

PAPER • OPEN ACCESS

Deep learning for the monitoring and process control of femtosecond laser machining

To cite this article: Yunhui Xie *et al* 2019 *J. Phys. Photonics* **1** 035002

View the [article online](#) for updates and enhancements.



PAPER

OPEN ACCESS

RECEIVED
29 January 2019REVISED
13 May 2019ACCEPTED FOR PUBLICATION
10 June 2019PUBLISHED
28 June 2019

Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](https://creativecommons.org/licenses/by/4.0/).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Deep learning for the monitoring and process control of femtosecond laser machining

Yunhui Xie , Daniel J Heath , James A Grant-Jacob , Benita S Mackay , Michael D T McDonnell , Matthew Praeger , Robert W Eason and Ben Mills

Optoelectronics Research Centre, University of Southampton, Southampton, SO17 1BJ, United Kingdom

E-mail: bm602@orc.soton.ac.uk

Keywords: laser machining, machine learning, real-time, depth detection

Abstract

Whilst advances in lasers now allow the processing of practically any material, further optimisation in precision and efficiency is highly desirable, in particular via the development of real-time detection and feedback systems. Here, we demonstrate the application of neural networks for system monitoring via visual observation of the work-piece during laser processing. Specifically, we show quantification of unintended laser beam modifications, namely translation and rotation, along with real-time closed-loop feedback capable of halting laser processing immediately after machining through a ~ 450 nm thick copper layer. We show that this approach can detect translations in beam position that are smaller than the pixels of the camera used for observation. We also show a method of data augmentation that can be used to significantly reduce the quantity of experimental data needed for training a neural network. Unintentional beam translations and rotations are detected concurrently, hence demonstrating the feasibility for simultaneous identification of many laser machining parameters. Neural networks are an ideal solution, as they require zero understanding of the physical properties of laser machining, and instead are trained directly from experimental data.

1. Introduction

Lasers and their applications are a global industry, with sales reaching \$11 billion in 2017, of which materials processing is valued at \$4 billion [1]. Laser-based processing is now widely applied to a range of materials processing challenges, such as the cutting of metals and plastics [2–6], the micro-patterning of ultrahard materials and medical devices [7–9], additive laser manufacturing [10–14], and laser cleaning and drilling [15, 16]. Often, in order to ensure that a particular process completes, the optimum exposure time or laser power is deliberately exceeded, erring on the side of caution in order to achieve a near-100% confidence level. This approach is inherently inefficient as well as causing possible damage to the work-piece through overheating, unintended material removal from an underlying layer, or compromising intended minimum feature sizes. By incorporating detection of process completion with feedback control, for example immediately ceasing laser machining when the material has been drilled through the entire thickness, both the efficiency and the precision of the laser processing can therefore be improved.

A further concern in laser machining is that instability in the laser output or in subsequent beam delivery optics may cause unwanted modifications to the beam position and shape. Ultimately, it is the shape of the laser beam at the work-piece that determines the quality of machining, and so a detection regime that relies on observation at this position will be of significant general applicability.

In this proof-of-principle approach, we demonstrate the application of neural networks for detection of beam translation and rotation, as well as for halting laser machining at precisely the point where a thin film is laser-machined through its entire depth. Both these demonstrations highlight the potential of feedback control for enabling improvements in both fabrication precision and reproducibility. These approaches are achieved via observation of the work-piece, as recorded by a camera during laser machining. Whilst detection of beam

position or rotation alone may be achievable with analytic methods, for example via phase-correlation algorithms [17], machine learning offers the potential for simultaneous detection of multiple laser machining parameters and hence offers a general and integrated solution to monitoring of laser machining processes. Predicting the point at which the entire depth of a thin film is machined through is challenging because of unexpected variations in material hardness, density, thickness, reflectivity and so on. An analytical approach based on the camera data alone would probably require a complex programmatical description of the appearance of debris and the color contrast during the machining process.

Laser control feedback is achieved in this work using machine learning, specifically neural networks (NNs), as this eliminates the need for an analytical description of the problem. NNs have the significant advantage of not needing to be programmed with a description of the underlying physical processes, as instead the neural network can be trained directly from the experimental data. Neural networks have used acoustic signatures for characterization of depth of weld penetration in laser welding [18, 19] and classification of melt-pool image in additive manufacturing [20]. Here, convolutional neural networks (CNNs) are used as they are particularly well suited to image analysis, as their architecture contains a hierarchy of convolutional processes that can identify the presence, or lack thereof, of specific features in an image [21]. CNNs have been widely used in areas such as medical diagnostics [22], language translation [23], pollution detection [24] and the development of AI opponents in computer games [25]. In relation to photonics, neural networks have enabled improvements in optical microscopy [26] and Ptychography [27], light scattering control through opaque media [28] and object classification through scattering media [29, 30], as well as for reconstructing ultrashort pulses, phase retrieval and holography [31, 32].

Machine learning has allowed for predictive control for self-tuning mode-locked lasers [33], and in our previous work machine learning has shown the application of CNNs to produce realistic and accurate depth profiles and surface appearance predictions that would result from femtosecond laser ablation of metals, despite the extremely nonlinear nature of the process [34, 35]. Other work with CNNs has shown identification of work-piece material, laser power, and number of pulses used to machine microstructures, directly from camera images of the work-piece during laser machining [35]. Here we demonstrate training data augmentation techniques to aid the detection of changes in beam translation and rotation, and real-time closed-loop feedback for efficient laser machining through thin films.

2. Experimental setup

150 fs, 800 nm wavelength, 1 mJ pulses from a Ti:sapphire amplifier were selected and routed to the experiment individually via a Pockels cell. A neutral density filter was used to attenuate the beam, followed by a spatial homogeniser (π -shaper model 6_6) that transformed the Gaussian spatial intensity profile into a top-hat spatial intensity that illuminated a central region of the mirror array on the digital micromirror device (DMD), which was acting as a binary spatial light modulator. The DMD was used to shape the spatial intensity profile of the laser pulses, as described previously [14, 36], in this case to an ellipsoid with ellipticity (ratio of major to minor axes) of a factor of two. In this manner, it was possible to artificially create undesired translation and rotation of the beam via the DMD, and an ellipsoid was chosen to demonstrate general applicability of the proposed technique to non-circularly symmetric beam shapes. The ellipsoidal spatially-shaped pulses were imaged onto the sample using a Nikon ELWD 50 \times microscope objective, where the major and minor axes of the ellipse-shaped intensity distribution at the sample were 14 μm and 7 μm respectively, resulting in an ablated structure with these dimensions. The minimum translation of the pattern on the DMD for this experiment setup, i.e. translating the image of the ellipse on the DMD by a single row of mirrors, corresponded to a translation of the position of the laser machined feature of 91 ± 15 nm on the surface of the sample, as measured via scanning electron microscopy. A dichroic mirror positioned above the objective allowed for real-time observation and recording of images of the sample during machining, via a CMOS camera (Thorlabs DC1545M). The sample was positioned in three dimensions via translation stages (Thorlabs LNR50S), with automated focal position corrections used to maintain the sample surface at the image plane [37]. Figure 1(a) shows a schematic of the real-time feedback loop, showing that the machined structures are imaged by the camera and that the camera images are then processed by the neural network, which subsequently provides monitoring of the transformation of machined structures and feedback loop that is capable of halting the laser in real time. The work here is split into two experiments. Firstly, as shown in figure 1(b), the DMD was used to simulate unintended modifications in the beam shape, namely translation and rotation. A set of neural networks was then used to determine these modifications for a series of random transformations, directly from the camera images of the machined sample. Secondly, as shown in figure 1(c), a neural network was used to provide real-time control of the machining of a thin film, specifically predicting the number of pulses remaining until breakthrough and stopping the machining at breakthrough, even though the number of pulses required was

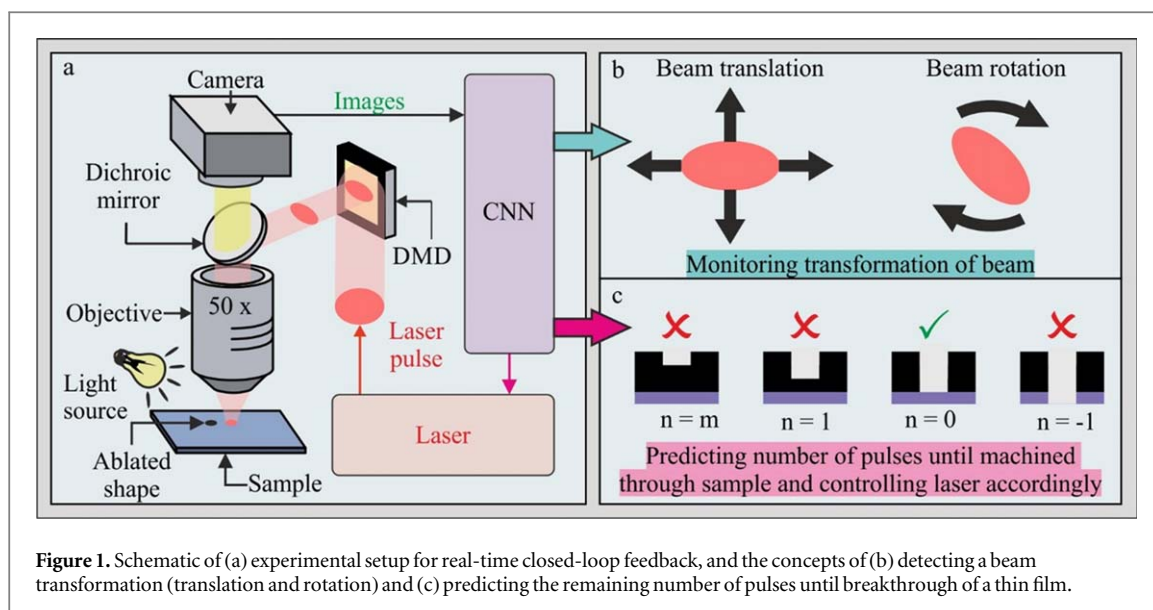


Figure 1. Schematic of (a) experimental setup for real-time closed-loop feedback, and the concepts of (b) detecting a beam transformation (translation and rotation) and (c) predicting the remaining number of pulses until breakthrough of a thin film.

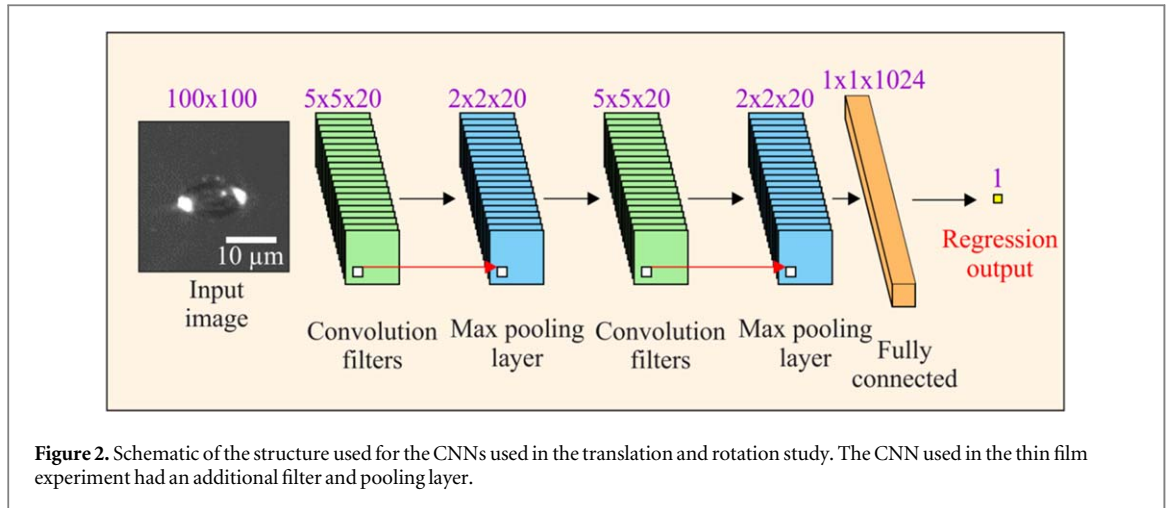
unknown for each trial. Although, in this work, the DMD was used to artificially introduce transformations in the beam that, the DMD was not integrated in the feedback loop. This was to demonstrate the capability for a CNN to monitor beam transformations directly from images of laser-machined structures. However, it is potentially feasible to build communication between the DMD and a CNN, hence enabling a real time feedback loop that can both halt the laser and correct for beam transformations caused by effects such as laser instability and experimental noise.

The target material for the beam translation and rotation experiment was a $5\ \mu\text{m}$ -thick electroless-nickel layer deposited on copper, and was laser processed with a fluence of $\sim 1.22\ \text{J cm}^{-2}$. The target material for the thin film machining experiment was a $\sim 450\ \text{nm}$ thick sputtered copper layer on silica, which was polished using abrasive paper to produce an uneven surface profile, and processed using a fluence of $\sim 1.83\ \text{J cm}^{-2}$. For the thin film processing, a non-translating and non-rotating elliptical beam profile was used for all exposures.

3. Neural networks

NNs are a computing paradigm that enable the representation of a complex transfer function, via a set of interconnected nonlinear functions [38]. Importantly, a neural network does not require a programmatical description of the physical processes underlying the transfer function, and instead the neural network can be trained directly from labelled experimental data [39]. In practice, this offers an alternative technique for solving problems that cannot easily be formulated in terms of analytical expressions or mathematical modelling problems. However, large amounts of labelled experimental data are generally needed, and hence the challenge often becomes the collection and labelling of suitable data. In this work, we use neural networks as transfer functions that can convert image data into numerical parameters that describe, for example, the rotation angle of the beam profile. In this case, once trained, such a NN has the capability to detect the rotation angle from camera images where the machined position differs from that used during training, and hence this enhances the robustness of the NN when applied to real-world laser processing problems.

The framework used for all CNNs in this work was Tensorflow, where each CNN comprised a series of convolutional layers, max pooling steps, and a final fully connected layer leading to the regression output. Each CNN took an input image size of 100×100 pixels, with layer parameters as shown in figure 2. In all cases, the activation function of each layer was rectified linear unit, the learning rate was 0.0001, and the optimiser was Adam [40]. The input images (100×100 pixels) were cropped from images recorded directly from the camera (native resolution 1280×1024 pixels). In the beam translation and rotation experiments two convolution layers were used with kernels of size 5×5 and 20 filters per layer. In the depth prediction experiments each convolution layer consisted of 64 filters with a kernel size of 3×3 . The convolutional layers were each followed by a max pooling layer with kernel size 2×2 and stride 2. Finally, a fully connected layer with 1024 units was used prior to the regression output.



4. Data augmentation

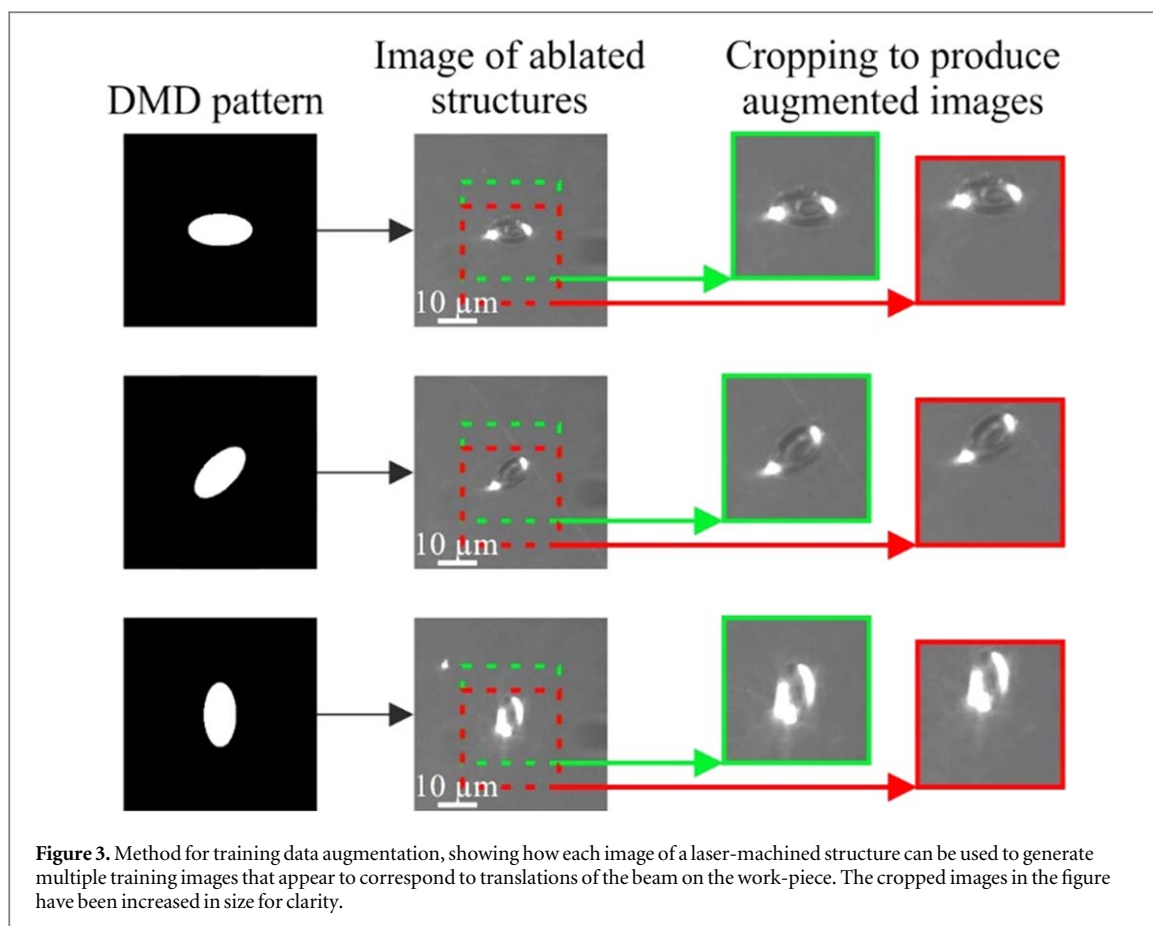
Successful training of a CNN generally requires a significant amount of labelled training data. Although the specific amount depends on many factors, tens of thousands of images are often regarded as a minimum. Therefore, collecting suitable training data can be a lengthy and challenging process, which in the case of laser machining would require the manufacturing of tens of thousands of machined structures. To alleviate the challenge of collecting such a large amount of training data, various methods for augmenting a smaller set of training data have been developed, such as stretching, cropping, translating, and even altering lighting and contrast levels [41, 42] in order to synthesise multiple variations of each original image. For this work, data augmentation took the form of cropping the images at different positions, hence offsetting the position of the laser-machined structure within each image. This was possible as the camera collected images at a resolution of 1280×1024 , whilst the input resolution of the CNN was set at 100×100 . An explanation is shown in figure 3, where the images of three laser-machined structures have been cropped, in order to give the appearance of a beam that has been translated on the work-piece. It is worth noting that as the size of the training data set increases, the time taken for a CNN to process the entire data set one time (commonly referred to as one iteration) is also increased. Therefore, there is a balance to be sought between the size of the data set and the training time and it is therefore of interest to determine the optimal degree of data augmentation.

5. Detection of single axis beam translation

An undesired shift in laser beam position relative to the target material could occur, for example, because of instability in the laser source or the beam delivery optics between the laser and the target material. As shown here, such a shifting can be identified via observation of the work-piece during laser machining. In this section, we show the effectiveness of a CNN to detect a one-dimensional translation in beam position when using different quantities of data augmentation.

For this work, the beam was translated from 0 to 100 DMD mirrors in steps of 1 (hence giving 101 positions), via changing the position of the elliptically shaped mask on the DMD by one mirror-width at a time. Due to the magnification within the experimental setup, a translation of four DMD mirrors approximately scaled to one camera pixel, hence providing the opportunity for sub-camera-pixel translation. Specifically, the translation of the elliptical pattern by the spacing of a single DMD mirror resulted in a translation of the position of the laser-machined structure by 91 ± 15 nm. This known translation was compared with the prediction from the neural network.

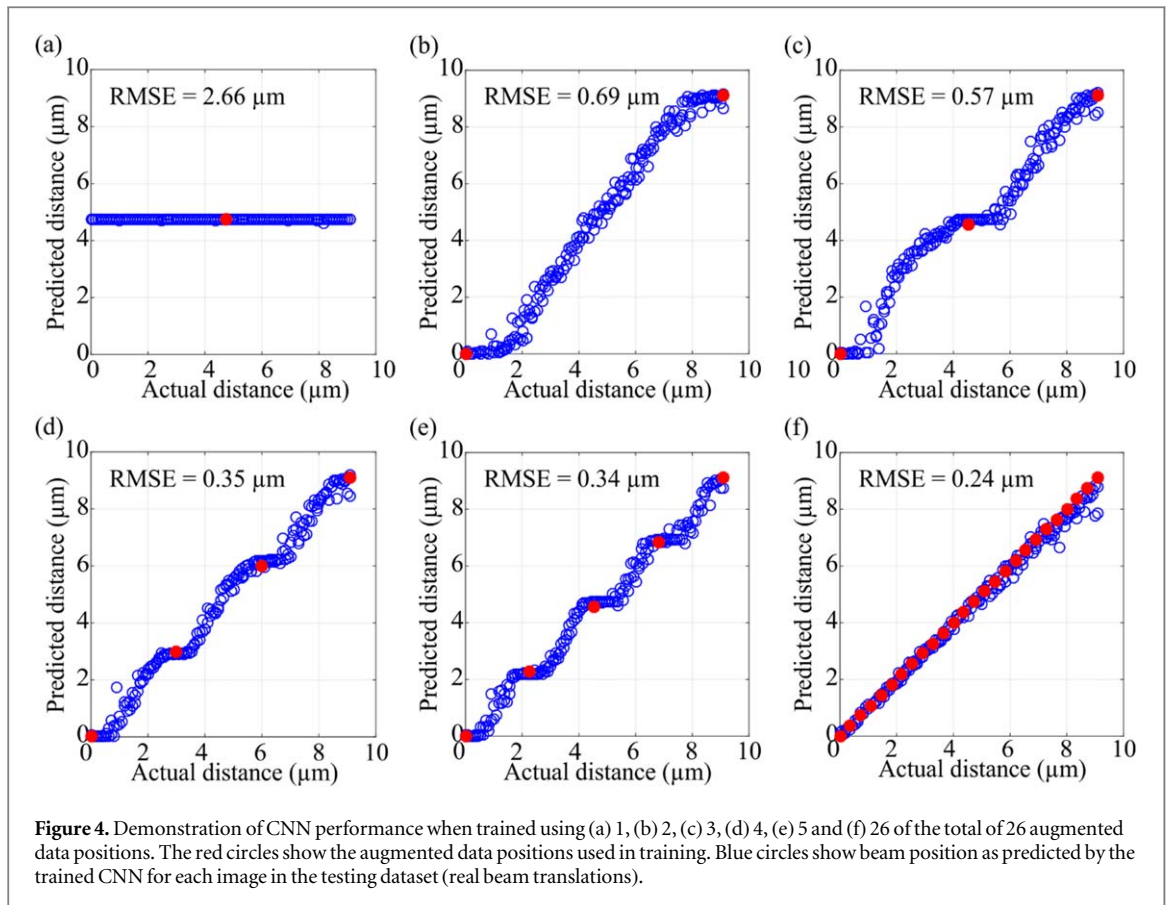
Firstly, 101 structures were machined, using the elliptical beam shape (i.e. using the same DMD pattern without any translation). These images were augmented in one dimension, by shifting the cropping window by an integer number of camera pixels, as illustrated in figure 3. The cropped images were shifted between 0 and 25 camera pixels, hence providing a total of $101 \times 26 = 2626$ images. Before training, 10% of these images were randomly chosen and were used to form the validation set, in which the neural networks would not observe during training. The remaining 90% of these images, were used to form the training set, and were repeatedly fed to the neural network during training. Secondly, a separate set of structures were laser-machined and imaged, in which the pattern on the DMD was translated, hence causing a real (i.e. not augmented) translation of the beam. A total of 202 images (two images for each of the 101 positions) were recorded, where this set of images is



referred to as the testing dataset. This was vital to prove that a CNN trained from purely augmented data could be used to recognise actual beam transformations. The CNN was therefore firstly trained on the training data in order to minimise the error on the validation data, then secondly tested on the testing data. The testing data was therefore not used during the training process.

To understand the benefit of training data augmentation, a series of CNNs was trained using differing numbers of the 26 augmented positions. Each training took 6 h to finish, and the trained CNN was then tested by feeding it images from the testing dataset (actual beam translations produced by the DMD). The results in figure 4 show the accuracy of each CNN at determining real-world beam translation distance when trained using (a) 1, (b) 2, (c) 3, (d) 4, (e) 5 and (f) 26 of the total of 26 augmented positions. The augmented positions are shown as red, and the testing data are shown as blue. The root mean square error (RMSE), corresponding to the average error in positional detection precision, is shown for the six cases. As observed in figure 4, data points tend to cluster around the points of augmentation, breaking the linearity of the scatter plot. The cause of this effect is likely related to the hypothesis that neural networks can be considered as polynomial regression models [43], and therefore their predictions will generally fit more strongly around clusters of data (i.e. in this case the augmented data) at the expense of more sparse regions of data. The results in figure 4 show that a CNN trained on purely augmented data, including those with small levels of data augmentation, is still capable of identifying real translations of features within an image, hence demonstrating the potential for a significant reduction in the amount of experimental data needed for training. Whilst the specific degree of augmentation required for a satisfactory level of precision will of course be dependent on the application, in this particular case we found that further improvements in precision were considerably smaller for five and higher numbers of augmented positions.

As observed in the figure 4(f), the prediction accuracy reduces for larger distance. A conjecture can be made that the causation of this phenomenon is related to the chosen structure of the neural network, which enabled a higher capability for prediction accuracy for smaller distances. It is important to realise that the images corresponding to the real translation had features that were not observed in the augmented training data, such as random debris positioning or variations in sample quality and laser power. Deliberately using the same images multiple times (i.e. augmentation) significantly reduced the difficulty of collecting a sufficient dataset. However, because the same images were used multiple times, the same experimental randomness in those images, such as debris position, sample quality and laser parameters, were also observed by the neural network multiple times,



hence producing a bias. Therefore, there is a trade-off between data collection and dataset bias. However, as shown here, despite the biased dataset, the neural network was capable of precisely determining the position. Each camera pixel corresponded to a distance of 360 nm on the surface of the target material. In the best case, the CNN was able to predict the distance to an accuracy of 240 nm, and hence below the camera imaging resolution. This result can be understood by considering the effect of translating an illuminating pattern on a camera by a distance corresponding to a fraction of a camera pixel. Although the image of the object would not have been translated by a single pixel, there would be some fraction of the image of the object that would have crossed into a different row of camera pixels [44], and this sub-pixel shift is believed to have enabled the sub-pixel capability of the CNN. It is worthwhile to mention that the error on the actual position measurement is likely related to the manufacturing precision of the DMD mirrors; the manufacturing error of machined structures caused by manufacturing precision of the DMD mirrors can be much smaller than one camera pixel in our case.

6. Detection of beam translation and beam rotation

To increase the complexity of the demonstration, the detection of beam translation in two dimensions (X and Y) and the rotation of the beam (0° – 179°) was investigated. Both the X and Y dimensions were augmented for all 26 camera positions, using the methodology in section 5. However, due to the directional nature of the white-light illumination on the sample, which resulted in shadows in one direction, augmenting the angle was not appropriate in this case. To enable the production of suitable training data, 180 structures were laser-machined and imaged, corresponding to one-degree rotations of the elliptical shape on the DMD. As the elliptical shape had rotational symmetry order 2, a half rotation was an equivalent test to a full rotation. These 180 images were augmented via cropping to give 26 different positions in both X and Y directions, hence giving a total of $180 \times 26 \times 26 = 121\,680$ images. 90% of these images were used to form the training dataset, and 10% were used for the validation dataset. As before, a separate set of structures was laser-machined and imaged in which the DMD pattern (and thus the laser beam) was simultaneously actually rotated and translated. This resulted in 227 images with a random distribution over all possible X and Y positions, and all possible rotations. This is referred to as the testing dataset. As before, each CNN was trained on the training dataset whilst minimising the error on the validation dataset. After 12 h of training for translation and 24 h of training for rotation, the trained CNNs were then applied to the testing dataset.

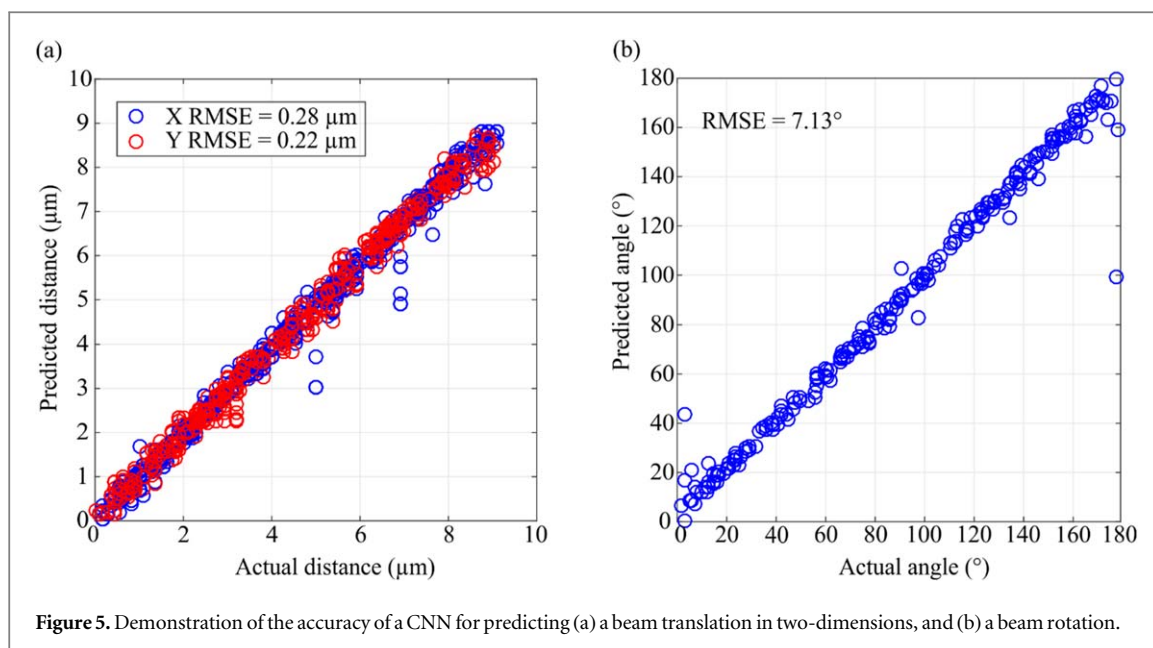


Figure 5. Demonstration of the accuracy of a CNN for predicting (a) a beam translation in two-dimensions, and (b) a beam rotation.

Figure 5 shows the precision for identification of (a) the X and Y position, and (b) the rotation, from the testing dataset. The X and Y axis prediction accuracy was 280 nm and 220 nm respectively, which is smaller than the 360 nm size of an imaged camera pixel. At the boundaries of 0° and 179° , angles were wrapped, resulting in an angular prediction accuracy of 7.13° . To identify the rotation angle of the beam, a minor addition to the CNN architecture was required. Here, two regression outputs were used in the final layer, with one trained to predict the sine of the rotation angle, and the other the cosine of the rotation angle [45]. This method of prediction was required in order to assign an appropriately low error when, for instance, an actual rotation angle of 1° resulted in a prediction by the CNN of 179° . In this case of rotational symmetry of order 2, this prediction would be close to the true rotation, but a direct comparison of angle would result in a very high error. Two sinusoids with an offset in phase were then required in order to uniquely define the predicted angle (as sine is a two-to-one mapping over 0° – 180°). It was found that training a separate CNN for each of the three parameters, rather than combining all three into a single CNN, was most efficient in terms of the total training time. In practice, a single image was sent simultaneously to each of the three CNNs, and three numbers were produced, namely the predicted X position, the predicted Y position, and the predicted rotation angle. Whilst here the results show the potential for simultaneous detection of the translation and rotation of a laser beam during machining, this technique could also be applied to the detection of other parameters such as laser pulse energy or laser mode quality. By measuring the total time for the three CNNs to process all 227 camera images in the testing dataset (i.e. predicting X , Y and angle), the time for processing a single image was calculated, at ~ 10 ms. This rapid processing speed is a fundamental advantage of using a CNN-based detection and is critical for enabling real-time monitoring of laser machining parameters.

7. Real-time control for thin film machining

In an industrial setting where quality control is paramount, when laser machining through a layer of material, a higher than expected total time of exposure is generally used to ensure an extremely low process failure rate. Variation in the required machining time can arise from the sample (e.g. uncertainty in thickness, surface quality, etc) and the uncertainty of the machining process itself (e.g. laser power variations, unexpected debris). Whilst excess machining time can ensure a lower wastage of material, the consequence will be a higher wastage of laser light energy. Ideally, a real-time closed-loop system would support the manufacturing process and halt the laser precisely when the layer of material is machined through. This experimental section describes the results of a CNN-based implementation that halts the laser when a ~ 450 nm copper film has been machined all the way through.

As discussed earlier, the copper film sample used in this experiment was deposited onto a silica substrate. Here, machining through the film refers to laser machining through the deposited copper and exposing the silica substrate underneath the copper film. This result therefore demonstrates not only the ability of machining through a layer of material but also the ability to machine through only a single layer of material in a multilayer

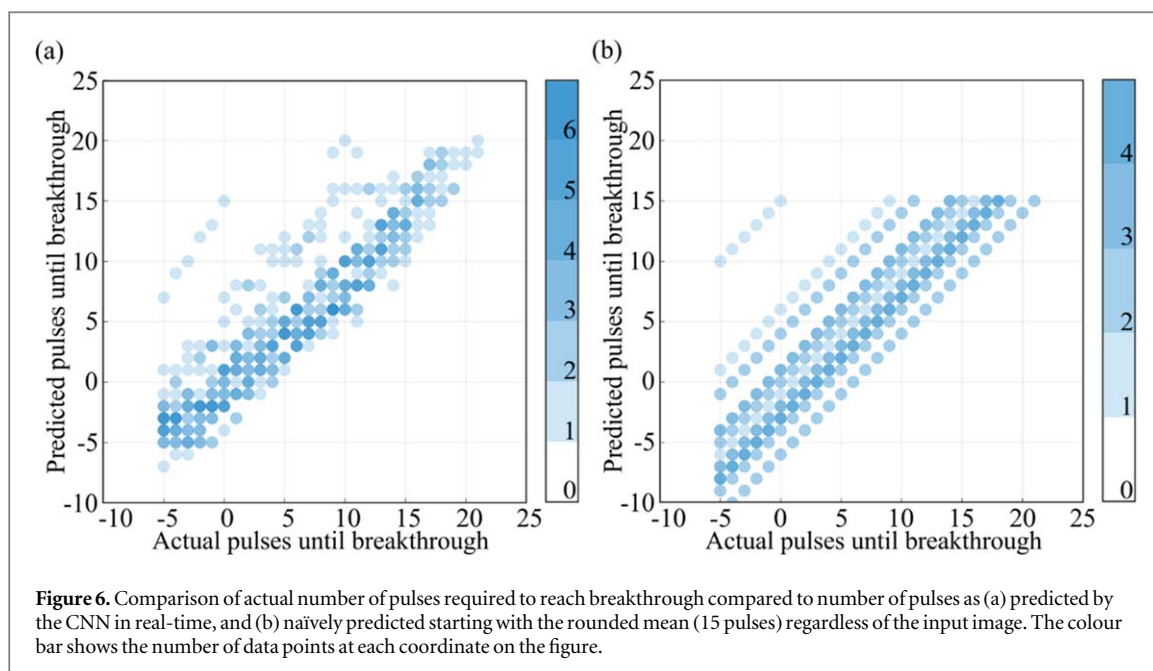


Figure 6. Comparison of actual number of pulses required to reach breakthrough compared to number of pulses as (a) predicted by the CNN in real-time, and (b) naively predicted starting with the rounded mean (15 pulses) regardless of the input image. The colour bar shows the number of data points at each coordinate on the figure.

structure. This could see applications in laser cleaning of surfaces, e.g. rust removal, where the thickness of the material is unknown but where laser damage to the second layer must be minimised.

In this experiment, a CNN-based feedback loop was implemented in the laser machining system which allowed the laser machining process to be stopped, in real-time, when the copper layer was machined through. After each incident laser pulse, a camera image of the sample was recorded and automatically sent to the CNN for processing. The CNN would then output a number corresponding to the predicted number of pulses remaining until the copper layer was machined all the way through. This number was then transmitted back to the automated laser machining system, which controlled the laser, therefore producing a feedback loop. If the number was positive, the laser machining would continue. If the number was zero or negative, the laser machining would halt, and a fresh, non-machined, region of the sample would be automatically found, in order to repeat the test. Here, a negative number of remaining pulses refers to using more pulses than was needed, for example, a prediction of ‘-1’ refers to the CNN predicting that 1 too many pulses had been used.

To introduce additional uncertainty in the thickness of the copper film, and hence uncertainty in the number of laser pulses required for machining, the surface of the copper film was coarsely polished. The mean in the number of pulses required to machine through the copper film could then be determined. The CNN would therefore have improved the efficiency of the process (by wasting fewer pulses) if its error in predicting the number of remaining pulses was smaller than the error if the value of prediction always started with the mean.

The training dataset was produced by exposing 50 sequential pulses at 156 well-separated positions on the copper film, and a camera image of the sample was taken after each pulse. Therefore, $50 \times 156 = 7800$ images were collected (10% of them were randomly chosen to be used for validation). The point of breakthrough for each position was determined through manual inspection and verified by checking the sample reflectivity and transmission using an optical microscope, as well as via an interferometric surface profiler measurement. Using this information, each of the images in the training dataset was labelled with the number of subsequent pulses required until breakthrough (i.e. +1 pulse, +2 pulses, +3 pulses, etc). The CNN training time for this dataset was 2 h. Once the CNN was trained, a new laser machining experiment, on a different region of the same sample, was conducted, where the CNN was used in real-time to control the experiment. In this real-time experiment, when predicting the number of remaining pulses corresponding to each camera image of the sample, the CNN only had access to a single image of the sample and hence each prediction was not influenced by any previous predictions.

Figure 6(a) shows the results for real-time prediction by the CNN of remaining pulses, for 21 positions on the copper sample. Each data point compares the predicted pulses based on a single image of the sample with the actual number of pulses remaining (determined manually after the experiment finished). The mean number of required pulses was 15.33, and the RMSE achieved by the CNN was 2.91. As a comparison, if the rounded mean value of pulses that was needed to machine through all 21 positions was used as the prediction at the start of the machining, as shown in figure 6 (b), the RMSE of this naïve guess was 3.78. Therefore, the CNN was able to predict the correct number of pulses $\sim 23\%$ more accurately in terms of RMSE than would be achieved by using the mean value of pulses appropriate for the average thickness across the film. Also of interest was the ability of

the CNN to make predictions of negative values, i.e. predicting the number of pulses that occurred after the point of breakthrough. We posit that the CNN had learned to recognise features in the edge-quality of the hole, or the distribution of surrounding debris, which may continue to change even after the exposed region had been machined away. For completeness, a separate region of the sample (16 positions) was used to successfully demonstrate halting of laser machining precisely at the breakthrough point, as controlled by the CNN-based feedback loop in real-time.

8. Conclusions

A set of CNNs that enable detection of various beam misalignments, as well as halting of laser machining immediately upon process completion, has been demonstrated. Firstly, the CNNs for monitoring of beam translation and rotation was demonstrated. Secondly, a CNN was applied in a real-time feedback loop, in order to halt laser machining at exactly the point when a thin film of copper was laser-machined through its entire depth. The results here demonstrate the potential for using neural networks to monitor and control laser-based material processing.

Acknowledgments

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) under grant numbers EP/N03368X/1 and EP/N509747/1. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan X GPU used for this research, donated through NVIDIA GPU Grant Program. Data supporting this manuscript is available at <https://doi.org/10.5258/SOTON/D0754>.

ORCID iDs

Yunhui Xie  <https://orcid.org/0000-0002-8841-7235>
Daniel J Heath  <https://orcid.org/0000-0003-3566-1813>
James A Grant-Jacob  <https://orcid.org/0000-0002-4270-4247>
Benita S Mackay  <https://orcid.org/0000-0003-2050-8912>
Michael D T McDonnell  <https://orcid.org/0000-0003-4308-1165>
Matthew Praeger  <https://orcid.org/0000-0002-5814-6155>
Robert W Eason  <https://orcid.org/0000-0001-9704-2204>
Ben Mills  <https://orcid.org/0000-0002-1784-1012>

References

- [1] Overton G, Noguee A, Belforte D and Holton C 2017 Annual laser market review & forecast: where have all the lasers gone? *Laser Focus World* **53** 1–24
- [2] Shin JS, Oh SY, Park H, Chung C M, Seon S, Kim T S, Lee L, Choi B S and Moon J K 2017 High-speed fiber laser cutting of thick stainless steel for dismantling tasks *Opt. Laser Technol.* **94** 244–7
- [3] Dutta Majumdar J and Manna I 2011 Laser material processing *Int. Mater. Rev.* **56** 341–88
- [4] Dutta Majumdar J and Manna I 2003 Laser processing of materials *Sadhana—Acad. Proc. Eng. Sci.* **28** 495–562
- [5] Dubey A K and Yadava V 2008 Laser beam machining—a review *Int. J. Mach. Tools Manuf.* **48** 609–28
- [6] Kumar Dubey A and Yadava V 2008 Multi-objective optimisation of laser beam cutting process *Opt. Laser Technol.* **40** 562–70
- [7] Mills B, Heath D J, Feinaeugle M, Grant-Jacob J A and Eason R W 2014 Laser ablation via programmable image projection for submicron dimension machining in diamond *J. Laser Appl.* **26** 041501
- [8] Zhang J, Tao S, Wang B and Zhao J 2017 Studies on nanosecond 532nm and 355nm and ultrafast 515nm and 532nm laser cutting super-hard materials *Proc. SPIE* **10091** 1009110
- [9] Gittard S D and Narayan R J 2010 Laser direct writing of micro- and nano-scale medical devices *Expert Rev. Med. Devices* **7** 343–56
- [10] Arnold C B, Serra P and Piqué A 2011 Laser direct-write techniques for printing of complex materials *MRS Bull.* **32** 23–31
- [11] Feinaeugle M, Heath D J, Mills B, Grant-Jacob J A, Mashanovich G Z and Eason R W 2016 Laser-induced backward transfer of nanoimprinted polymer elements *Appl. Phys. A* **122** 398
- [12] Gu D D, Meiners W, Wissenbach K and Poprawe R 2012 Laser additive manufacturing of metallic components: materials, processes and mechanisms *Int. Mater. Rev.* **57** 133–64
- [13] Khairallah S A, Anderson A T, Rubenchik A and King W E 2016 Laser powder-bed fusion additive manufacturing: Physics of complex melt flow and formation mechanisms of pores, spatter, and denudation zones *Acta Mater.* **108** 36–45
- [14] Heath D J, Feinaeugle M, Grant-Jacob J A, Mills B and Eason R W 2015 Dynamic spatial pulse shaping via a digital micromirror device for patterned laser-induced forward transfer of solid polymer films *Opt. Mater. Express* **5** 1129
- [15] Tam A C, Leung W P, Zapka W and Ziemlich W 1992 Laser-cleaning techniques for removal of surface particulates *J. Appl. Phys.* **71** 3515–23
- [16] Von Allmen M 1976 Laser drilling velocity in metals *J. Appl. Phys.* **47** 5460–3
- [17] Zitová B and Flusser J 2003 Image registration methods: a survey *Image Vis. Comput.* **21** 977–1000

- [18] Huang W and Kovacevic R 2011 A neural network and multiple regression method for the characterization of the depth of weld penetration in laser welding based on acoustic signatures *J. Intell. Manuf.* **22** 131–43
- [19] Günther J, Pilarski P M, Helfrich G, Shen H and Diepold K 2014 First steps towards an intelligent laser welding architecture using deep neural networks and reinforcement learning *Proc. Technol.* **15** 474–83
- [20] Kwon O, Kim H G, Ham M J, Kim W, Kim G H, Cho J H, Kim N I and Kim K 2018 A deep neural network for classification of melt-pool images in metal additive manufacturing *J. Intell. Manuf.* **2018** 1–12
- [21] Krizhevsky A, Sutskever I and Hinton G E 2012 ImageNet classification with deep convolutional neural networks *Proc. 25th Int. Conf. on Neural Information Processing Systems-NIPS'12* vol 1 (Red Hook, NY: Curran Associates) pp 1097–105
- [22] Milletari F, Navab N and Ahmadi S-A 2016 V-net: fully convolutional neural networks for volumetric medical image segmentation *2016 Fourth International Conference on 3D Vision (3DV)* (Piscataway, NJ: IEEE) pp 565–71
- [23] Kalchbrenner N, Grefenstette E and Blunsom P 2014 A Convolutional Neural Network for Modelling Sentences *Proc. 52nd Annual Meeting of the Association for Computational Linguistics* vol 1 pp 655–65
- [24] Grant-Jacob J A, Mackay B S, Baker J A G, Heath D J, Xie Y, Loxham M, Eason R W and Mills B 2018 Real-time particle pollution sensing using machine learning *Opt. Express* **26** 27237–46
- [25] Silver D et al 2016 Mastering the game of Go with deep neural networks and tree search *Nature* **529** 484–9
- [26] Rivenson Y, Göröcs Z, Günaydin H, Zhang Y, Wang H and Ozcan A 2017 Deep learning microscopy *Optica* **4** 1437–43
- [27] Nguyen T, Xue Y, Li Y, Tian L and Nehmetallah G 2018 Deep learning approach for Fourier ptychography microscopy *Opt. Express* **26** 26470–84
- [28] Turpin A, Vishniakou I and Seelig J D 2018 Light scattering control in transmission and reflection with neural networks *Opt. Express* **26** 30911–29
- [29] Satat G, Tancik M, Gupta O, Heshmat B and Raskar R 2017 Object classification through scattering media with deep learning on time resolved measurement *Opt. Express* **25** 17466–79
- [30] Li Y, Xue Y and Tian L 2018 Deep speckle correlation: a deep learning approach toward scalable imaging through scattering media *Optica* **5** 1181–90
- [31] Wu Y, Rivenson Y, Zhang Y, Wei Z, Günaydin H, Lin X and Ozcan A 2018 Extended depth-of-field in holographic imaging using deep-learning-based autofocusing and phase recovery *Optica* **5** 704–10
- [32] Rivenson Y, Zhang Y, Günaydin H, Teng D and Ozcan A 2018 Phase recovery and holographic image reconstruction using deep learning in neural networks *Light Sci. Appl.* **7** 17141
- [33] Baumeister T, Brunton S L and Kutz J N 2018 Deep learning and model predictive control for self-tuning mode-locked lasers *J. Opt. Soc. Am. B* **35** 617–26
- [34] Heath D J, Grant-Jacob J A, Xie Y, Mackay B S, Baker J A G, Eason R W and Mills B 2018 Machine learning for 3D simulated visualization of laser machining *Opt. Express* **26** 4984–8
- [35] Mills B, Heath D J, Grant-Jacob J A and Eason R W 2018 Predictive capabilities for laser machining via a neural network *Opt. Express* **26** 17245–53
- [36] Heath D J, Grant-Jacob J A, Feinaeugle M, Mills B and Eason R W 2017 Sub-diffraction limit laser ablation via multiple exposures using a digital micromirror device *Appl. Opt.* **56** 6398–404
- [37] Heath D J, Mackay B S, Grant-jacob J A, Xie Y, Oreffo R O C, Eason R W and Mills B 2018 Closed-loop corrective beam shaping for laser processing of curved surfaces *J. Micromech. Microeng.* **28** 127001
- [38] Hornik K, Stinchcombe M and White H 1989 Multilayer feedforward networks are universal approximators *Neural Netw.* **2** 359–66
- [39] Hecht-Nielsen R 1992 Theory of the backpropagation neural network *Neural Networks for Perception* (Amsterdam: Elsevier) pp 65–93
- [40] Kingma D P and Ba J 2014 Adam: a method for stochastic optimization arXiv:1412.6980
- [41] Ding J, Chen B, Liu H and Huang M 2016 Convolutional neural network with data augmentation for SAR target recognition *IEEE Geosci. Remote Sens. Lett.* **13** 364–8
- [42] Hernández-García A and König P 2018 Do deep nets really need weight decay and dropout? arXiv: 1802.07042
- [43] Cheng X, Khomtchouk B, Matloff N and Mohanty P 2018 Polynomial Regression As an Alternative to Neural Nets arXiv:1806.06850
- [44] Khademi S, Darudi A and Abbasi Z 2010 A sub pixel resolution method arXiv:1211.2221 643–6
- [45] Hara K, Vemulapalli R and Chellappa R 2017 Designing Deep Convolutional Neural Networks for Continuous Object Orientation Estimation arXiv:1702.01499