabilisation

When the mode stabilisation system was not in use, the output purity varied quite significantly as the 1907 nm pump power was increased (red line in Figure 6 and Figures 7a–e). At pump powers below 28 W (but above threshold), the laser emission was a combination of the fundamental HG_{00} mode and the radially polarised LG_{01} mode (Figure 7a), dominated by the presence of HG_{00}




Fig. 6 Variation of the radial polarisation purity as a function of incident pump power when the machine learning stabilisation is, or is not, being used.




Fig. 7 Ho:YAG output beam profiles through a linear polarisation filter at a selection of 1907 nm pump powers (21–49 W) for the situations with: (a)–(e) no modal stabilisation in use, (f)–(j) radial LG_{01} stabilisation in use and (k)–(o) HG_{00} stabilisation in use.

and giving P values as low as 0.21. At pump powers close to 28 W, excellent radial polarisation purity up to $P = 1.00$ was achieved as the laser cavity had been initially aligned here. When the pump power increased beyond 28 W, the radial polarisation purity fell rapidly (down to $P = 0.54$ at 52 W of pump power) and the output was increasingly formed of the randomly polarised LG_{01} transverse mode due to uncompensated changes in the Ho:YAG thermal lensing.

At 28 W of pump power, the unstabilised laser output demonstrated M^2 values of 2.0 and 2.1 in the x and y directions respectively, whilst at 49 W the M^2 measurement yielded 2.1 and 2.0 in x and y respectively. The stable M^2 values across the pump power range suggest that the beam radius has remained largely consistent throughout operation of the laser, even when the active stabilisation system is not being used.

In Figure 8, the situation where no modal stabilisation was used has produced the highest output slope efficiency (67%) and the highest output power level (22 W) as the laser was allowed to adapt its modal content when the thermal lens strength was increasing, finding other higher-order transverse modes that can operate with greater gain extraction than the radially polarised LG_{01} mode.




Fig. 8 Output power curves for the Ho:YAG laser with either no machine learning stabilisation, radial LG_{01} stabilisation or fundamental HG_{00} stabilisation.

5.2 With radial LG_{01} stabilisation

When utilising the machine learning feedback system to provide stabilisation of the radially polarised LG_{01} transverse mode, the polarisation purity was maintained above $P = 0.86$ for the full range of pump powers (blue line in Figure 6). The maintenance of high quality radial polarisation can also be seen from images (f)–(j) in Figure 7. The radially-stabilised polarisation purity reached a maximum value of $P = 0.95$ across the full range of pump powers. While this maximum polarisation purity is lower than the unstabilised situation, we attribute this to the continuous dithering of the intra-cavity lens, causing the system to oscillate around the optimal value rather than locking onto it.

With increased pump power, the range of intra-cavity lens positions that support radially polarised emission is narrowed. As a result, it is preferable to have very fine movements of the intra-cavity lens at high pump powers to ensure that the optimal position is not missed and the gradient locking approach can still operate appropriately. To this end, it is the relatively large step size of our system that we believe to be the likely cause of the gradual reduction in polarisation purity, P , as a function of pump power in the case of active mode stabilisation. Even so, the beam propagation factor remained effectively unchanged when using radial stabilisation ($M_x^2 = 2.0$ and $M_y^2 = 2.1$ at 28 W of pump versus $M_x^2 = 2.0$ and $M_y^2 = 2.0$ at 49 W of pump). To this end, we conclude that the machine learning stabilisation system has maintained a stable beam radius across the full range of available pump powers.

5.3 With HG_{00} stabilisation

To explore the versatility of transverse mode stabilisation with the CNN developed here, the system was al-

tered to stabilise the laser output on the fundamental HG_{00} mode. Such a modification only required the desired mode vector, \vec{D} , to be changed (a single line of code), with no changes to the training data. In Figure 7(k)–(o), it can be seen that HG_{00} emission was maintained at all pump powers. At 49 W of pump power, the M^2 of the HG_{00} -stabilised beam was measured to be 1.3 and 1.4 in the x and y axes respectively.

The somewhat poor mode quality can be attributed to the large spatial mismatch between the HG_{00} cavity mode and the ring-shaped 1907 nm pump spot. Here, the pump spot is much larger than the HG_{00} transverse mode, producing diffraction losses of the cavity mode as well as gain extraction by the higher-order modes that have more spatial overlap with the ring-shaped inversion profile. Furthermore, as the HG_{00} transverse mode significantly overlaps with the central un-pumped region of the Ho:YAG crystal, the HG_{00} mode has reduced energy extraction efficiency, a lower output power (18 W) and a poorer slope efficiency (55%), as shown in Figure 8.

Despite the reduced beam quality observed during the HG_{00} stabilisation (a mode which the cavity is not optimised for), the system has shown the capability to lock onto an arbitrary transverse mode that it has been conditioned to identify from the training data set. Such modal flexibility and stabilisation is attractive for use in complex materials processing applications where an assortment of beam profiles and polarisation states can be used in quick succession for the highest throughput and minimum additional expense. For more complex vector beam profiles, it may be necessary to use more than two beam imaging cameras in order to view the beam through multiple linear polarisation filters orientated at different angles simultaneously, while also expanding the machine learning code and stabilisation loop to incorporate the additional data inputs.

6 Conclusion

Stress-induced birefringence in solid-state gain media is a simple method for generating radially polarised transverse laser mode emission, but is typically only optimised for excellent modal purity and polarisation purity across a narrow range of output powers.

Here, we have incorporated a real-time feedback system to automatically correct for degradation of the radial polarisation as a function of pump power. The output beam profile from a radially polarised Ho:YAG laser was imaged with and without a linear polariser, and analysed by an in-house-developed convolutional neural network (CNN). Based on the discrepancy between the laser modal composition and the desired mode of operation, the CNN automatically re-positioned an intra-cavity lens via a gradient descent algorithm.

Stabilisation of the radially polarised LG_{01} output has been demonstrated across a 30 W pump power range,

maintaining polarisation purity above $P = 0.86$ at all times. In the absence of transverse mode stabilisation, good polarisation purity was only achieved around one pump power before degrading into the randomly polarised LG_{01} mode, with P dropping to 0.54 as the pump power increased. The laser has also been stabilised on the HG_{00} mode across the full range of pump powers, following a simple modification to the stabilisation code. Utilising the rapid image processing capabilities that machine learning can achieve, the ability to generate high purity radially polarised beams at a range of output power levels without manual cavity re-alignment offers a new level of flexibility not currently available from radially polarised laser sources targeted at laser materials processing applications.

Supplementary material

The data underpinning this publication can be found in the University of Southampton repository at <https://doi.org/10.5258/SOTON/D2142>.

Acknowledgements

T. L. Jefferson-Brain acknowledges financial support from EPSRC (1921150). M. J. Barber acknowledges financial support from EPSRC (2115206) and Leonardo UK.

Disclosures

The authors declare no conflicts of interest.

References

1. L. Novotny, M. R. Beversluis, K. S. Youngworth, and T. G. Brown. Longitudinal field modes probed by single molecules. *Physical Review Letters*, 86(23):5251, 2001.
2. Y. Liu, D. Cline, and P. He. Vacuum laser acceleration using a radially polarized CO_2 laser beam. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 424(2–3):296–303, 1999.
3. M. Meier, V. Romano, and T. Feurer. Material processing with pulsed radially and azimuthally polarized laser radiation. *Applied Physics A*, 86(3):329–334, 2007.
4. V. G. Niziev and A. V. Nesterov. Influence of beam polarization on laser cutting efficiency. *Journal of Physics D: Applied Physics*, 32(13):1455, 1999.
5. A. C. Butler, R. Uren, D. Lin, J. R. Hayes, and W. A. Clarkson. Simple technique for high-order ring-mode selection in solid-state lasers. In *2015 Conference on Lasers and Electro-Optics Europe and European Quantum Electronics Conference*, page CA-7-5. Optical Society of America, 2015.
6. A. Liu, T. Lin, H. Han, X. Zhang, Z. Chen, F. Gan, H. Lv, and X. Liu. Analyzing modal power in multi-mode waveguide via machine learning. *Optics Express*, 26(17):22100–22109, 2018.
7. T. L. Jefferson-Brain, A. D. Coupe, M. D. Burns, W. A. Clarkson, and P. C. Shardlow. Alignment of higher-order mode solid-state laser systems with machine learning diagnostic assistance. In *2019 Conference on Lasers and Electro-Optics Europe and European Quantum Electronics Conference*, page CA-p-48. Optical Society of America, 2019.
8. Y. An, L. Huang, J. Li, J. Leng, L. Yang, and P. Zhou. Learning to decompose the modes in few-mode fibers with deep convolutional neural network. *Optics Express*, 27(7):10127–10137, 2019.
9. Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
10. T. Baumeister, S. L. Brunton, and J. N. Kutz. Deep learning and model predictive control for self-tuning mode-locked lasers. *J. Opt. Soc. Am. B*, 35(3):617–626, 2018.
11. A. Sinha, J. Lee, S. Li, and G. Barbastathis. Lensless computational imaging through deep learning. *Optica*, 4(9):1117–1125, 2017.
12. B. Rahmani, D. Loterie, G. Konstantinou, D. Psaltis, and C. Moser. Multimode optical fiber transmission with a deep learning network. *Light: Science & Applications*, 7(69):1–11, 2018.
13. Y. Rivenson, Z. Göröcs, H. Günaydin, Y. Zhang, H. Wang, and A. Ozcan. Deep learning microscopy. *Optica*, 4(11):1437–1443, 2017.
14. H. Tünnermann and A. Shirakawa. Deep reinforcement learning for coherent beam combining applications. *Optics Express*, 27(17):24223–24230, 2019.
15. J. W. Nicholson, A. D. Yablon, S. Ramachandran, and S. Ghalmi. Spatially and spectrally resolved imaging of modal content in large-mode-area fibers. *Optics Express*, 16(10):7233–7243, 2008.
16. J. Demas and S. Ramachandran. Sub-second mode measurement of fibers using C^2 imaging. *Optics Express*, 22(19):23043–23056, 2014.
17. Y. Z. Ma, Y. Sych, G. Onishchukov, S. Ramachandran, U. Peschel, B. Schmauss, and G. Leuchs. Fiber-modes and fiber-anisotropy characterization using low-coherence interferometry. *Applied Physics B*, 96(2–3):345–353, 2009.
18. T. Kaiser, D. Flamm, S. Schröter, and M. Duparré. Complete modal decomposition for optical fibers using CGH-based correlation filters. *Optics Express*, 17(11):9347–9356, 2009.
19. M. Lyu, Z. Lin, G. Li, and G. Situ. Fast modal decomposition for optical fibers using digital holography. *Scientific Reports*, 7(6556):1–9, 2017.
20. O. Shapira, A. F. Abouraddy, J. D. Joannopoulos, and Y. Fink. Complete modal decomposition for optical waveguides. *Physical Review Letters*, 94(14):143902, 2005.
21. R. Brüning, P. Gelszinnis, C. Schulze, D. Flamm, and M. Duparré. Comparative analysis of numerical methods for the mode analysis of laser beams. *Applied Optics*, 52(32):7769–7777, 2013.
22. L. Li, J. Leng, P. Zhou, and J. Chen. Multimode fiber modal decomposition based on hybrid genetic global

optimization algorithm. *Optics Express*, 25(17):19680–19690, 2017.

- 23. L. Huang, S. Guo, J. Leng, H. Lü, P. Zhou, and X. Cheng. Real-time mode decomposition for few-mode fiber based on numerical method. *Optics Express*, 23(4):4620–4629, 2015.
- 24. J. W. Kim, J. I. Mackenzie, J. R. Hayes, and W. A. Clarkson. High power Er:YAG laser with radially-polarized Laguerre-Gaussian (LG_{01}) mode output. *Optics Express*, 19(15):14526–14531, 2011.