

# Live imaging of laser machining via plasma deep learning

JAMES A. GRANT-JACOB,<sup>1,\*</sup> BEN MILLS,<sup>1</sup> AND MICHALIS N. ZERVAS<sup>1</sup>

<sup>1</sup>*Optoelectronics Research Centre, University of Southampton, SO17 1BJ, UK*

*\*J.A.Grant-Jacob@soton.ac.uk*

**Abstract:** Real-time imaging of laser materials processing can be challenging as the laser generated plasma can prevent direct observation of the sample. However, the spatial structure of the generated plasma is strongly dependent on the surface profile of the sample, and therefore can be interrogated to indirectly provide an image of the sample. In this study, we demonstrate that deep learning can be used to predict the appearance of the surface of silicon before and after the laser pulse, in real-time, when being machined by single femtosecond pulses, directly from camera images of the generated plasma. This demonstration has immediate impact for real-time feedback and monitoring of laser materials processing where direct observation of the sample is not possible.

## 1. Introduction

Lasers are used widely across manufacturing [1], and have rapidly become the standard technique for applications such as cutting [2,3], marking [4–6], welding [7,8], deposition [9–11], and 3D printing [12–14]. However, the physical processes inherent to laser materials processing mean that the technique is generally highly nonlinear, and hence even small changes in the underlying conditions can lead to a significantly different manufacturing outcome. There is therefore great interest in the development of methods for the automation and control of lasers in manufacturing via real-time feedback mechanisms, with the goal of improving speed, accuracy, and efficiency. However, laser materials processing can result in the creation of a plasma that prevents the direct observation of the sample during machining [15,16], hence making real-time feedback considerably more complicated. There is therefore clear motivation for developing methods for indirect imaging of the sample.

Deep learning is a subcategory of machine learning that has rapidly gained popularity in recent years, due to its ability to identify structure in data that might otherwise have been too complicated to identify via traditional numerical and algorithmic techniques. Deep learning therefore offers the potential for a data-driven approach to scientific research, where solutions to complex problems can be identified automatically through the processing of large amounts of experimental or simulated data. The word “deep” refers to the use of neural networks that have multiple layers, which unlocks the ability to progressively extract higher-levels of feature abstraction in the input data [17–19].

Convolutional neural networks (CNNs) have been applied to a wide range of laser applications, including spectroscopy [20], particulate sensing [21], laser welding [22], monitoring laser ablation [23–26], laser powder bed fusion [27], monitoring audio information from laser machining [28,29], and classifying laser melt pools [30]. Conditional generative adversarial neural networks (cGANs) [31], which can be used for image-to-image transformations [32], have seen application in microstructure prediction of laser sintering [33], generating images in laser welding [34] and interference patterns from fibres [35], and modelling the outcome of laser machining [36–38], including topographical predictions for fibre laser cutting of steel [39].

Previous work has demonstrated the use of plasma sensing for predicting the pulse energy and the size of laser machined craters [28], for predicting the surface morphology of laser

machined silica due to machining [40], and for real-time control for preventing machining off the edge of material boundaries [41]. In this work, we show how deep learning can be applied in real-time to produce live images of the machined surface, through the application of femtosecond laser machining of silicon. Critically, we show that analysis of the network predictions provides evidence that the generated plasma is only correlated with the sample morphology before the laser pulse is incident, and that to predict the appearance of the sample after machining, the neural network creates an internal model of laser machining. All results shown here correspond to real-time predictions.

## 2. Experimental methods

### 2.1 Setup

Figure 1 displays a concept of the experimental setup. A Light Conversion Pharos SP laser was used to generate 190 fs, 1 mJ pulses with a central wavelength of 1030 nm. The pulses were focused onto the surface of a silicon sample using a Nikon 20× objective (TU Plan ELWD, 0.40 NA) to a spot size of  $\sim 30 \mu\text{m}$ . The sample was a  $\sim 1 \text{ cm}^2$  piece of p-type (100) silicon that had been glued to a  $0.75 \text{ mm} \times 25 \text{ mm} \times 75 \text{ mm}$  borosilicate slide. The sample was attached to Zaber XYZ motorized translation stages (LSM050A-T4) with a maximum travel distance of 5 cm, which enabled automated translation of the sample relative to the laser focus. The surface of the silicon was imaged along the laser axis using a Basler acA4112-20uc camera (1914  $\times$  1200, RGB). The emitted plasma was imaged using a Basler daA1920-160uc camera (4096  $\times$  3000, RGB) coupled with an Olympus 50× long working distance objective (SLMPLN, 0.35 NA) that was oriented perpendicular to the laser axis. Single pulses from the laser were triggered using Python software using a REST interface. A Microsoft Windows 10 workstation with an NVIDIA Titan Xp (12 GB), an Intel® Core™ i7-7700 CPU @ 3.60 GHz 3.60 GHz and GB 64 GB RAM was used to automate the experimental setup and run both neural networks used for this work.

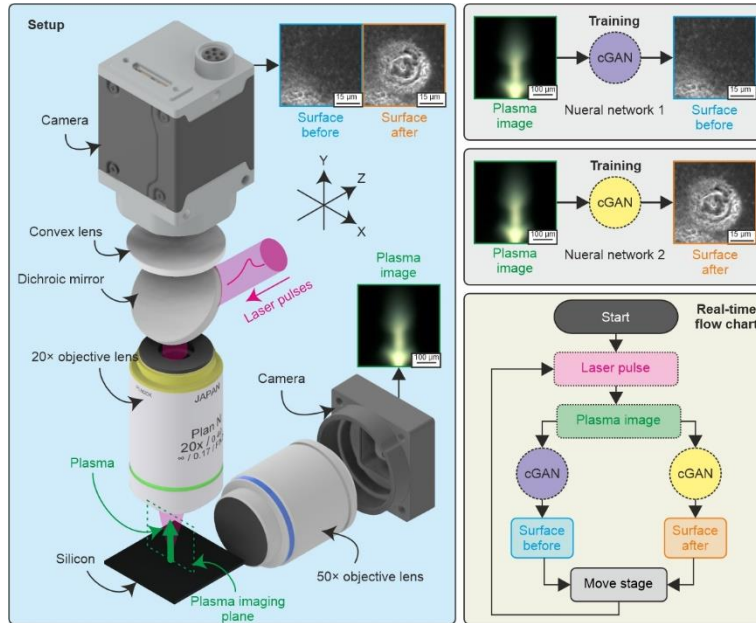


Fig. 1. Schematic of the experimental setup along with an example set of experimental plasma images and associated experimental images of the laser machined sample before and after the laser pulse. For this work, the two neural networks were run in real-time, hence providing a live image of the sample during machining.

## 2.2 Data collection and processing

The experimental automation, which combined stage translation, recording of images from the two cameras, a shutter connected to a white light source, and triggering of single laser pulses, was written in Python. To collect training data, camera images of the sample were recorded before and after laser pulses, and images of the plasma was recorded during the pulses. The white light source was blocked when recording plasma images. Images of the surface were recorded with an integration time of 500 ms, and the images of the plasma were recorded with an integration time of 250 ms. The long integration time on the plasma image was to compensate for the random latency resulting from communication via the REST server when triggering a single pulse, where this latency issue could be reduced by using an external signal that triggers the laser pulse and the camera simultaneously. After each set of before, during and after images were recorded, the sample was translated by a random distance and direction, in the approximate range of 10-30  $\mu\text{m}$ , to ensure that cases where subsequent pulses overlapped with the position of previous pulses were included in the training data. A total of 4326 sets of image data were collected (i.e., before, during, after), and all images were cropped and resized down to  $256 \times 256 \times 3$  pixels.

## 2.3 Neural networks and training

Two separate cGANs, which were based on the “pix2pix” architecture [32], were used to transform images of plasma into images of the surface of the silicon sample. The first network (network 1) was trained to transform a plasma image into an image of the surface before laser machining, and the second network (network 2) was trained to transform a plasma image into an image of the surface after laser machining, as shown by the blue and orange boxes in Fig. 1. Both neural networks had a plasma image as the input, and a prediction of the sample surface as the output, with all images of size  $256 \times 256 \times 3$ . The U-Net architecture of the generator consists of two paths: a contracting path, depicted in mauve, and an expansive path, illustrated in green (see Fig. 2(a)). The down-sampling convolutions are represented by orange arrows, while the up-sampling convolutions are denoted by grey arrows. Skip connections, shown as blue arrows, combine the feature maps from the expansive path with those from the corresponding layer in the contracting path. The multi-channel feature maps are symbolized by coloured rectangular boxes, with the dimensions of each map indicated inside the box and the number of channels specified below. During training, the generated images were automatically compared to the associated experimental image using a discriminator network. The discriminator’s task is to classify whether an image is real (i.e., from the training set) or fake (i.e., created by the generator). If the generated image is different from the real image, the discriminator may be able to judge that the image is fake and subsequently will output a value close to 0. This result is then used to calculate the loss for both the generator and the discriminator. The generator loss is calculated based on how successfully it fooled the discriminator. The discriminator’s loss, on the other hand, is based on how accurately it classified the images. The loss function for the discriminator was the sigmoid cross-entropy. These loss values are then used to adjust the model parameters (including the convolutional filters in both the generator and discriminator) via backpropagation. This process is repeated many times over the epochs until the generator becomes so accurate at creating images that the discriminator cannot distinguish between the real and the fake images. The generator loss for the neural networks predicting the (b) before and (c) after images is shown in Fig. 2. The root-mean-square-error (RMSE) of the predicted images compared with actual was 43.4 for before and 42.3 for after, with the maximum RMSE being 79.0 for before and 68.9 for after. The RMSE values were calculated using a built-in RMSE function in Matlab, which compares the pixel values of each image element-wise and returns the square root of the average of the squared differences. The mean of all the RMSE values for each pair of images (experimental and predicted) were then calculated.

The combined loss, as shown in Figure 2(b-c), represents the sum of the L1 and the adversarial loss. The adversarial loss measures the ability of the generator to fool the discriminator into judging that a generated image is real, and the L1 loss measures the pixel-wise mean absolute error between an experimental and predicted image. Therefore, the combined loss conveys a complex picture of convergence. In this work, the training was stopped at 100 epochs, which met the requirements for supporting the hypothesis that deep learning could be applied for real-time imaging of laser machining. To confirm that no further improvements in predictive accuracy occurred after 100 epochs, Figure 2(d) presents an example of predictions for 100, 150, 200, and 250 epochs, with the L1 average error loss and standard deviation (of all test samples) labelled on the images.

Both networks were trained using the same parameters, including a minibatch size of 2, a generator and discriminator learn rate of  $2 \times 10^{-4}$ , an L1-to-GAN loss ratio of 100:1, an ADAM optimizer, 100 epochs, and taking approximately 10.5 hours. The generator and discriminator learning rates, the L1-to-GAN loss ratio, and the ADAM optimiser, were chosen based on previous work that was found to produce accurate prediction of images [40]. Each neural network was trained on a total of 4132 sets of image data, and 194 sets of data were used to validate the accuracy of the neural networks before real-time implementation. No data augmentation was carried out on the datasets. MATLAB was used to train both neural networks, using a Microsoft Windows 10 computer workstation with an Intel® Xeon® Gold 5222 CPU @3.80 GHz and 3.79 GHz, and 192 GB RAM. The workstation was equipped with 3× NVIDIA A4050s, each with a capacity of 20 GB.

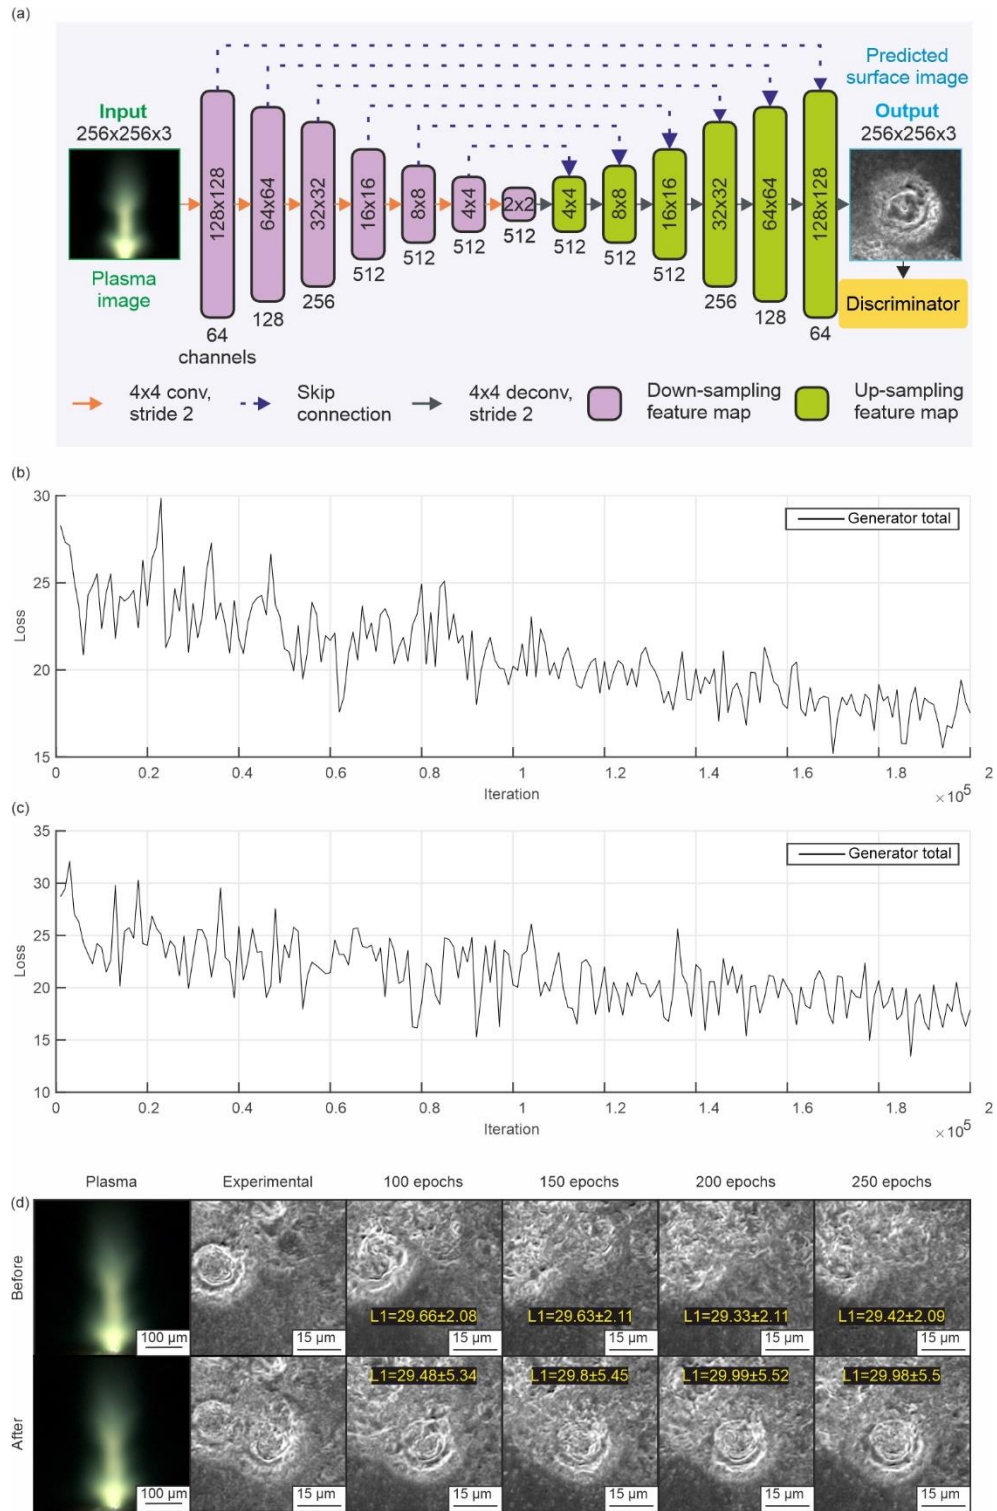


Fig. 2. (a) Schematic of the U-net architecture used for the generator for both neural network models used in this work. Loss for the generator for predicting the (b) before and (c) after images during the training process. There were 2000 iterations per epoch. An example of (d) plasma and

corresponding experimental and predicted images before and after ablation for 100, 150, 200 and 250 epochs, with the average of all test data L1 losses labelled on the images.

## *2.4 Real-time implementation*

The day after the neural networks were trained, the network weights were transferred to the computer workstation that operated the experimental automation. Each neural network was executed in a separate MATLAB environment, running in an infinite loop, and waiting for plasma images to be saved in the target folder. Upon a new plasma image being saved, the two neural networks immediately processed the image and saved the two output images, corresponding to the predictions for the appearance of the sample before and after the laser pulse, into another folder. In this case, due to a range of latency sources, the time taken between saving a plasma image and generating the images of the surface was approximately 1 second. Through a specifically designed automation architecture, including dedicated hardware and a refined and smaller neural network, this time could be likely reduced towards the tens of milliseconds, with the key bottlenecks needing to be solved likely being data transfer of camera data and neural network inference.

## **3. Results and discussions**

Figure 3 shows a schematic of the process for using a single image of the generated plasma for real-time prediction of the appearance of the sample before and after the laser pulse. The plasma image was used as input to both neural networks, and the two predicted images were generated. To demonstrate the prediction accuracy for this example, the associated experimental images are also shown in the figure, with the final row showing all four images masked at the spatial extent of the laser pulse, with the results showing strong agreement. Size scales are not included on the predicted images of the sample surface.

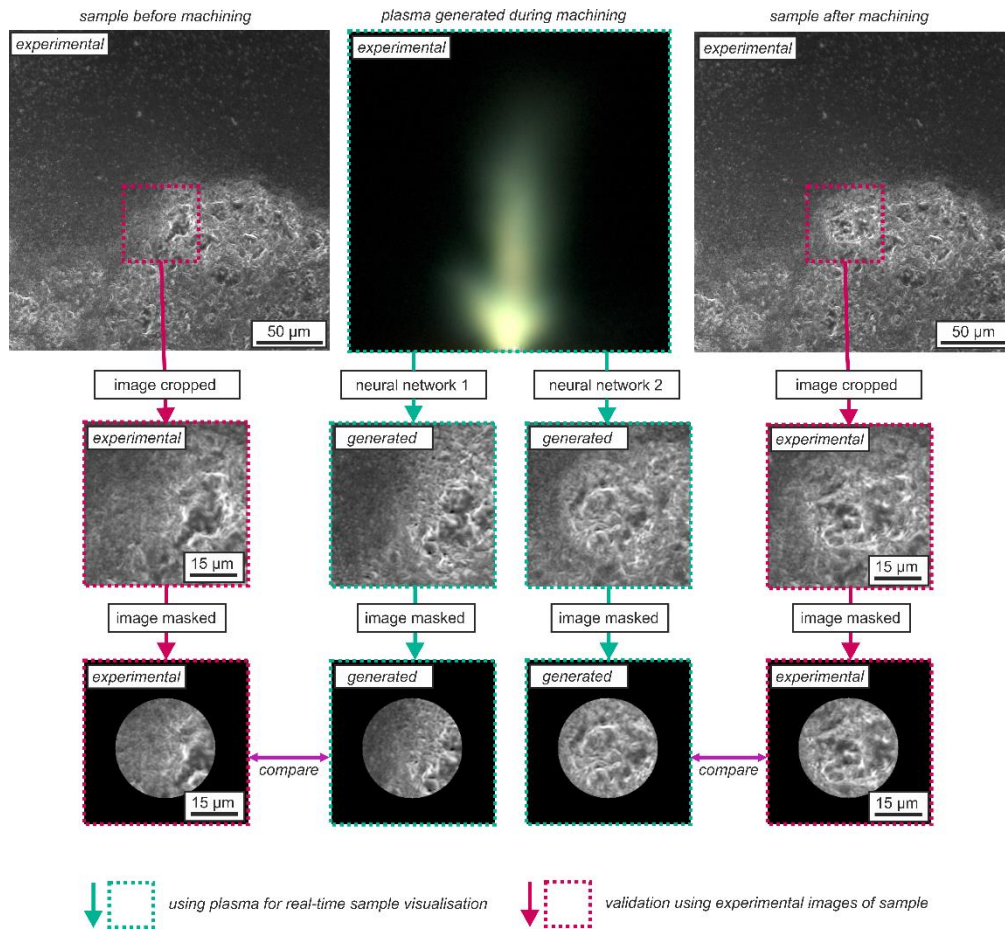


Fig. 3. A single example of a real-time prediction, with a comparison to the associated experimental result, shown as a process flowchart. Neural network 1 predicts the appearance of the sample before the laser pulse, and neural network 2 predicts the appearance of the sample after the laser pulse.

Figure 4 shows the results for ten consecutive pulses, and presents the generated plasma, and the associated experimental and predicted images for before and after the laser pulse, for both full images and masked images. The orange squares and dotted lines illustrate the relative position of the sample between subsequent pulses for the first few cases, showing that the sample was translated by a random distance and direction, in the approximate range of 10-30 µm. The neural network predictions are generally very similar to the experimental results, and the presentation of ten sequential pulses in the figure highlights the robustness and reliability of this technique on a real experimental setup.






Fig. 4. Ten sequential laser pulses and the associated experimental and generated before and after images of the sample, with and without masking of the region corresponding the spatial extent of the laser pulse. Pulse 10 in this figure was used for the Fig. 3 schematic.

The ability to image a 2D surface from a single perpendicular projection (i.e., only requiring a single plasma image to predict the appearance of the sample surface) is possibly due to the spatial distribution of the imaged plasma being related to the integration along the imaging axis of the interference effects resulting from the surface modulations on the sample. This would imply that each point on the plasma image contains information about the appearance of multiple regions on the sample. Indeed, previous work has shown that a neural network trained to identify the laser pulse energy directly from an image of the associated plasma tends to focus on the more strongly varying regions of the generated plasma, rather than treating all parts of the plasma image equally (see Fig. 3 in [40]). For reference, and to illustrate the variation of plasma images, Fig. 5 shows 100 examples of plasma images from sequential laser pulses.



The plasma is a result of laser ablation, and hence the shape and structure of the plasma contains information related to the sample surface at the position where the laser pulse is incident on the sample. However, an interesting observation is that the neural networks in this work can also predict, in many cases, the appearance of the sample outside this region. Given that there is no information regarding this outer region in the plasma image, this observation is attributed to a neural network extrapolation (i.e., a guess) based on the appearance of the sample within the spatial extent of the laser pulse. In other words, the neural network attempts to predict the appearance of the outer region of the sample, using the predicted appearance of the inner region of the sample. The experimental data for this work were collected by moving the translation stages by a uniform random distance and angle between laser pulses for ~250 pulses, before the sample was translated to an unmachined region. As a result, there was therefore some degree of correlation between the appearance of the inner region of the sample and the outer regions, which the neural networks appear to have learnt.




Fig. 5. One hundred examples of experimental plasma images, taken from sequential laser pulses, with the pulse number and scale bar included in each image.

Figure 6 shows a comparison of the average differences between the experimental and predicted results, for the before and after cases. The average is determined by mean absolute difference between all of 1004 of the experimental and predicted test images for (a) before and (b) after a laser pulse is incident on the surface. The prediction error in the central region for (b) the “after image” is notably higher than the prediction for (a) the “before image”, implying that the neural network that predicts the “after image” is not provided with the required

information to describe the appearance of the sample after the laser pulse. This observation could be explained by considering the timescales involved in femtosecond laser materials processing, where the plasma is a result of sample ionisation and hence a femtosecond time scale event, but where the subsequent formation of the modulations on the silicon surface due to melting of the surface are longer than a femtosecond time scale. As both neural networks used the same architecture, the same amount of data, and were trained for the same amount of time, it is unlikely that the difference in prediction accuracy is related to the networks themselves. If this is true, then a likely consequence is that the neural network tasked with predicting the “after image” is actually predicting the “before image” and then simulating the effect of femtosecond machining on this predicted “before image”, all in a single step. The use of neural networks for simulating the effect of laser machining has been demonstrated previously [36,37,42] and hence such results could support this hypothesis. The figure shows additional evidence, as (d) the average difference in the predicted “before images” and “after images” is very similar in spatial distribution and magnitude to (c) the average difference in the experimental “before images” and “after images”. It seems therefore plausible that the network may be applying a general set of learnt rules that can predict the distribution (but not the exact positions) of the surface modulations after melting, as the information describing the positions of the surface modulations after machining does not exist in the plasma images. For reference, the figure also shows the average pixel values for (e) experimental before, (f) experimental after, (g) predicted before, and (h) predicted after.

Whilst the (a) difference between the experimental “before image” and the predicted “before image” confirms that the neural network prediction error is generally smaller inside the region where the pulse is incident, it is important to note that the network is also able to predict the appearance of the sample outside this region. Given that sample is generally not modified outside this region, as indicated by (c), it is possible that the neural network is predicting the area in the outer region through a statistical extrapolation of the appearance of the sample in the inner region. This proposed extrapolation could also explain why (b) shows a prediction error that is lower in the outer region as compared to the inner region, as it may be that the neural network trained to predict the “after image” is firstly predicting the appearance of “before image”, then extrapolating this information to predict the outer region in the “after image”. If this hypothesis is true, then this result provides an application of a neural network for identification of some of the time scales associated with the process of femtosecond laser machining, exclusively from camera images that have longer than millisecond integration times.

To provide further evidence of this hypothesis, a third neural network was trained, to transform predicted before images into predicted after images, as shown by the flowchart in Fig. 7(a). The flowchart shows the difference between the “direct route”, where the plasma image is directly transformed into a prediction for the appearance of the sample after machining, and the “indirect route”, where the appearance of the sample before machining is predicted from the plasma first, before this is used to predict the appearance of the sample after machining. The motivation for this analysis was to examine whether there was a difference in the prediction accuracy between the direct and the indirect route, with the assumption that if there was no noticeable difference then the plasma did not provide any information about the appearance of the machined sample that was not present in the appearance of the sample before machining. In other words, this approach would judge whether the plasma only provided information about the appearance of the sample before machining, and that a neural network would have to create an internal model to predict the effect of laser machining. The additional neural network was trained on 950 sets of images and tested on 50 sets of images, with analysis of the results presented in Fig. 7, which shows average images for (b) plasma, (c) predicted after via the direct route, (d) predicted before, and (e) predicted after via the indirect route. The average predictions via the direct and indirect route show very strong similarity, hence providing evidence for this hypothesis. The average prediction errors for (f) the direct and (g)

indirect routes, as compared to the experimental data, also are very similar, hence providing further evidence. Calculations show that the indirect route was 3.3% more accurate in predicting the appearance of the sample after machining, which is smaller than the standard deviation of 4.5% and hence considered to be below the statistical significance level for this measurement. This leads back to the previous conclusion that the direct and indirect routes are equivalent in terms of prediction accuracy, and hence that, for the experimental conditions presented here, the plasma only contains information about the appearance of the sample before machining.

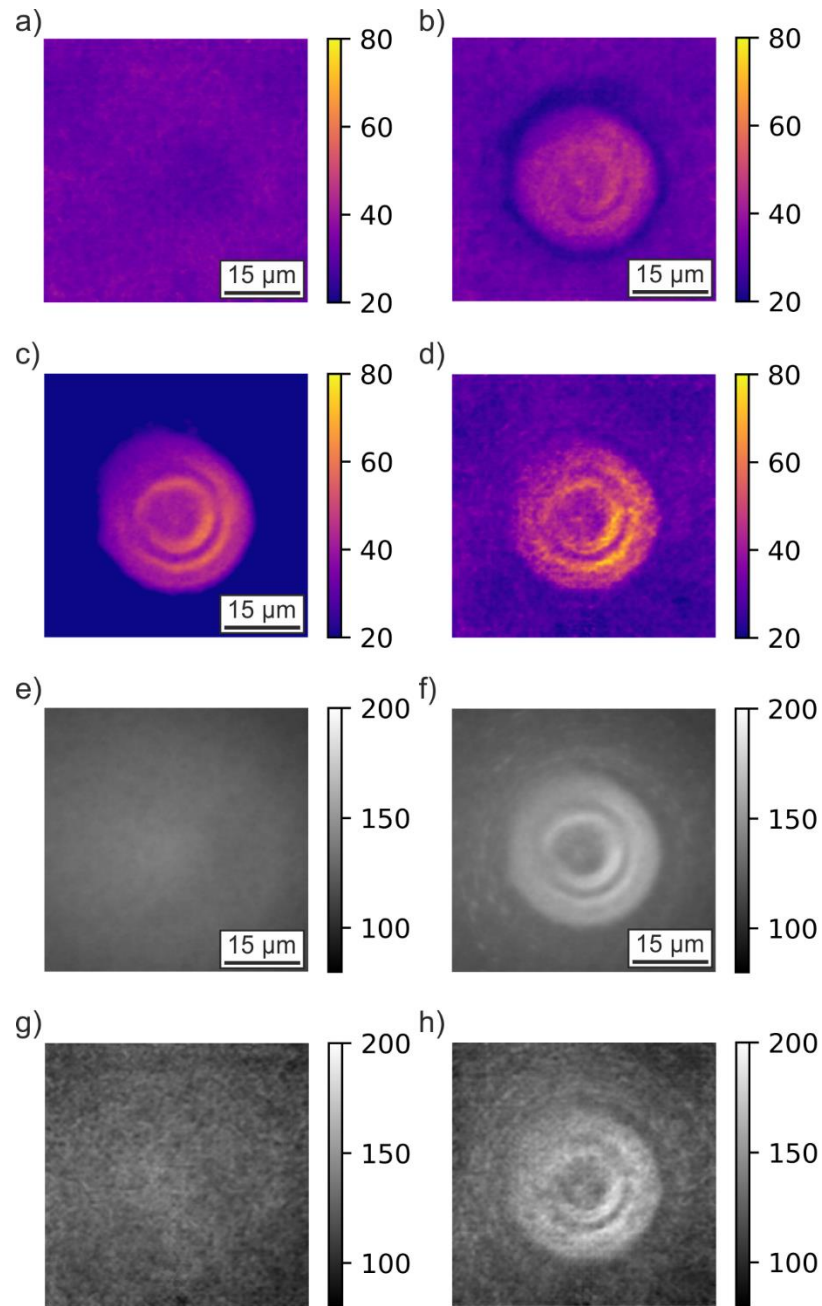


Fig. 6. Average absolute difference between (a) E1 and P1, (b) E2 and P2, (c) E1 and E2, and (d) P1 and P2 (where E1 = experimental before, E2 = experimental after, P1 = predicted before, P2 = predicted after). The figure therefore shows the prediction error for (a) before and (b) after the laser pulse, and (c) the real change and (d) the predicted change in the sample appearance due to the laser pulse. Also show for reference are the average images for (e) E1, (f) E2, (g) P1, and (h) P2. The sets of figures are shown using the same colour scale to assist in comparison.




Fig 7. Comparison of neural network capability in predicting the after image via a direct and indirect route. Showing (a) a flowchart describing the direct and indirect prediction route, the average images for (b) plasma, (c) direct after prediction, (d) before prediction and (e) indirect after prediction, and prediction errors for the (f) direct and (g) indirect routes.

Although plasma emissions can have lifetimes of anywhere between ns to several ms [43,44], the experimental plasma images in this work were recorded on a camera using the much longer integration time of 250 ms, and hence all plasma images corresponded to an integration over the whole time period of plasma emission. Similarly, whilst the melting and cooling of silicon, and subsequent morphology formation, can occur up to several nanoseconds after an incident pulse [45], the image of the sample after machining was recorded 400 ms after the laser pulse, and hence was after the sample had cooled. The training data provided to the neural network therefore contained no information about the temporal nature of the plasma emission, or the time scales for the melting and cooling of the silicon sample. Therefore,

although analysis of the neural network predictions provided evidence that the information in the plasma corresponded to the morphology of the sample before machining, it was not possible to use this neural network approach to identify the absolute time scales associated with plasma emissions and surface melting and cooling. However, a different detection approach, such as one linked to fast photodiodes, could be used in future for providing temporal information during the plasma emission to the neural network.

Whilst the results presented here correspond to single pulse ablation of silicon, this work could be extended to other materials and other laser conditions, such as different pulse lengths. A major challenge, however, would be the collection of experimental data that covers the desired set of material and laser conditions, and this would likely benefit from the support of experimental automation and robotics. In addition, it is likely that such a neural network would need to contain a larger number of parameters, and hence require additional computing hardware. However, assuming that it is possible to collect sufficient experimental data across different materials, it is plausible that a neural network could become able to learn a set of fundamental material properties that allow the network to make predictions for materials unseen during training. This capability could be enhanced further using physics informed neural networks [46].

Whilst the plasma generated by femtosecond laser machining may include a wide spectrum of emissions, including UV and IR wavelengths [47,48], the results in this manuscript were limited to those that could be recorded using the silicon-based CCD camera. Since a neural network generally becomes more accurate as the amount of information present during the training data is increased [49], it is plausible that the use of cameras that can detect wavelengths outside of this range may further improve the prediction accuracy of this approach.

#### 4. Conclusion

In conclusion, we have demonstrated the application of neural networks for indirectly imaging the surface of a silicon sample before and after machining with femtosecond laser pulses, directly from images of the generated plasma. This work could find use in a range of industrial applications where the plasma generated during laser materials processing prevents the observation of the work piece.

**Funding.** Engineering & Physical Sciences Research Council (EP/W028786/1, EP/T026197/1, EP/P027644/1).

**Disclosures.** The authors declare no conflicts of interest.

**Data availability.** Data underlying the results presented in this paper are available in Ref. [50].

#### References

1. M. Sparkes and W. M. Steen, "'Light' industry: an overview of the impact of lasers on manufacturing," *Advances in Laser Materials Processing* 1–22 (2018).
2. A. Wetzig, P. Herwig, J. Hauptmann, R. Baumann, P. Rauscher, M. Schlosser, T. Pinder, and C. Leyens, "Fast laser cutting of thin metal," *Procedia Manuf* 29, 369–374 (2019).
3. A. N. Fuchs, M. Schoeberl, J. Tremmer, and M. F. Zaeh, "Laser cutting of carbon fiber fabrics," *Phys Procedia* 41, 372–380 (2013).
4. S. Valette, P. Steyer, L. Richard, B. Forest, C. Donnet, and E. Audouard, "Influence of femtosecond laser marking on the corrosion resistance of stainless steels," *Appl Surf Sci* 252(13), 4696–4701 (2006).
5. J. Diaci, D. Bračun, A. Gorkič, and J. Možina, "Rapid and flexible laser marking and engraving of tilted and curved surfaces," *Opt Lasers Eng* 49(2), 195–199 (2011).
6. C. Velotti, A. Astarita, C. Leone, S. Genna, F. M. C. Minutolo, and A. Squillace, "Laser marking of titanium coating for aerospace applications," *Procedia CIRP* 41, 975–980 (2016).
7. T. Tamaki, W. Watanabe, and K. Itoh, "Laser micro-welding of transparent materials by a localized heat accumulation effect using a femtosecond fiber laser at 1558 nm.," *Opt Express* 14(22), 10460–10468 (2006).
8. J. Górká, W. Suder, M. Kciuk, and S. Stano, "Assessment of the Laser Beam Welding of Galvanized Car Body Steel with an Additional Organic Protective Layer," *Materials* 16(2), (2023).
9. J. A. Grant-Jacob, S. J. Beecher, J. J. Prentice, D. P. Shepherd, J. I. Mackenzie, and R. W. Eason, "Pulsed laser deposition of crystalline garnet waveguides at a growth rate of 20  $\mu\text{m}$  per hour," *Surf Coat Technol* (2018).



10. E. Morintale, C. Constantinescu, and M. Dinescu, "Thin films development by pulsed laser-assisted deposition," *Physics AUC* 20(1), 43–56 (2010).
11. A. Muniyallappa, H. Chandra, and D. Marla, "Numerical modeling to predict threshold fluence for material ejection in Laser-Induced forward transfer of metals," *Phys Scr* 98(9), (2023).
12. N. T. Goodfriend, S. Y. Heng, O. A. Nerushev, A. V Gromov, A. V Bulgakov, M. Okada, W. Xu, R. Kitaura, J. Warner, H. Shinohara, and E. E. B. Campbell, "Blister-based-laser-induced-forward-transfer: A non-contact, dry laser-based transfer method for nanomaterials," *Nanotechnology* 29(38), 385301 (2018).
13. P. Serra, M. Colina, J. M. Fernández-Pradas, L. Sevilla, and J. L. Morenza, "Preparation of functional DNA microarrays through laser-induced forward transfer," *Appl Phys Lett* 85(9), 1639–1641 (2004).
14. M. Feinaeugle, R. Pohl, T. Bor, T. Vaneker, and G. Römer, "Printing of complex free-standing microstructures via laser-induced forward transfer (LIFT) of pure metal thin films," *Addit Manuf* 24, 391–399 (2018).
15. Y. Kawahito, N. Matsumoto, M. Mizutani, and S. Katayama, "Characterisation of plasma induced during high power fibre laser welding of stainless steel," *Science and Technology of Welding and Joining* 13(8), 744–748 (2008).
16. J. Greses, P. A. Hilton, C. Y. Barlow, and W. M. Steen, "Plume attenuation under high power Nd: YAG laser welding," in *International Congress on Applications of Lasers & Electro-Optics* (Laser Institute of America, 2002), 2002(1), p. 47727.
17. A. HajiRassouliha, A. J. Taberner, M. P. Nash, and P. M. F. Nielsen, "Suitability of recent hardware accelerators (DSPs, FPGAs, and GPUs) for computer vision and image processing algorithms," *Signal Process Image Commun* 68, 101–119 (2018).
18. Y. Sun, N. B. Agostini, S. Dong, and D. Kaeli, "Summarizing CPU and GPU Design Trends with Product Data," (n.d.).
19. Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature* 521, 436 (2015).
20. C.-S. Ho, N. Jean, C. A. Hogan, L. Blackmon, S. S. Jeffrey, M. Holodniy, N. Banaei, A. A. E. Saleh, S. Ermon, and J. A. Dionne, "Rapid identification of pathogenic bacteria using Raman spectroscopy and deep learning.," *Nat Commun* 10(1), 4927 (2019).
21. J. A. Grant-Jacob, S. Jain, Y. Xie, B. S. Mackay, M. D. T. McDonnell, M. Praeger, M. Loxham, D. J. Richardson, R. W. Eason, and B. Mills, "Fibre-optic based particle sensing via deep learning," *Journal of Physics: Photonics* 1(4), 44004 (2019).
22. D. Ma, P. Jiang, L. Shu, and S. Geng, "Multi-sensing signals diagnosis and CNN-based detection of porosity defect during Al alloys laser welding," *J Manuf Syst* 62, 334–346 (2022).
23. Y. Xie, D. J. Heath, J. A. Grant-Jacob, B. S. Mackay, M. D. T. McDonnell, M. Praeger, R. W. Eason, and B. Mills, "Deep learning for the monitoring and process control of femtosecond laser machining," *JPhys Photonics* 1(3), (2019).
24. B. Mills, D. J. Heath, J. A. Grant-Jacob, Y. Xie, and R. W. Eason, "Image-based monitoring of femtosecond laser machining via a neural network," *Journal of Physics: Photonics* 1(1), 015008 (2018).
25. S.-K. Park, K.-H. Song, S. Y. Oh, J. S. Shin, and H. Park, "Improving Image Monitoring Performance for Underwater Laser Cutting Using a Deep Neural Network," *International Journal of Precision Engineering and Manufacturing* 24(4), 671–682 (2023).
26. N. Contuzzi and G. Casalino, "On modelling Nd:Yag nanosecond laser milling process by neural network and multi response prediction methods," *Optik (Stuttg)* 284, 170937 (2023).
27. L. Scime and J. Beuth, "A multi-scale convolutional neural network for autonomous anomaly detection and classification in a laser powder bed fusion additive manufacturing process," *Addit Manuf* 24, 273–286 (2018).
28. J. A. Grant-Jacob, B. Mills, and M. N. Zervas, "Acoustic and plasma sensing of laser ablation via deep learning," *Opt. Express* 31(17), 28413–28422 (2023).
29. W. Liu, Y. Rong, X. Fan, G. Xin, and Y. Huang, "Crack growth analysis of ultraviolet nanosecond laser scanning glass with acoustic emission," *Ultrasonics* 132, 106997 (2023).
30. W. Xing, X. Chu, T. Lyu, C.-G. Lee, Y. Zou, and Y. Rong, "Using convolutional neural networks to classify melt pools in a pulsed selective laser melting process," *J Manuf Process* 74, 486–499 (2022).
31. I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative Adversarial Networks," *Commun. ACM* 63(11), 139–144 (2020).
32. P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-Image Translation with Conditional Adversarial Networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (IEEE, 2017), pp. 5967–5976.
33. J. Tang, X. Geng, D. Li, Y. Shi, J. Tong, H. Xiao, and F. Peng, "Machine learning-based microstructure prediction during laser sintering of alumina," *Sci Rep* 11(1), 10724 (2021).
34. C. Liu, J. Shen, S. Hu, D. Wu, C. Zhang, and H. Yang, "Seam tracking system based on laser vision and CGAN for robotic multi-layer and multi-pass MAG welding," *Eng Appl Artif Intell* 116, 105377 (2022).
35. 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* 2022 12:1 12(1), 1–12 (2022).
36. S. Tani and Y. Kobayashi, "Ultrafast laser ablation simulator using deep neural networks," *Sci Rep* 12(1), 5837 (2022).
37. K. Shimahara, S. Tani, H. Sakurai, and Y. Kobayashi, "A deep learning-based predictive simulator for the optimization of ultrashort pulse laser drilling," *Communications Engineering* 2(1), 1 (2023).

38. M. D. T. McDonnell, J. A. Grant-Jacob, Y. Xie, M. Praeger, B. S. Mackay, R. W. Eason, and B. Mills, "Modelling laser machining of nickel with spatially shaped three pulse sequences using deep learning," *Opt. Express* 28(10), 14627–14637 (2020).
39. A. F. Courtier, M. Praeger, J. A. Grant-Jacob, C. Codemard, P. Harrison, M. Zervas, and B. Mills, "Predictive visualization of fiber laser cutting topography via deep learning with image inpainting," *J Laser Appl* 35(3), 032007 (2023).
40. J. A. Grant-Jacob, B. Mills, and M. N. Zervas, "Visualizing laser ablation using plasma imaging and deep learning," *Opt. Continuum* 2(7), 1678–1687 (2023).
41. J. A. Grant-Jacob, B. Mills, and M. N. Zervas, "Real-time control of laser materials processing using deep learning," *Manuf Lett* 38, 11–14 (2023).
42. B. Mills, D. J. Heath, J. A. Grant-Jacob, and R. W. Eason, "Predictive capabilities for laser machining via a neural network," *Opt Express* 26(13), 17245–17253 (2018).
43. D. J. Hwang, H. Jeon, C. P. Grigoropoulos, J. Yoo, and R. E. Russo, "Femtosecond laser ablation induced plasma characteristics from submicron craters in thin metal film," *Appl Phys Lett* 91(25), 251118 (2007).
44. M. Park, Y. Gu, X. Mao, C. P. Grigoropoulos, and V. Zorba, "Mechanisms of ultrafast GHz burst fs laser ablation," *Sci Adv* 9(12), eadf6397 (2023).
45. A. A. Ionin, S. I. Kudryashov, L. V. Seleznev, D. V. Sinitsyn, A. F. Bunkin, V. N. Lednev, and S. M. Pershin, "Thermal melting and ablation of silicon by femtosecond laser radiation," *Journal of Experimental and Theoretical Physics* 116(3), 347–362 (2013).
46. G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nature Reviews Physics* 2021 3:6 3(6), 422–440 (2021).
47. V. Narayanan and R. K. Thareja, "Emission spectroscopy of laser-ablated Si plasma related to nanoparticle formation," *Appl Surf Sci* 222(1), 382–393 (2004).
48. K. Zehra, S. Bashir, S. A. Hassan, A. Hayat, and M. Akram, "Spectroscopic and morphological investigation of laser ablated silicon at various laser fluences," *Optik (Stuttg)* 164, 186–200 (2018).
49. J. Cho, K. Lee, E. Shin, G. Choy, and S. Do, "How much data is needed to train a medical image deep learning system to achieve necessary high accuracy?," *arXiv preprint arXiv:1511.06348* (2015).
50. J. A. Grant-Jacob, B. Mills, and M. N. Zervas, "Dataset to support the publication 'Live imaging of laser machining via plasma deep learning'," University of Southampton (2023), <https://doi.org/10.5258/SOTON/D2764>.