



# Modelling laser machining of nickel with spatially shaped three pulse sequences using deep learning

M. D. T. McDONNELL,<sup>\*</sup>  J. A. GRANT-JACOB,  Y. XIE, M. PRAEGER, B. S. MACKAY,  R. W. EASON, AND B. MILLS 

*Optoelectronics Research Centre, University of Southampton, SO171BJ, UK*

*<sup>\*</sup>m.d.t.mcdonnell@soton.ac.uk*

**Abstract:** Femtosecond laser machining is a complex process, owing to the high peak intensities involved. Modelling approaches for the prediction of final sample quality based on photon-atom interactions are therefore challenging to extrapolate up to the microscale and beyond. The problem is compounded when multiple exposures are used to produce a final structure, where surface modifications from previous exposures must be taken into consideration. Neural network approaches allow for the automatic creation of a model that accounts for these challenging processes, without any physical knowledge of the processes being programmed by a specialist. We present such a network for the prediction of surface quality for multi-exposure femtosecond machining on a 5µm electroless nickel layer deposited on copper, where each pulse is uniquely spatially shaped using a spatial light modulator. This neural network modelling method accurately predicts the surface profile after three, sequential, overlapping exposures of dissimilar intensity patterns. It successfully reproduces such effects as the sub-diffraction limit machining feasible with multiple exposures, and the smoothing effect on edge-burr from previous exposures expected in multi-exposure laser machining.

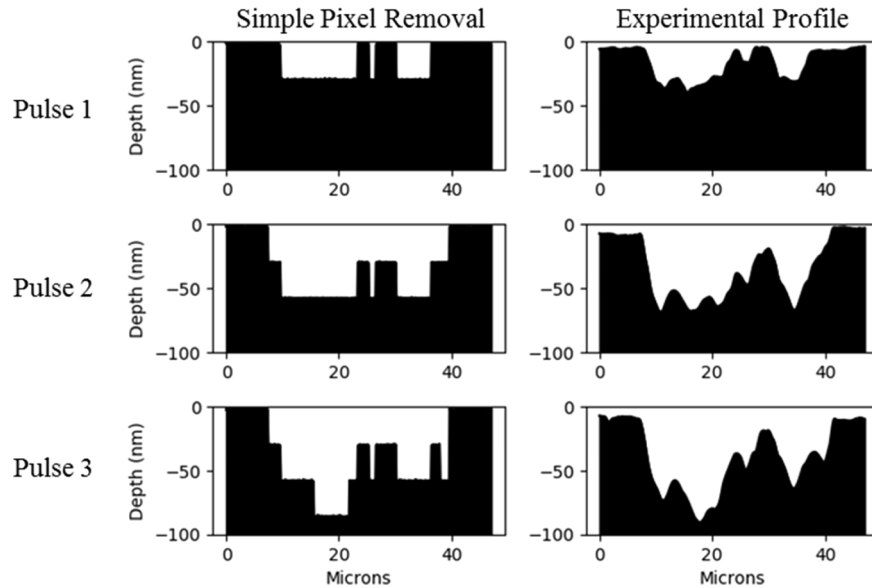
Published by The Optical Society under the terms of the [Creative Commons Attribution 4.0 License](#). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

## 1. Introduction

In general, when modeling femtosecond laser machining, a series of complex mechanisms must be accounted for. The first step is the multi-photon absorption and subsequent avalanche ionisation, which transfers energy from the incident light field to the work piece. While the effect of this is less apparent in metals than dielectrics, heat conduction, sample material phase-changes and the concomitant fluid and plasma dynamics, and eventual material redeposition as debris must all be taken into consideration [1–3]. Scaling this to three-dimensions, with all the complexity found in these interactions, is therefore a large task to undertake. This is compounded by interactions taking place on the ultrafast timescale and, in the case reported here, the two-dimensional intensity profile incident at the sample surface being subject to change. Photon-atom interactions have formed the basis of approaches by others for a strict theoretical understanding [4–6]. However, to expand the examined volume up to experimentally useful dimensions would require either an unrealistic level of computing power or simplifying assumptions to be made.

The challenge is further increased when the sample is exposed to more than one laser pulse. This is because the surface will only be uniform and smooth (compared to the depth of machined features) prior to the initial illumination. After the first pulse is used to machine the surface, several changes will have occurred. The simplest of these is that the surface will no longer be smooth and perpendicular to the beam. Taking the example of a typical Gaussian laser spot, the resultant feature after machining is expected to resemble a crater. The sloping bottom of this crater would cause non-normal incidence for subsequent laser pulses, changing the detail of the

light-matter interaction. Other factors could include changing material properties, incidence position compared to optimal focus, and the material temperature [7,8]. The importance of these effects, along with other femtosecond pulse features, is illustrated in Fig. 1 where profile predictions based on a naive method of removing material evenly where the laser is incident on the material are shown to be very different to those observed experimentally.




**Fig. 1.** Comparison of simple pixel removal, where material is removed evenly where the laser is incident on the surface, and experimentally measured profiles. The three pulses represent three sequential laser pulses with different spatial intensity profiles.

Previously, we have demonstrated a deep learning approach for the simulation of a sample surface after a single, shaped, ultrafast laser pulse exposure [9]. Experimental examples of spatially shaped intensity profiles and their resulting interferometrically depth-mapped structures in electroless nickel were used with a deep learning approach to train a neural network (NN). This NN was then used to predict the sample's appearance after being machined with binary intensity profiles that it had not seen. This approach required no coding of the physical principles involved in machining, a distinct advantage of the technique when attempting to model such a complex process.

Here, we show an extension of this principle, and demonstrate a neural network that can predict the machining result after multiple overlapping exposures, where the intensity profile is unique during each exposure. While this work focuses on femtosecond exposures, the method would of course apply equally to patterning via longer pulses. The extension from 1 pulse to 3 pulses was chosen to provide enough permutations ( $2^3 = 8$ ) in order to demonstrate the effectiveness of this technique for predicting the effect of the order of multiple laser pulses during laser machining. Based on the accuracy of the results presented here, it is anticipated that the result of laser machining with a much larger number of pulses could similarly be predicted by a neural network, provided that appropriate training data was provided. The inclusion of additional pulses may also require modifications to the neural network architecture, for example increasing the number of parameters. The additional physical complexity of exposing an already-machined surface strengthens our previous results. Here, we show that this approach offers a template for modelling processes whose underlying principles are too complicated to model accurately or that are entirely unknown.

## 2. Experimental setup

An appropriate training dataset is a fundamental requirement for training a NN. Here, a NN was used to convert spatial intensity profiles into depth profiles, and hence there were two experimental setup requirements. A laser machining setup with beam-shaping functionality allowed the input intensity patterns to be projected onto the work-piece, while an interferometric 3D surface profiler allowed accurate measurement of the depth profile of the laser machined substrates. Schematics of the setups and training set examples are shown in Fig. 2. Depth values in the figure are in nanometres, with the non-machined surface set to 0, hence any points with positive values are associated with burr and debris. Lateral values are measured in ‘pixels’, to highlight that these generated profiles are not real measurements, but rather generated by the NN.



**Fig. 2.** Experimental schematic used to collect all training and testing data. Showing a) schematic of the laser machining setup, b) illustration of the sample characterisation method.

A digital micromirror device (DMD) acted as an intensity spatial light modulator with a spatial resolution of  $840 \times 480$  pixels. This was used to spatially shape individual 150 fs, 800 nm, 200  $\mu$ J laser pulses from a Ti:sapphire amplifier. In this work, the laser was operating in single pulse mode (i.e. pulses on demand) at approximately one pulse per second. This was a requirement for the beam shaping device and associated automation. As a result, the training data did not include any information regarding thermal or other accumulation effects that would generally occur for high repetition rate lasers. In order to include such effects in the neural network, appropriate data from a high repetition rate laser would need to be used. A  $\pi$ Shaper 6\_6 and neutral density filter were used to produce an approximately uniform fluence at the DMD surface.

A single diffraction order from the DMD was imaged via a 50x objective onto the nickel work piece, at a calculated fluence of  $\sim 0.5 \text{ J cm}^{-2}$ . As the DMD has a fixed position compared to the laser, pointing instabilities are converted primarily into pulse to pulse intensity variations; however, since the DMD is smaller than the beam size and since the beam has a nominally top-hat profile at this location, there were no observable effects from laser pointing instability. A range of applications for this setup have been discussed elsewhere [10–13]. An electroless nickel mirror (5  $\mu$ m electroless nickel layer deposited on copper, LBP Optics Ltd.), which had an amorphous structure and was chosen to reduce grain boundary effects, was used as the work piece. A Zygo Zescope, white light interferometer [14] was used to measure the depth profiles (accurate to  $< 1$  nm) after all desired structures had been machined.

Here, the spatial resolution of the DMD was dependent on the magnification between the DMD and the work piece. Specifically, the DMD mirrors were each  $10 \mu\text{m} \times 10 \mu\text{m}$ , and each mirror was geometrically imaged down to a theoretical size of  $\sim 90$  nm. Each overall structure produced via up to 3 pulses was around 30 microns across. The patterns displayed on the DMD were generated

randomly by combining different geometric shapes including line segments, circles and arcs. The sizes of the features ranged both above and below the diffraction limit of the setup when the patterns were imaged. (The microscope objective used for focusing has an NA of 0.42 and the central wavelength of the laser is 800 nm giving an Abbé diffraction limit of approximately 952 nm). As all training data was collected at this specific resolution, the neural network was trained to operate at this particular resolution. In order to predict the result of laser machining for a different magnification, additional training data would be required.

### 3. Application of the neural network


NNs aim to emulate the processing structure of the human brain, using a set of interconnected virtual neurons which each receive, process and output information. It can be shown that a NN can be used to approximate *any* function [15], and is typically formed of an input layer, one or more hidden layers, and an output layer. The hierarchical series of neuronal connections that form the basis of the overall transform learned by the NN are iteratively adjusted via a training algorithm known as backpropagation [16,17]. This relies on pairs of input-output examples (“training data”) to demonstrate ‘correct’ transforms to the network. 404 input-output pairs were recorded and used for NN training.

The NN variant chosen in this work was a conditional generative adversarial network (cGAN) [18]. Previous work has shown that, when applied to laser machining, a cGAN can transform a spatial intensity profile of a laser pulse into a predicted depth profile that includes realistic randomly generated effects, such as debris [9,19]. Critically, as experimental data was used as the training data, all laser instabilities, such as variations in pulse energy (determined to be  $\sim 0.5\%$ ) and spatial intensity profile inhomogeneities, were also included in the learning process for the neural network. Here, the cGAN was designed as a combination of multiple neural networks, namely a ‘generator’, which is trained to learn the physical transform, and a ‘discriminator’, which is trained to judge whether the produced output is realistic [20]. The cGAN training process encourages the generator towards producing depth profile predictions that the discriminator judges to be realistic. At the same time, the discriminator is shown a mixture of experimentally measured depth profiles and predicted depth profiles that were produced by the generator and is trained to identify the differences. The generator therefore becomes more effective at making predictions, as the discriminator simultaneously becomes more effective at identifying the differences between experimental depth profiles and predicted depth profiles. The presence of a discriminator comprises the major difference between a cGAN and a traditional convolutional neural network, and enables the generation of realistic random features, rather than being trained to produce an ‘average’ result that may appear blurred, as a blurred result would be easily identified by the discriminator and hence discouraged during training.

The network used in this work is based on a recent development of the pix2pixHD network [21], called SPADE [22]. The network uses an upscaling generator, eschewing the down sampling path of a traditional U-net structure [23] to decrease network size. Instead, the network incorporates a down sampled version of the original image at each resolution. The network also employs feature loss where a large pre-trained network (VGG16) is used to extract features of both real and fake images to be compared. The network was trained for a total of 300 epochs, with an epoch being defined as the network being trained on each training pair exactly once. The network was trained on a PC with an Nvidia Quadro 6000 GPU with 24GB RAM.

For this work each input was a set of three images representing the patterns to be machined (Fig. 3). As explained in Section 2, these were uploaded to a digital micro-mirror device (DMD), which was used to spatially shape the laser intensity profiles. Experimentally measured depth profiles of the final surface structures, captured using a Zygo Zescope, formed the output data. Each structure was matched with a single input to form the data pairs. Importantly, the order of images in this three-dimensional stack represents the order in which the exposures are made

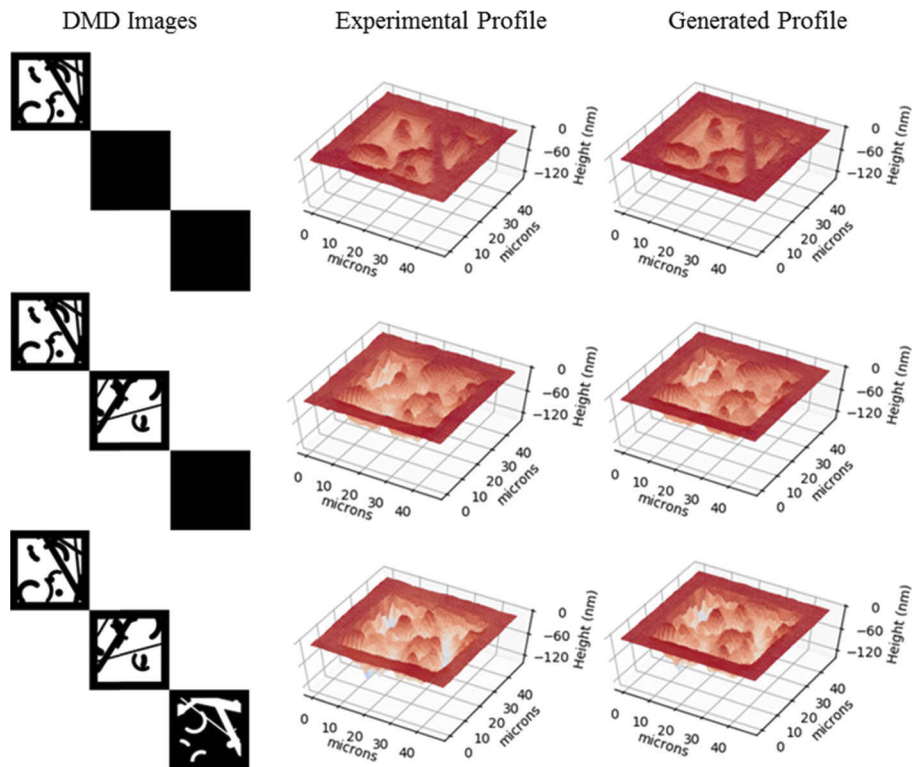
at the sample, with the first channel in the NN representing the first pulse fired. The NN then performs a transformation on the series of intensity profiles to predict the 3D surface profile of the laser machined surface.



**Fig. 3.** In this work, up to three DMD images are fed into a neural network, which then outputs a predicted depth profile.

#### 4. Validation of the neural network

A dataset held separate from the training data is typically used to test the accuracy of a NN model and is generally referred to as the validation dataset. Here, a set of 80 image stack and surface profile pairs that were not used during training, were used in order to observe the quality of the




**Fig. 4.** Demonstrating the ability of the NN to produce realistic depth maps. The left-hand column is the series of DMD images in order of how they were laser machined with white pixels denoting the 'on' state. Fully black patterns show where a pattern was not machined. The middle column shows the experimentally measured depth maps as measured by the Zscope. The right column contains depth maps generated by the neural network.



transform achieved by the cGAN on data not seen by the network during training. In addition to randomly generated images, input data with user-defined structures, such as grids, were also included in the validation set. Examples of both experimentally measured and generated depth profiles with their corresponding input maps are shown in Fig. 4.

As the experimental setup was subject to random noise (for example, laser pulse energy fluctuation, inhomogeneities in the sample, positional accuracy of sample surface relative to the image plane), small disparities are to be expected between the experimentally measured and generated outputs. These discrepancies are most evident in the un-machined surface as the cGAN has no knowledge of the initial surface profile and so cannot include effects such as scratches present on the surface. This is because while the cGAN has learnt a realistic distribution of random noise to apply, it cannot have knowledge of the exact parameters at the point of the validation example's creation. However, the fact that these random perturbations are represented at all in the cGAN model of laser ablation is an indicator of its potential. None of the patterns shown in Fig. 4 were included in the training data for the neural network.

Figure 5 shows a comparison between the experimental measurements and generated profiles using the same machining patterns as those seen in Fig. 1. This shows clearly that while the GAN was not able to account for all the random effects of laser machining, it was able to create a realistic model, producing results that differ little from experimental measurements. The percentage mean absolute errors of the generated profiles in this case were 2%, 5%, and 11%, for the one, two, and three pulse case respectively.



**Fig. 5.** A comparison between the experimentally measured profile and the GAN generated profile, for the same profiles as those in Fig. 1.

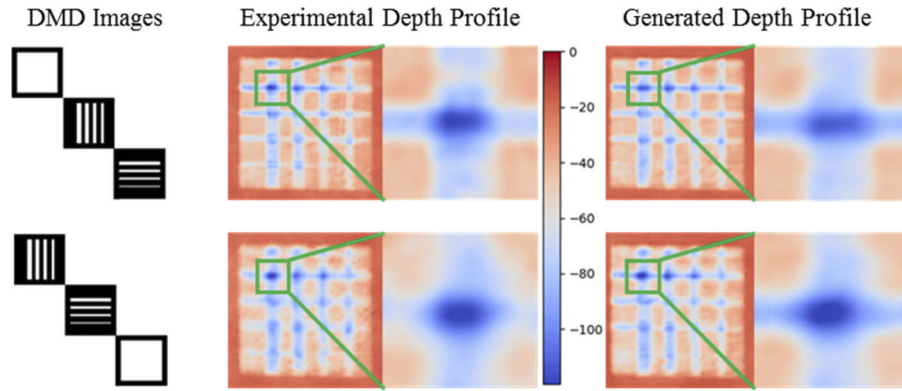
## 5. Analysis of the capability of the neural network

The cGAN enables prediction of the approximate surface profile after up to three exposures with unique intensity profiles, in a chosen order, even if the cGAN did not observe those exact profiles during training. When the aim is simply to estimate the final structure, the cGAN offers clear advantages over the experimental approach of positioning the sample, triggering laser exposures, transferring the sample to a surface profiler, and so on. While such a manual procedure may

take minutes to hours per iteration of the intensity masks used, depending on the particular experimental setup used, the transformation of intensity profiles to depth profile using the cGAN required only 29 ms.

### 5.1. Effect of the order of pulses on resolution

Of particular interest in this work is the effect of *order* on the sequence of exposures used. Suppose one wished to machine a recessed grating with a certain edge pattern, would it be better to machine the recess first and then the grating or the other way around. This situation is examined in Fig. 6 using validation data where two possible orderings of a solid square, vertical grating and horizontal grating have been machined (one on each row). The image stack sent to the DMD for three sequential pulses is shown, in addition to the experimentally measured and predicted depth profiles from the validation set. Beside each depth profile is a magnified view showing the edge sharpness between machined and un-machined areas. As can be seen, in the experimental data, where the gratings have been machined last (upper row), the edge sharpness is much higher than in the case where the solid square is machined last (lower row). This effect can also be seen in the depth profile generated by the NN.



**Fig. 6.** The results of laser machining a recessed grid using two different machining orders. In the first (top), a solid square is machined first, followed by the grid lines. In the second (bottom), the grid lines are machined before the solid square. The former is shown to produce sharper edges, in both the experimental case and for the NN prediction.

To numerically quantify the effect of the pattern ordering, the sharpness of each depth profile was calculated by finding the corresponding Laplacian energy [24], where higher energy values correspond to sharper images. The Laplacian energy was calculated using Eq. (1) where  $g(x,y)$  corresponds to either the experimental or predicted depth profile.

$$L = \iint \left[ \frac{\partial^2 g(x,y)}{\partial x^2} + \frac{\partial^2 g(x,y)}{\partial y^2} \right]^2 dx dy \quad (1)$$

Table 1 gives the mean Laplacian energy for the four cases shown in Fig. 6, (including the profiles with the ordering of the horizontal and vertical gratings reversed). The sharpness is clearly higher in both the experimental and the predicted depth profiles when the solid square is machined first. Furthermore, the magnitude of the order-dependent change in sharpness is similar for both the experimental, and generated profiles. This confirms that the NN was able to accurately predict the effect of the order of machining on the sharpness of the corresponding profile. It is interesting to note that the standard deviation of sharpness is lower for the NN generated profiles; this is perhaps because the NN is not subject to sample-to-sample experimental errors such as small changes in focusing.

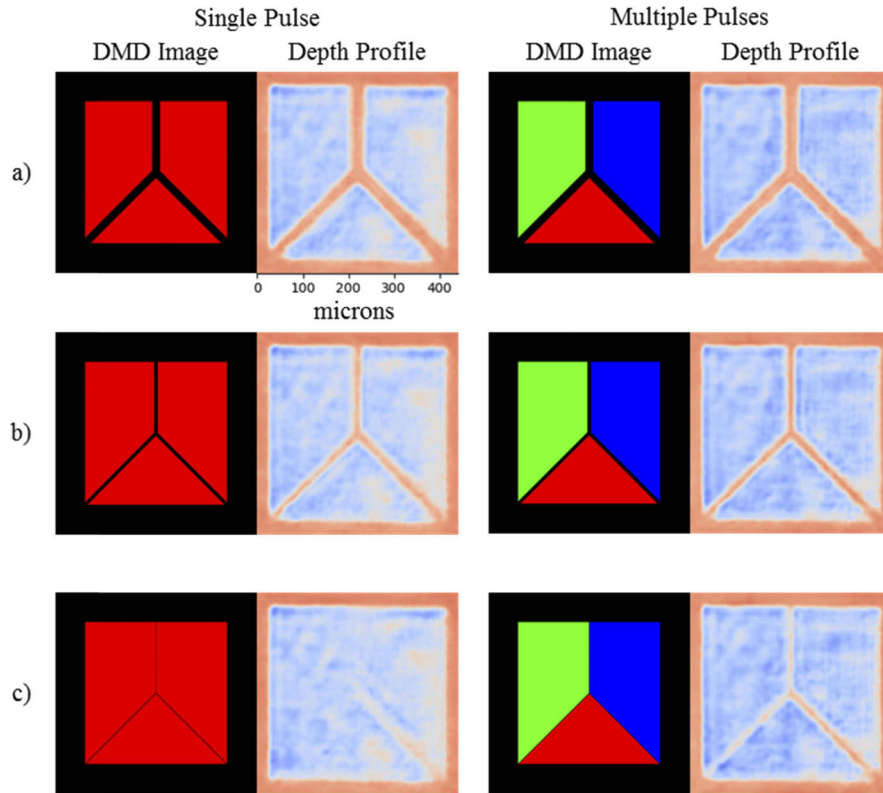
**Table 1. The effect of ordering on the sharpness of the experimental and generated depth profiles.**

	Laplacian energy	
	Experimental profile	Generated profile
<i>Solid square first</i>	$1.15 \times 10^{-3}$	$1.11 \times 10^{-3}$
<i>Solid square last</i>	$0.840 \times 10^{-3}$	$0.987 \times 10^{-3}$

### 5.2. Effect of multiple pulses on resolution

Of course, the analysis possible with the cGAN extends beyond prototype design. The actual physical response of the sample may also be investigated. Previously, it was discovered that by overlapping multiple exposures with small displacements, final features with dimensions that would be below the diffraction limit for a single exposure could be achieved [10]. A similar experiment can be performed by looking at the minimum size of a non-machined feature left between multiple laser exposed regions.

Figure 7 shows predictions from the cGAN for both single pulse and three pulse examples for three different separations of the exposed regions. In the figure, each pulse is represented by a



**Fig. 7.** Demonstrating the ability of the network to differentiate between the effects of a single pulse and multiple pulses, even with no overlap. On the left are DMD patterns that represent a single pulse, while on the right the DMD patterns represent three time-separated pulses. The RGB channels each show a different pulse, red representing pulse 1, green pulse 2, and blue pulse 3,. The spatial separations between the exposed regions are 1200 nm, 900 nm, and 100 nm for a), b) and c) respectively.




separate colour: red for pulse 1, green for pulse 2, and blue for pulse 3. Therefore, the pure red figures consist of a single pulse while the RGB represents three pulses.

In a) the separation between the machined areas was 1200 nm, above the diffraction limit of the system. The network correctly predicts that there should be very little difference between the profile machined by a single pulse as opposed to that machined by multiple pulses.

In b) the separation between the two machined blocks was 900 nm, just below the diffraction limit for a single pulse. The effect of this can be seen in the prediction for single pulse machining where the height of the non-machined area between the exposed regions is lower than the bordering non-machined area (as indicated by the fainter colouring). In the multi-pulse case, while the ridge is narrower, the height of the non-machined area is consistent with the bordering material, showing less of an impact from the diffraction limit.

In c) the separation between the two exposed regions was just 100 nm, well below the diffraction limit. In the single pulse case, the network predicts that the machining will be heavily impacted by the diffraction limit. In some areas this is shown by the ridge height being small compared to the machining depth while in others the non-machined ridge has completely disappeared. In contrast to this, the predicted outcome with three time-separated pulses is very similar to that predicted for separations above the single pulse diffraction limit.

The difference in behavior between single pulse and three pulse cases can be seen clearly in Fig. 8, where the height of the non-machined ridge is plotted. In previous work [19], the effect of the diffraction limit of the system (952 nm) was discussed. It was seen that when trying to leave a non-machined region, a feature was still seen down to a quarter of the diffraction limit, although with decreased maximum feature height past the diffraction limit. The cGAN used here predicts the same effect, with the peak of the ridge remaining level with the bordering material until the diffraction limit was approached. After this, the ridge height decreases until it almost disappears. With the multiple pulse case the height of the ridge stayed largely unaffected until approximately half the diffraction limit, as shown in [10]. This shows that the network has learned that the single pulse minimum feature size can be beaten by use of multiple pulses.



**Fig. 8.** The effect of separation on the height of remaining non-machined area for both single and multiple pulse cases. The vertical line represents the diffraction limit of the system.

## 6. Conclusions

We have expanded on our previous demonstration that a NN can learn physical characteristics of laser ablation by observing experimental data. In this work, the NN also learned to include the effect of pulse order on laser machining. Structures produced using multiple exposures with different spatial intensity patterns can now be modelled accurately, and used for rapid process optimisation, without any first-principles understanding of the machining process. Furthermore,

it was shown that the order of exposure of the different spatial intensity patterns influences the final machined structure and that this subtlety is successfully captured by the NN model. The network accurately predicted the blurring of features that occurs when large, coarse features are machined on top of fine features, whilst in the reverse case the edges remain sharp. Once trained, the network was also used to interrogate the capabilities of the machining setup. In this work, cGAN predictions were used to estimate the minimum feature size achievable with a single laser pulse, and how this could be bettered by using multiple time-separated pulses; these estimates accurately match previous experimental results but can be obtained more quickly and with lower cost.

For the approach presented here, the neural network is only able to predict the outcome of laser machining with the parameters that are intentionally varied in the training data. For example, the neural network can predict the effect of changing the spatial intensity profile, but is unable to predict the effect of changing the pulse energy or the laser wavelength. Whilst this proof of principle demonstration shows clearly the capability and precision that can be achieved via a neural network, future work should focus on the development of techniques that will enable the training of a neural network capable of predicting laser machining for all possible parameters.

## Funding

Engineering and Physical Sciences Research Council (EP/N03368X/1).

## Acknowledgments

We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Quadro 6000 GPU used for this research.

## Disclosures

The authors declare no conflicts of interest.

## References

1. E. G. Gamaly, A. V. Rode, B. Luther-Davies, and V. T. Tikhonchuk, "Ablation of solids by femtosecond lasers: Ablation mechanism and ablation thresholds for metals and dielectrics," *Phys. Plasmas* **9**(3), 949–957 (2002).
2. M. S. Amer, M. A. El-Ashry, L. R. Dosser, K. E. Hix, J. F. Maguire, and B. Irwin, "Femtosecond versus nanosecond laser machining: Comparison of induced stresses and structural changes in silicon wafers," *Appl. Surf. Sci.* **242**(1–2), 162–167 (2005).
3. O. Albert, S. Roger, Y. Glinec, J. C. Loulergue, J. Etchepare, C. Boulmer-Leborgne, J. Perrière, and E. Millon, "Time-resolved spectroscopy measurements of a titanium plasma induced by nanosecond and femtosecond lasers," *Appl. Phys. A: Mater. Sci. Process.* **76**(3), 319–323 (2003).
4. B. Rethfeld, D. S. Ivanov, M. E. Garcia, and S. I. Anisimov, "Modelling ultrafast laser ablation," *J. Phys. D: Appl. Phys.* **50**(19), 193001 (2017).
5. H. O. Jeschke, M. E. Garcia, M. Lenzner, J. Bonse, J. Krüger, and W. Kautek, "Laser ablation thresholds of silicon for different pulse durations: theory and experiment," *Appl. Surf. Sci.* **197–198**, 839–844 (2002).
6. J. K. Chen and J. E. Beraun, "Modelling of ultrashort laser ablation of gold films in vacuum," *J. Opt. A: Pure Appl. Opt.* **5**(3), 168–173 (2003).
7. I. H. Chowdhury and X. Xu, "Heat transfer in femtosecond laser processing of metal," *Numer. Heat Transfer, Part A* **44**(3), 219–232 (2003).
8. C. B. Schaffer, J. F. García, and E. Mazur, "Bulk heating of transparent materials using a high-repetition-rate femtosecond laser," *Appl. Phys. A: Mater. Sci. Process.* **76**(3), 351–354 (2003).
9. D. J. Heath, J. A. Grant-Jacob, Y. Xie, B. S. Mackay, J. A. G. Baker, R. W. Eason, and B. Mills, "Machine learning for 3D simulated visualization of laser machining," *Opt. Express* **26**(17), 21574–21584 (2018).
10. D. J. Heath, J. A. Grant-Jacob, M. Feinaeugle, B. Mills, and R. W. Eason, "Sub-diffraction limit laser ablation via multiple exposures using a digital micromirror device," *Appl. Opt.* **56**(22), 6398–6404 (2017).
11. D. J. Heath, B. Mills, M. Feinaeugle, and R. W. Eason, "Rapid bespoke laser ablation of variable period grating structures using a digital micromirror device for multi-colored surface images," *Appl. Opt.* **54**(16), 4984–4988 (2015).
12. D. J. Heath, J. A. Grant-Jacob, R. W. Eason, and B. Mills, "Single-pulse ablation of multi-depth structures via spatially filtered binary intensity masks," *Appl. Opt.* **57**(8), 1904–1909 (2018).

13. B. Mills, D. J. Heath, M. Feinaeugle, J. A. Grant-Jacob, and R. W. Eason, "Laser ablation via programmable image projection for submicron dimension machining in diamond," *J. Laser Appl.* **26**(4), 041501 (2014).
14. Zygo, "Optical Profiler Basics," <https://www.zygo.com/?/met/profilers/opticalprofilersabout.htm>.
15. K. Hornik, M. Stinchcombe, and H. White, "Multilayer feedforward networks are universal approximators," *Neural Networks* **2**(5), 359–366 (1989).
16. D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature* **323**(6088), 533–536 (1986).
17. Hecht-Nielsen, "Theory of the backpropagation neural network," in *International Joint Conference on Neural Networks* (IEEE, 1989), pp. 593–605 vol. 1.
18. 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.
19. 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).
20. I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative Adversarial Nets," in *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, eds. (Curran Associates, Inc., 2014), pp. 2672–2680.
21. T. C. Wang, M. Y. Liu, J. Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro, "High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs," *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, 8798–8807 (2018).
22. J.-Y. Park, Taesung and Liu, Ming-Yu and Wang, and Ting-Chun and Zhu, "Semantic Image Synthesis with Spatially-Adaptive Normalization," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2019), pp. 2337–2346.
23. O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)* **9351**, 234–241 (2015).
24. C.-Y. Wee and R. Paramesran, "Measure of image sharpness using eigenvalues," *Inf. Sci. (N. Y.)* **177**(12), 2533–2552 (2007).