

PAPER • OPEN ACCESS

## Deep learning-based reconstruction of irregular laser ablation crater geometry from single plasma projection

To cite this article: Yuchen Liu *et al* 2026 *Mach. Learn.: Eng.* **2** 015010

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

You may also like

- [Machine learning-based prediction of ultrafast laser ablation in nickel using early-stage transient density evolution](#)  
Rafiqul Haq, Tasnuva Islam, Md Mustakim Hayder et al.
- [Ultrafast laser filamentation classification and analysis via neural networks](#)  
James A Grant-Jacob and Ben Mills
- [Generative adversarial networks for creating realistic training data for machine learning-based segmentation of FIB tomography data](#)  
Trushal Sardhara, Christian J Cyron, Martin Ritter et al.

# MACHINE LEARNING

## Engineering

### PAPER

## Deep learning-based reconstruction of irregular laser ablation crater geometry from single plasma projection

Yuchen Liu\*, Fedor Chernikov, James A Grant-Jacob, Yunhui Xie, Michalis N Zervas and Ben Mills

Optoelectronics Research Centre, University of Southampton, United Kingdom

\* Author to whom any correspondence should be addressed.

E-mail: [yl22u22@soton.ac.uk](mailto:yl22u22@soton.ac.uk)

**Keywords:** ultrafast laser-matter interaction, laser ablation, plasma imaging, machine learning

Supplementary material for this article is available [online](#)

### Abstract

During ultrafast laser ablation, a plasma plume forms as the laser interacts with the material, while the resulting crater geometry reflects how energy is deposited. We present a machine learning framework that reconstructs the 2D ablation pattern using only a single side-view image of the laser-induced plasma. A conditional generative adversarial network is trained to map the plasma projection directly to the corresponding crater morphology, incorporating an edge-aware loss function to improve the reconstruction of irregular contours. Despite relying on a single projection, the model successfully recovers the dominant spatial structure of the ablation pattern. Analysis of the learned representations indicates that spatial features within the plasma image contain information beyond the primary viewing axis, enabling the reconstruction of laterally irregular geometries. While fine surface textures remain challenging to resolve from a single view, the method accurately predicts the dominant crater spatial morphology. The approach complements spectral techniques by providing morphological information from indirect imaging, supports real-time *in situ* monitoring, and offers a practical route toward predictive diagnostics in laser-material processing. The framework could be extended to multi view fusion potentially when higher reconstruction fidelity is required.

## 1. Introduction

Ultrafast laser technology is indispensable for precision material processing as it enables highly controlled surface modification. This capability arises from nonlinear light-matter interactions, where the energy of a high-intensity pulse is deposited within a confined volume and an ultrashort timescale. The deposited energy drives excitation, ionization, and subsequent nonthermal phase transitions that lead to rapid material removal and plasma formation [1, 2], a phenomenon known as laser ablation. The precision of this process has made it a cornerstone technique in applications such as micromachining [3, 4], surface texturing [5], and precision fabrication [6, 7], where controlled ablation is essential for high-quality structures [6]. Laser ablation also produces characteristic features such as craters, rims, and scattered particles that contain valuable information about both the laser parameters and the material response.

The morphology of ablation patterns carries important physical information. In fundamental studies, the geometry and periodicity of these patterns reveal the route of energy deposition and the subsequent thermal and mechanical responses [1]. Rethfeld *et al* [8] related crater geometry and edge formation to the generation of laser-induced periodic surface structures through nonthermal excitation and interference-driven localization. Liu *et al* [6] linked morphological features observed during plasma-assisted processing of transparent materials with transient absorption enhancement and back-surface coupling. While the ablation crater records the cumulative result of laser-matter coupling, the transient



### OPEN ACCESS

#### RECEIVED

15 January 2026

#### REVISED

20 March 2026

#### ACCEPTED FOR PUBLICATION

30 April 2026

#### PUBLISHED

14 May 2026

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

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



laser-induced plasma provides a dynamic optical view of the process. Plasma emission is highly sensitive to parameters such as focal position, target alignment, and beam profile, which contributes to the diversity of surface outcomes.

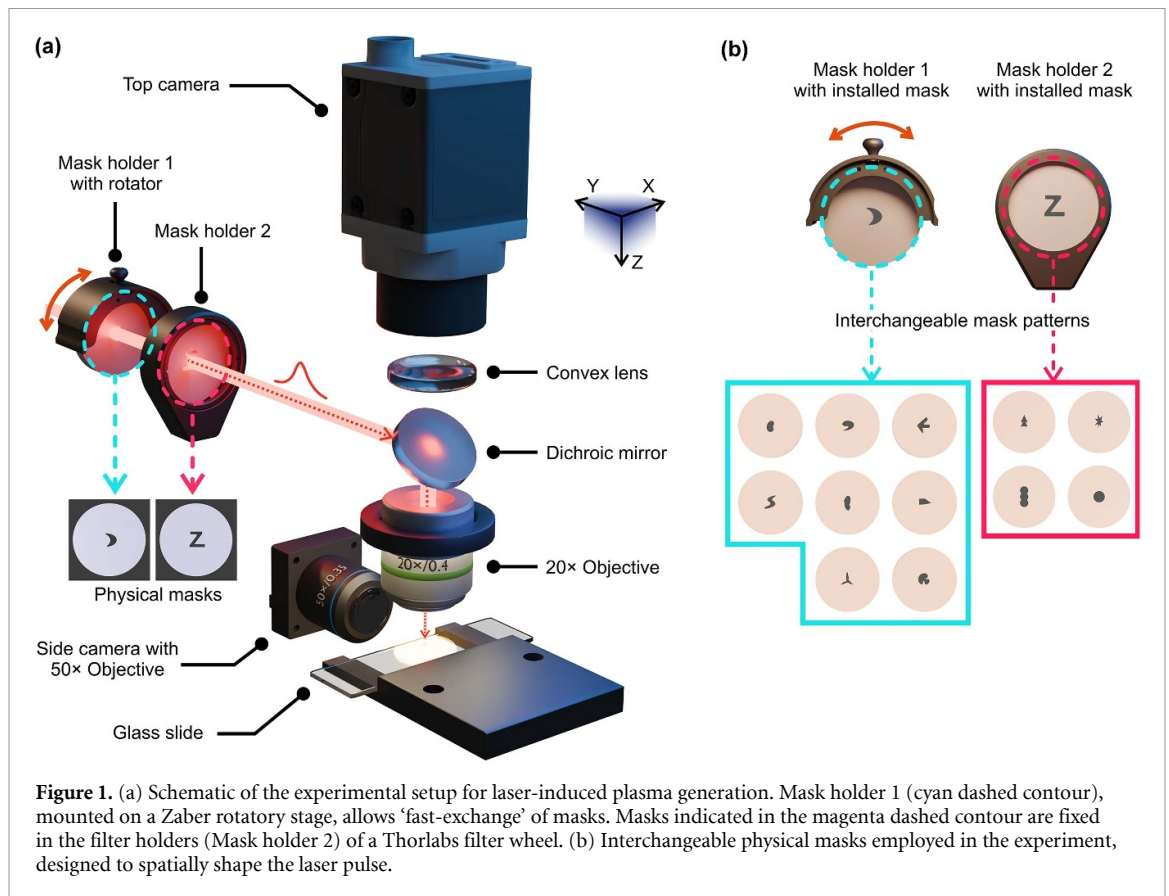
Reconstructing ablation crater geometry directly from plasma observations remains challenging. A single side-view projection provides only partial spatial information, and the relationship between plasma emission and crater morphology is complex and nonlinear. When the beam profile or material response is asymmetric, classical reconstruction methods such as Abel inversion [9] become invalid because they assume an axially symmetric plasma emission and ablation geometry. Acquiring multiple projections would provide additional information to help recover the missing spatial dimensions, but this would require precise alignment and more complex instrumentation. These limitations motivate the use of data-driven approaches that can infer spatial information directly from a single projection.

Recent advances in data-driven methods have shown strong potential for representing complex physical processes. Neural networks have been widely adopted in laser-based manufacturing to monitor process quality [10], detect defects [11], and predict fabrication outcomes [12, 13]. Conditional generative adversarial networks (cGAN) [14, 15] have shown a broad applicability. For example, You *et al* [16] used a cGAN as a powerful segmentation-enhancement framework for breast ultrasound imaging, improving segmentation accuracy and tumor delineation; while Liu *et al* [17] mapped pre- and post-cleaning data to predict the evolving arrangement of microbeads for adaptive control. Wang *et al* [18] used a cGAN to achieve super-resolved images with high perceptual quality. Grant-Jacob *et al* [19] demonstrated that a cGAN can reconstruct ablation craters from plasma images, with the reconstructed crater size correlating strongly with the laser pulse energy under a Gaussian beam profile.

Previous studies have investigated the prediction of laser ablation outcomes using physics-based simulations and data-driven approaches. Physics-based models have been developed to simulate laser-material interaction and estimate ablation characteristics such as crater depth or shape under specific processing conditions [20, 21]. Data-driven approaches have also been explored, where machine learning models are trained to predict crater geometry or surface characteristics [22, 23]. However, these approaches generally focus on estimating scalar quantities such as crater depth or diameter rather than reconstructing spatially irregular ablation morphology. Building on these developments, the present work explores whether deep learning can extend beyond symmetrical ablation conditions to handle spatially irregular beam profiles. To address this question, we apply a cGAN framework to reconstruct diverse and irregular ablation patterns from a single side-view image of the laser-induced plasma. Correlations between plasma emission and ablation pattern geometry are examined quantitatively [24] through principal component (PC) analysis of paired plasma-ablation pattern images (section 3), verifying that the plasma contains sufficient spatial information for prediction. These findings justify the use of a neural network (section 4) to predict ablation outcomes by learning the transformation from plasma appearance to crater morphology. Because a single projection is inherently ambiguous, the model incorporates geometric priors through an edge-aware loss and orientation-sensitive features to promote accurate crater geometry. Although fine textures are more difficult to reproduce, the approach reliably recovers the global shape and spatial distribution of ablation features, enabling the potential for real-time *in situ* monitoring and providing a pathway toward predictive diagnostics of laser-material interactions. In practical laser machining environments, the plasma plume and ejecta may partially obscure the machining surface, making direct surface inspection difficult, whereas plasma emission can be captured continuously during the interaction, providing an alternative route for real-time monitoring of ablation outcomes.

## 2. Experimental setup and data collection

Figure 1(a) depicts a schematic of the experimental setup employed in this study. A Light conversion Pharos SP femtosecond laser (190 fs, 6 W average power, 6 kHz repetition rate, central wavelength 1030 nm) delivered single pulses at each sample position (with a pulse energy of 400  $\mu\text{J}$ ) that were focused onto the target surface using a Nikon 20x objective (MUE21200). The target material was a glass microscope slide of dimensions 25 mm  $\times$  75 mm  $\times$  1 mm (Erie JME, 2950WX-003). Before introducing the physical masks used to modify the beam profile, the laser fluence at the target surface under the chosen focusing conditions was approximately 37.67 J cm<sup>-2</sup>. When masks were inserted, the fluence decreased due to partial beam blocking; for example, a representative mask combination resulted in a fluence of approximately 16.13 J cm<sup>-2</sup>. The interaction was observed simultaneously from two orthogonal directions; the surface of the sample was imaged along z-axis with a high-resolution CMOS camera (Basler, a2A5320-23ucPRO, 5320  $\times$  3032 pixels, RGB), while plasma emission was captured in the

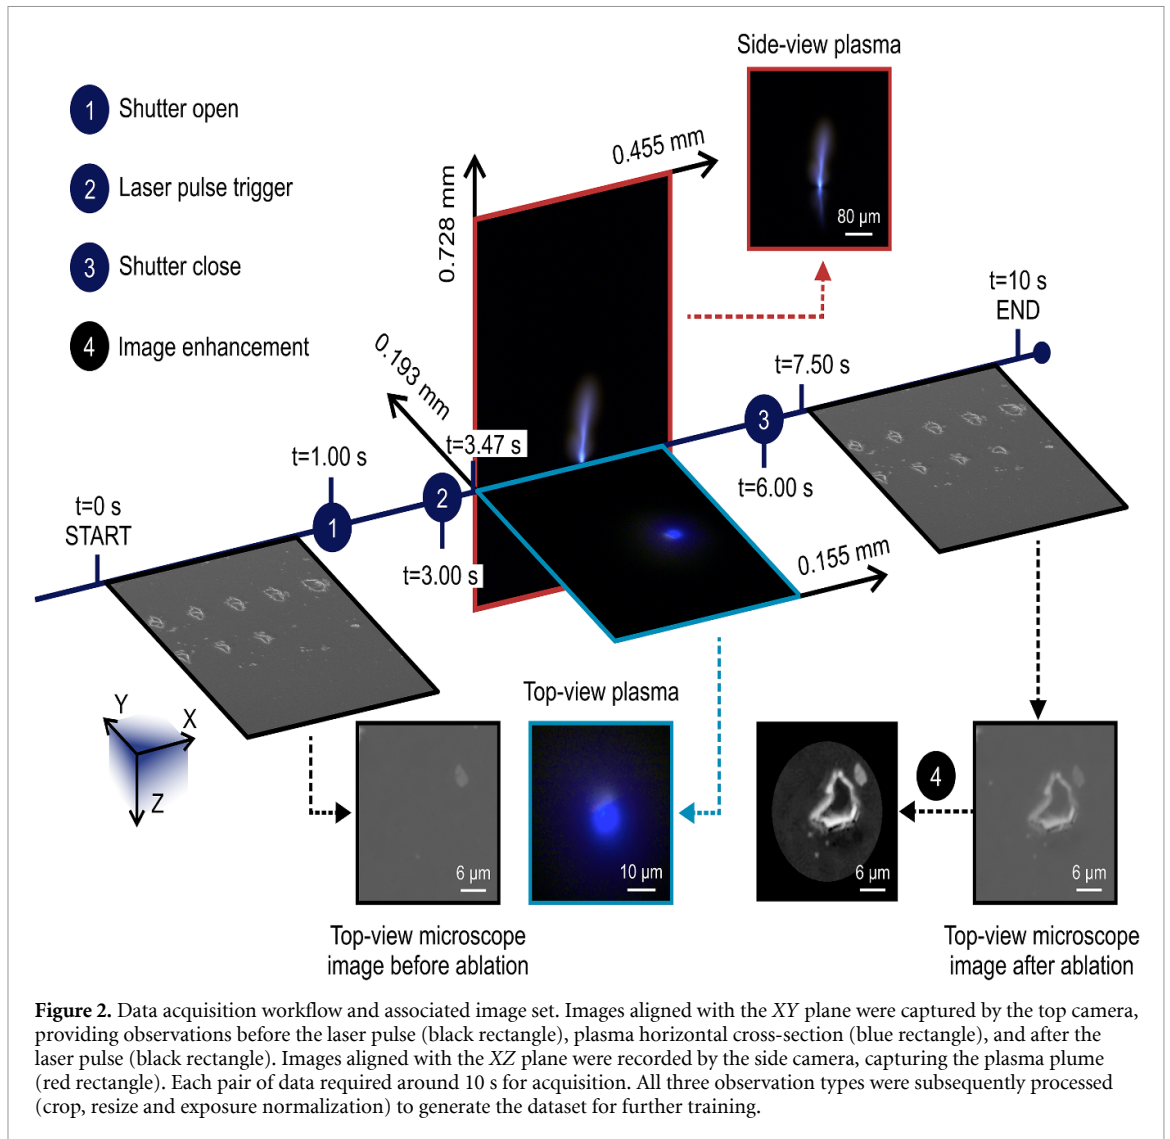


**Figure 1.** (a) Schematic of the experimental setup for laser-induced plasma generation. Mask holder 1 (cyan dashed contour), mounted on a Zaber rotatory stage, allows ‘fast-exchange’ of masks. Masks indicated in the magenta dashed contour are fixed in the filter holders (Mask holder 2) of a Thorlabs filter wheel. (b) Interchangeable physical masks employed in the experiment, designed to spatially shape the laser pulse.

side view (along the  $y$ -axis) through an Olympus SLMPlan N 50x objective coupled to a second CMOS camera (Basler, daA1920-160uc,  $1920 \times 1200$ , RGB). Both cameras were operated using fixed exposure settings during the experiments to ensure consistent image acquisition. The exposure times were set to 150 ms for the top-view camera and 90 ms for the side-view camera. Prior to data collection, a one-time white balance calibration was performed to maintain consistent color representation.

The sample was mounted on a motorized XYZ translation stage (Zaber, 3x X-LSM050A), offering a travel range of 5 cm per axis with  $0.048 \mu\text{m}$  command resolution, which enabled accurate positioning relative to the point where laser beam was focused. To control the incident beam profile, two mask holders were inserted upstream of the microscope system (figure 1(b)) to hold two physical masks. The masks were intended to distort the incoming beam and result in a highly irregular (azimuthally non-symmetric) ablation crater. The first holder (Mask holder 1) was attached to a rotatory translation stage (Zaber, X-RSW60A) that allowed full  $360^\circ$  rotation of the mounted mask, while the second (Mask holder 2) was a six-slot cage filter wheel (Thorlabs, CFW6) accommodating alternative masks. During the experiment, the mask in the first holder could be exchanged via a fast-change lens tube filter holder (Thorlabs, SM1QU), and the second holder could be rotated to position the selected mask into the beam path. All the physical masks were printed from 3D printer (Prusa i3 MK3S+) using PLA (Polylactic Acid, a biodegradable thermoplastic widely used for 3D printing) filament. This arrangement enabled flexible shaping of the incident laser profile for ablation studies.

To explore the correlation between the laser-induced plasma and the resulting ablation pattern, it was necessary to record both phenomena under synchronized experimental conditions. Each data point was therefore acquired through a fixed recording sequence, shown in figure 2, comprising several steps with overall duration of 10 s. Throughout the experiments, both the side-view and top-view cameras were mounted in fixed and calibrated positions to ensure consistent projection geometry across all recordings. The sequence begins after the translation stage positioned the glass sample to a new position, the top camera ( $1500 \times 1200$  pixels, field of view  $0.155 \text{ mm} \times 0.193 \text{ mm}$ , RGB) first observed the sample before laser irradiation. The shutter was then opened to prepare for exposure, after which the laser was triggered. Image acquisition by both the side camera ( $1980 \times 1200$  pixels, field of view  $0.728 \text{ mm} \times 0.455 \text{ mm}$ , RGB) and the top camera was performed during the laser-plasma event, after which the shutter was closed. Following this, the top camera recorded the sample after laser irradiation.

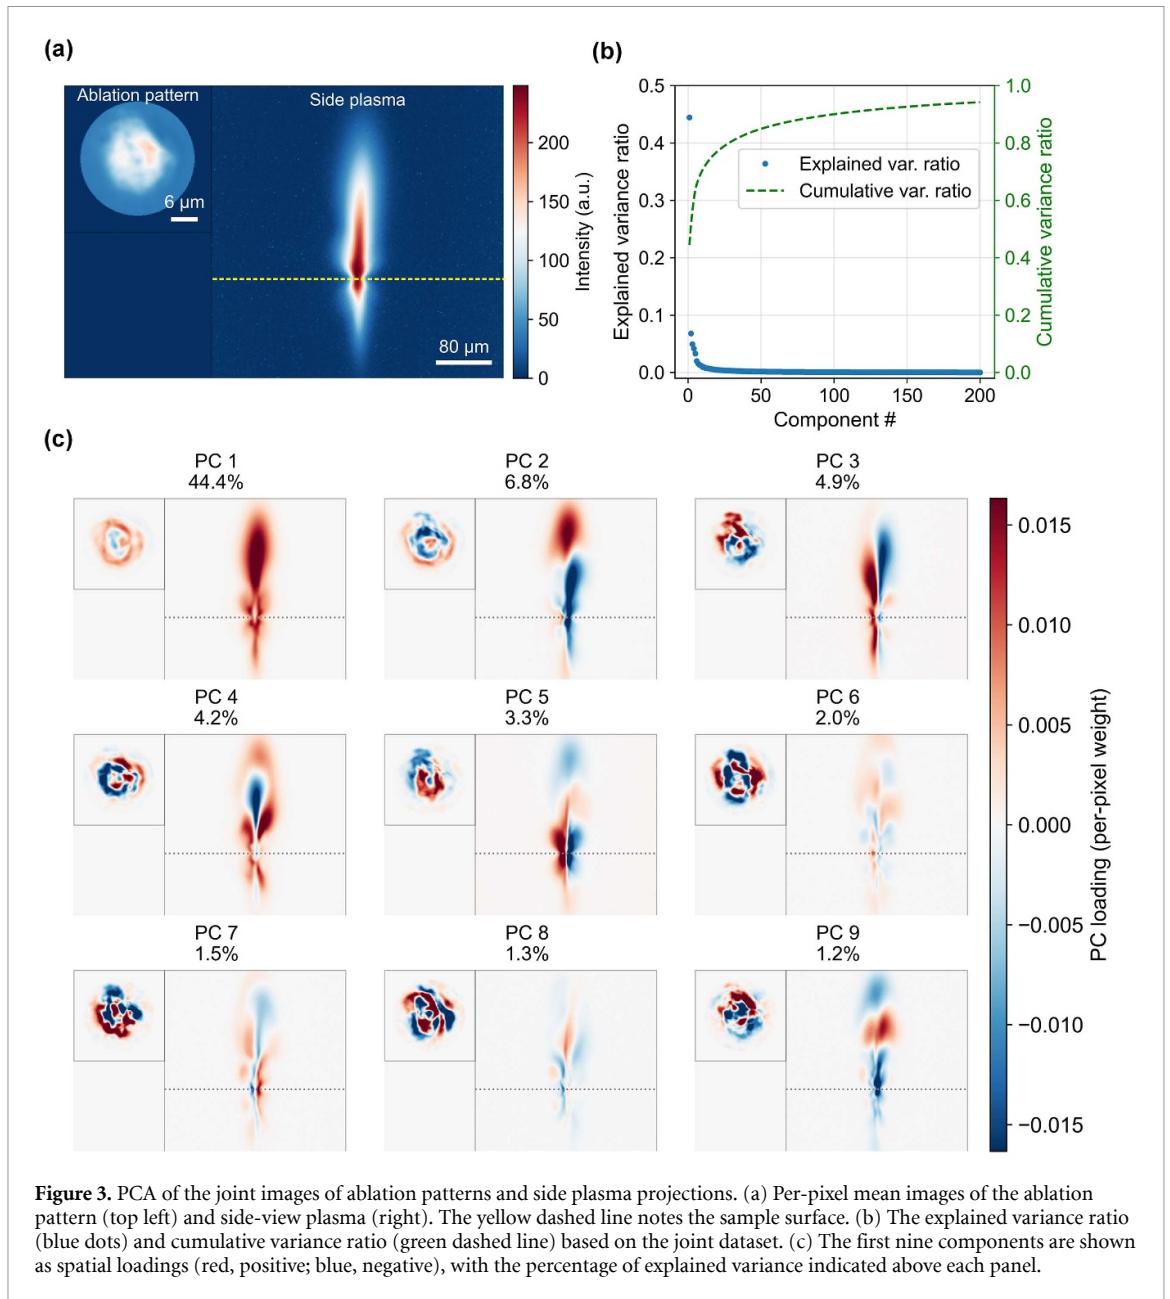


In total, nine different masks were available in holder 1 and five in holder 2. Since holder 1 was mounted on a rotator, we collected data by fixing a mask in each slot of holder 2 while rotating the mask in holder 1 every  $10^\circ$  across  $360^\circ$ . This procedure resulted in 45 unique mask combinations with 36 pairs of data for each. From these, we obtained 1620 experimental plasma-ablation data pairs. After excluding invalid recordings caused by mistimed captures, 1487 valid pairs remained for subsequent processing. All raw images were cropped and resized to  $256 \times 256$  pixels (RGB) prior to analysis and training.

Exposure normalization was applied to camera recordings to highlight the ablation pattern, where pixel intensities were rescaled using percentile-based normalization: regions with brighter pixel intensities were compressed by a soft-knee function and faint rims with darker pixel intensities were enhanced through adaptive gamma correction [25], followed by contrast-limited adaptive histogram equalization (see supplementary material SI). To reduce the influence of debris on the sample, the background was masked and replaced with a constant grayscale value of zero. This procedure produced consistent, exposure-stable images of the ablation features while minimizing saturation at the edge.

### 3. Principal components analysis

Principal component analysis (PCA) [26, 27] is a statistical method that reduces high-dimensional data into a smaller set of orthogonal patterns, called PCs, which capture the main modes of variation [28]. Instead of analyzing thousands of pixel values directly, PCA summarizes each image by numerical scores that describe how strongly these modes of variation are expressed. This approach highlights the dominant, shared changes across images, separates them from smaller view-specific variations, and provides quantitative measures such as explained variance. Grant-Jacob *et al* [24] demonstrated that the spatial

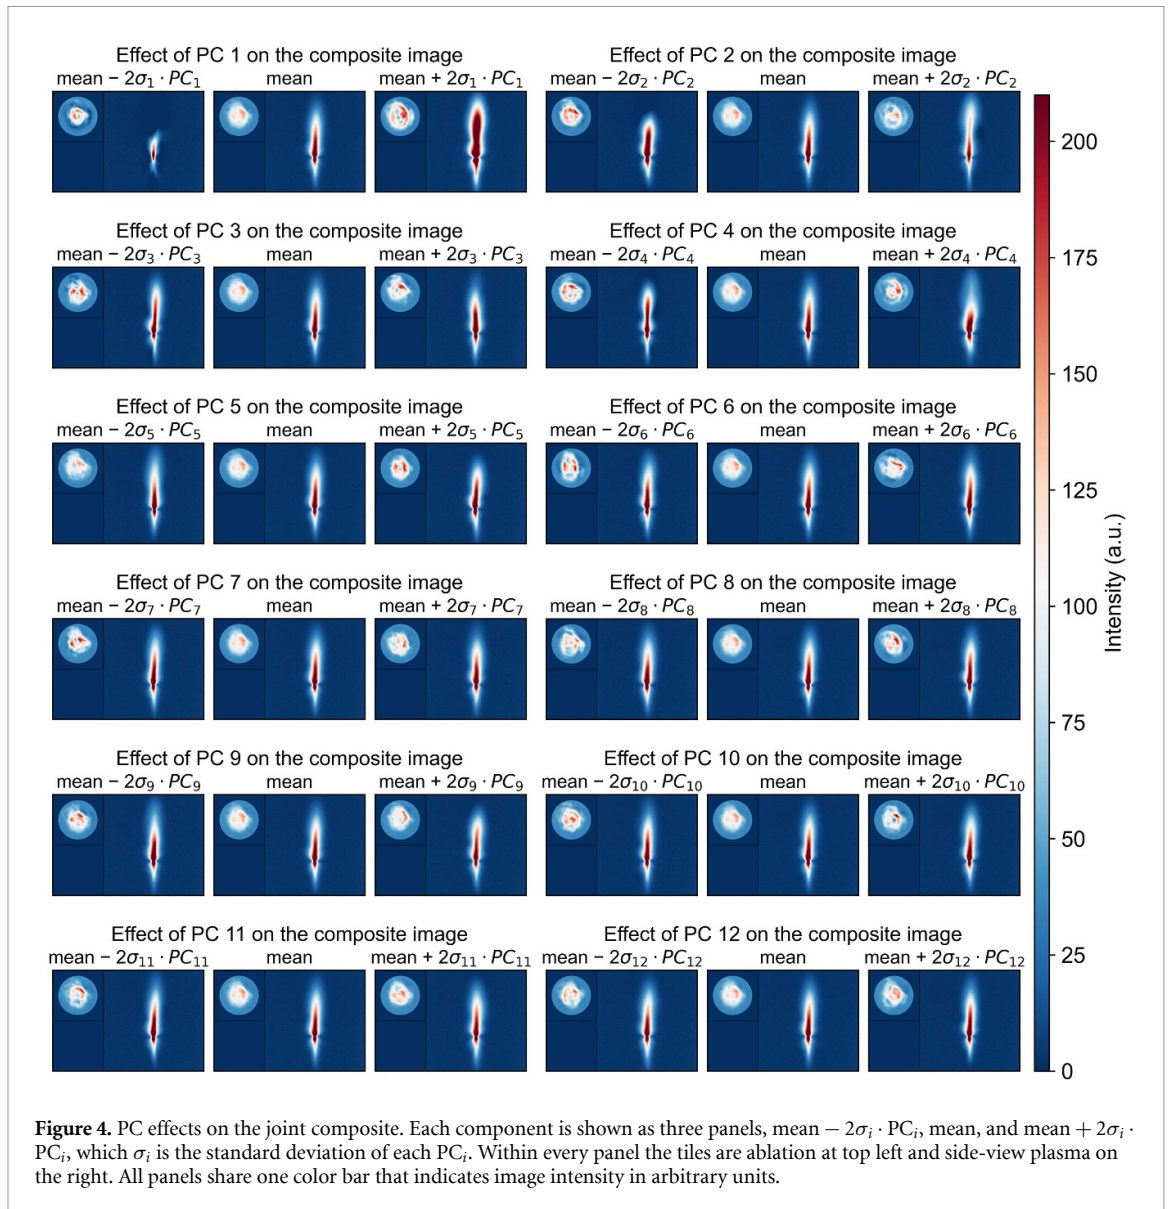


**Figure 3.** PCA of the joint images of ablation patterns and side plasma projections. (a) Per-pixel mean images of the ablation pattern (top left) and side-view plasma (right). The yellow dashed line notes the sample surface. (b) The explained variance ratio (blue dots) and cumulative variance ratio (green dashed line) based on the joint dataset. (c) The first nine components are shown as spatial loadings (red, positive; blue, negative), with the percentage of explained variance indicated above each panel.

PCs derived from PCA can be directly interpreted as indicating where variations occur in both the sample and plasma images, thus revealing correlations between spatial features in ablation patterns and plasma emission patterns.

In this study, PCA was not applied as a reconstruction tool but rather as a diagnostic framework to evaluate the extent to which the plasma projections (side and top views) and the ablation patterns share correlated modes of variation (figure 3). The joint dataset was decomposed into orthogonal components, with figures 3(a) and (b) showing the dataset mean images as a baseline together with the variance explained by each component. The leading components account for the majority of the variability, whereas later components contribute progressively less. In figure 3(c), the corresponding loading maps are presented as per-pixel weight distributions, where red and blue denote positive and negative contributions and white indicates values near zero. When dominant components are expressed simultaneously in both plasma and ablation images, this implies that information about the ablation pattern is, in theory, embedded in the plasma emission. The fact that such shared variation can be captured by a linear decomposition serves as supporting evidence that ablation outcomes are basically predictable from single plasma projections using a deep learning framework.

Using the PCA basis defined in figure 3, we visualized the effect of each component on the images (figure 4). The leading PCs (PC1—PC8) modulate plume intensity, width, length or focus which appear simultaneously in the plasma and ablation tiles, indicating a shared, low-dimensional structure that is



theoretically predictable. Higher PCs (after PC9) are predominantly ablation-dominated and capture fine-scale ring textures and local variations that are only weakly represented in the plasma images (see Supplementary material SII). Given that these modes contribute only weakly to the total variance and are spatially confined, they most likely reflect secondary or post-ablation effects such as surface inhomogeneity, or post-ablation material dynamics, that are less visible in plasma emission.

Effect along PC1 corresponds to a longer and brighter side plume and a stronger ablation ring, which may indicate changes in focus or laser energy deposition. PC2 primarily modifies the axial profile of the plume, where plume length and waist trade off, and the ablation image alternates in contrast between ring and core, which shows defocus effect from axial misplacement of the target. PC3 corresponds to a directional and pointing component in which the side plume tilts slightly, the ablation contour brightens asymmetrically. PC4 shows plume tightening with a cleaner central spine and a thinner ablation ring, suggesting energy concentration along the axis. PC5 reflects a core-shell trade-off, with the ablation ring thickening and the outer contour gaining contrast, consistent with slight misalignment. PC6 indicates a small standoff shift, where the plume waist and ablation ring radius vary modestly, resembling a minor focus height adjustment. These PCs indicate that beyond the dominant modes, finer variations arise from small changes in focus, alignment, or beam quality, which alter both plume morphology (to a lesser extent) and ablation pattern (to a larger extent), with the top spot remaining largely unaffected. Because the experiments employ single femtosecond laser pulses on a glass substrate, energy deposition occurs on a time scale much shorter than thermal diffusion [8]. As a result, the overall crater morphology is largely governed by the spatial energy distribution of the incident beam. Across PC1-6,

correlated variations are evident simultaneously in the side plasma and ablation pattern, suggesting that these features are shared and, in principle, can be inferred from plasma observations.

#### 4. Neural network background and results analysis

The neural network model we used is based on the pix2pix [29] architecture (structure can be found in [30]). The core U-Net generator and discriminator follow the standard structure, while the objective function was modified to better suit our work. Specifically, we introduced a warm-up strategy [31] and label smoothing in the model script to stabilize learning on small dataset, and augmented the generator objective with an additional edge-aware loss term (equation (1) to emphasize contour fidelity,


$$\mathcal{L}_{G_{\text{new}}} = \mathcal{L}_{\text{cGAN}} + \lambda_{\text{L1}} \cdot \mathcal{L}_{\text{L1}} + \mathcal{L}_{\text{edge}}. \quad (1)$$

Detailed objective functions can be found in the Supplementary material SIII. The dataset was split at an approximate ratio of 97:3, determined by the number of samples in the chosen mask combination (36 pairs of data as default). To minimize bias introduced by any specific mask combination, five additional neural networks were independently trained and tested for analysis. In each case, the validation set consisted of data from a mask combination excluded from the training pool, ensuring that the evaluation was performed on previously unseen geometries rather than memorized templates. While some plasma structures generated by different beam configurations may appear visually similar, this split reduces the possibility that the network simply associates a specific mask configuration with a fixed ablation pattern. Each model was trained for 150 epochs with a mini-batch size of four using the Adam optimizer [32]. The generator was configured with an encoding-decoding depth of eight (detailed NN structure see Supplementary material SV), corresponding to 61 layers including the input and output, while the discriminator had a 3-depth (10 layers with input and output) structure. Learning rates were set to  $1 \times 10^{-4}$  for the generator and  $5 \times 10^{-5}$  for the discriminator. Training was conducted on an Nvidia A100 (80 GB) through IRIDIS high performance computing facility served by University of Southampton.

After training and testing with five independently initialized neural network models, 24 test pairs were randomly selected from test set for demonstration (figure 5). The results show that the method reproduces the overall size and morphology of the ablation craters. For irregular or fragmented structures, however, predictions sometimes oversimplify the geometry toward smoother shapes or introduce artifacts based on the learned mapping. Because only side-view plasma images are provided as input, certain ablation features may be weakly encoded or ambiguous, leading to reconstruction uncertainty.

Although the training relies on the 2D image pairs (side-view plasma and top-view ablation patterns), the ablation images occasionally capture more than lateral morphology. In several cases with larger craters (e.g. data #6-9, #12-14, #17, #21), bright reflections are observed within the crater interior. These are likely associated with variations in crater depth or subsurface scattering, indicating that the dataset may contain indirect morphological cues related to depth. While pix2pix is fundamentally a 2D image translation model, the observed correlations between plasma structure and ablation crater geometry (size and shape), as well as the independence from mask rotation angle (see supplementary material SIV), indicate sensitivity to global morphological features. Depth-related variations, if present, are manifested as intensity variations within the ablation pattern (detail can be found in Supplementary material SV). Nevertheless, the depth information present is limited, indirect, and not systematically calibrated, so at the current stage the reconstructions should be regarded as qualitative rather than quantitative indications of depth. Nonetheless, their presence suggests that with improved imaging modalities or multi-view datasets, this framework could be extended toward three-dimensional ablation prediction rather than lateral morphology alone. Compared with multi-view or repeated measurements, the present approach demonstrates that predictive reconstruction is feasible from a single plasma projection, which is advantageous for *in situ* monitoring. Furthermore, once trained, the cGAN generates predictions rapidly ( $\sim 13$  ms per image), highlighting its potential for real-time and high-throughput applications.

The accuracy of the predicted ablation patterns was evaluated using metrics designed to quantify agreement at the pixel level as well as perceptual relevance. Table 1 reports prediction quality in terms of pixel-wise accuracy, perception-based image quality evaluator (PIQE) [33], structural similarity index (SSIM) [34], and learned perceptual image patch similarity (LPIPS) [35] for models trained with different objective functions. The cGAN + L1 configuration follows the original pix2pix setup; cGAN + L1 (low weight) retains the GAN term but reduces the L1 weight, encouraging sharper, more textured outputs; cGAN + L1 + Edge augments the baseline with our newly introduced edge-consistency term.



**Figure 5.** Examples of ablation pattern prediction from side-view plasma using the trained neural networks. Each set of data shows the input side-view plasma (left), the ground-truth ablation pattern (middle), and the prediction (right), with 24 test samples displayed in total.

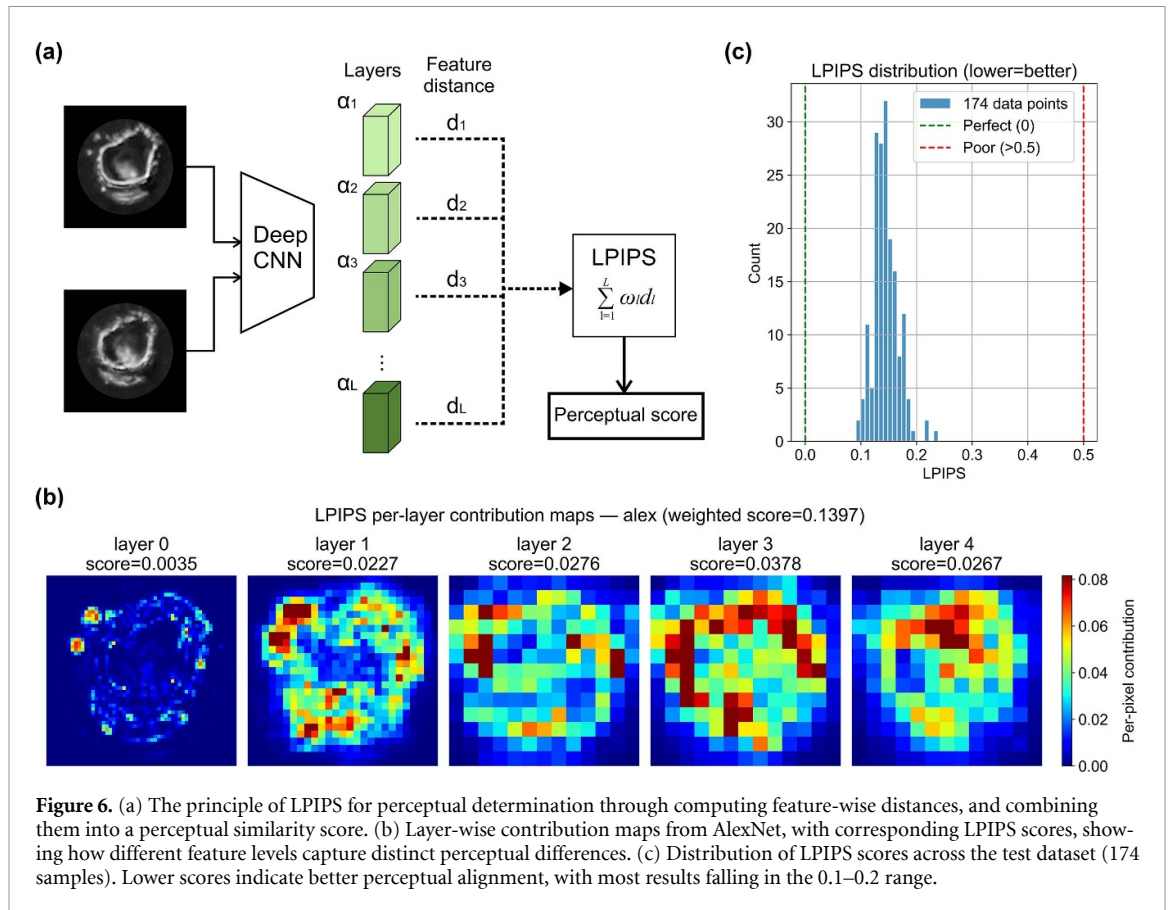
**Table 1.** Comparison of loss function variants on accuracy and perceptual quality metrics.

Objective function	Pixel-wise accuracy (0–1)	PIQE (0–100)	SSIM (0–1)	LPIPS (0–1)
cGAN + L1 (original)	0.69	<b>33.56 (Good)</b>	0.80	0.14
cGAN + L1 (low weight)	0.65	37.86 (Fair)	0.78	0.15
cGAN + L1 + Edge	<b>0.71</b>	34.46 (Good)	<b>0.82</b>	<b>0.14</b>

Pixel-wise accuracy was calculated using equation (2), defined as the average agreement between predicted and ground-truth images across all pixels, providing a direct measure of intensity similarity. Amongst the three analyzed models, the one incorporating the edge-aware loss term achieved the highest pixel accuracy. By focusing on absolute intensity differences (equation (1)), this metric yields a rigid assessment of reconstruction fidelity, while perceptual metrics offer a more representative measure of human-vision relevance,

$$\text{Relative accuracy} = 1 - \frac{|\text{Pre} - \text{GT}|}{\text{GT}} \quad (2)$$

PIQE is a no-reference image quality metric that estimates perceptual quality by analyzing spatial regions of significance based on the human visual system. It operates in a block-wise manner to detect artifacts, noise, and blur, and outputs a score on a 0–100 scale. According to the MATLAB implementation, scores below 20 correspond to excellent quality, 21–35 to good quality, 36–50 to fair quality, and



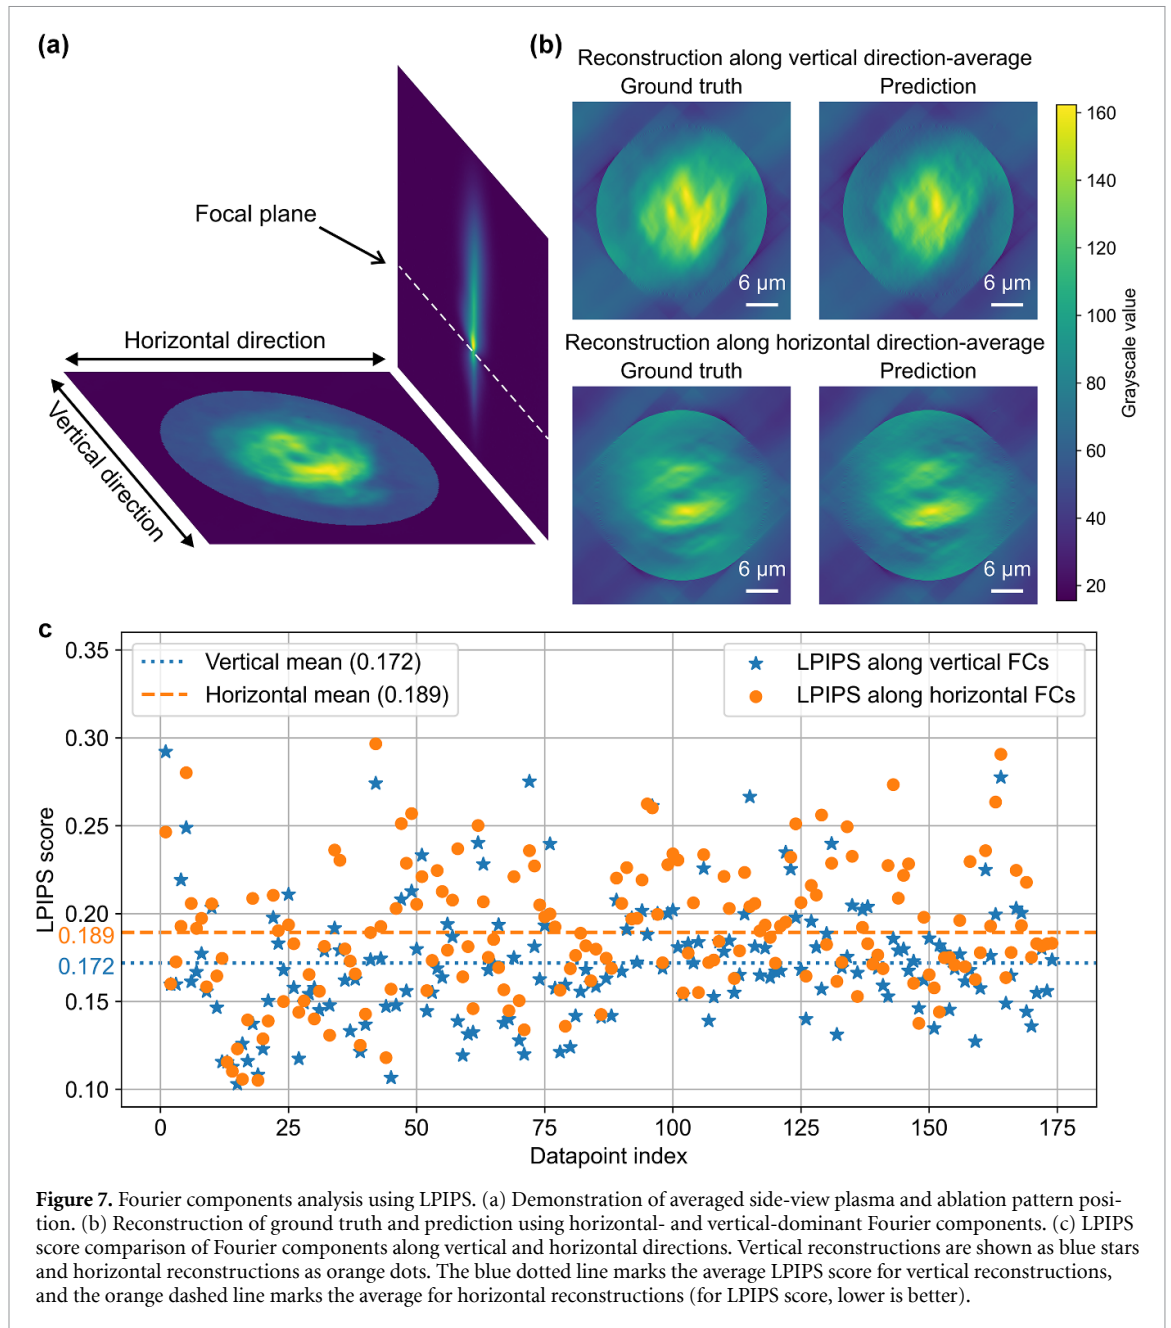
**Figure 6.** (a) The principle of LPIPS for perceptual determination through computing feature-wise distances, and combining them into a perceptual similarity score. (b) Layer-wise contribution maps from AlexNet, with corresponding LPIPS scores, showing how different feature levels capture distinct perceptual differences. (c) Distribution of LPIPS scores across the test dataset (174 samples). Lower scores indicate better perceptual alignment, with most results falling in the 0.1–0.2 range.

values above 51 indicate poor quality, with scores over 81 considered very poor. As shown in table 1, neural network models trained with the original L1 and the cGAN losses as the main objective function generally achieved scores corresponding to good perceptual quality.

SSIM is a full-reference image quality metric that evaluates perceptual similarity between two images by quantifying differences in luminance, contrast, and local structural information, based on principles of the human visual system. It is particularly effective for assessing the preservation of structural features, although it may assign relatively high scores to smoothed or blurred images, since such distortions reduce high-frequency differences while maintaining global structure. In this study, SSIM was calculated using the equation (3), in the formula, a reference (ground-truth) image  $GT$  and a comparison image  $pre$ ,  $\mu_{GT}$  and  $\mu_{pre}$  are the local mean intensities (luminance terms),  $\sigma_{GT}^2$  and  $\sigma_{pre}^2$  are the corresponding local variances (contrast terms)  $C_1$  and  $C_2$  are small positive stabilizing constants (refer to SSIM function in Scipy library, data range was set as 255). The scores obtained with the three objective functions were comparable, model with the original loss augmented by the edge term achieved the highest SSIM value, as summarized in table 1,

$$SSIM(GT, pre) = \frac{(2\mu_{GT}\mu_{pre} + C_1)(2\sigma_{GT,pre} + C_2)}{(\mu_{GT}^2 + \mu_{pre}^2 + C_1)(\sigma_{GT}^2 + \sigma_{pre}^2 + C_2)} \quad (3)$$

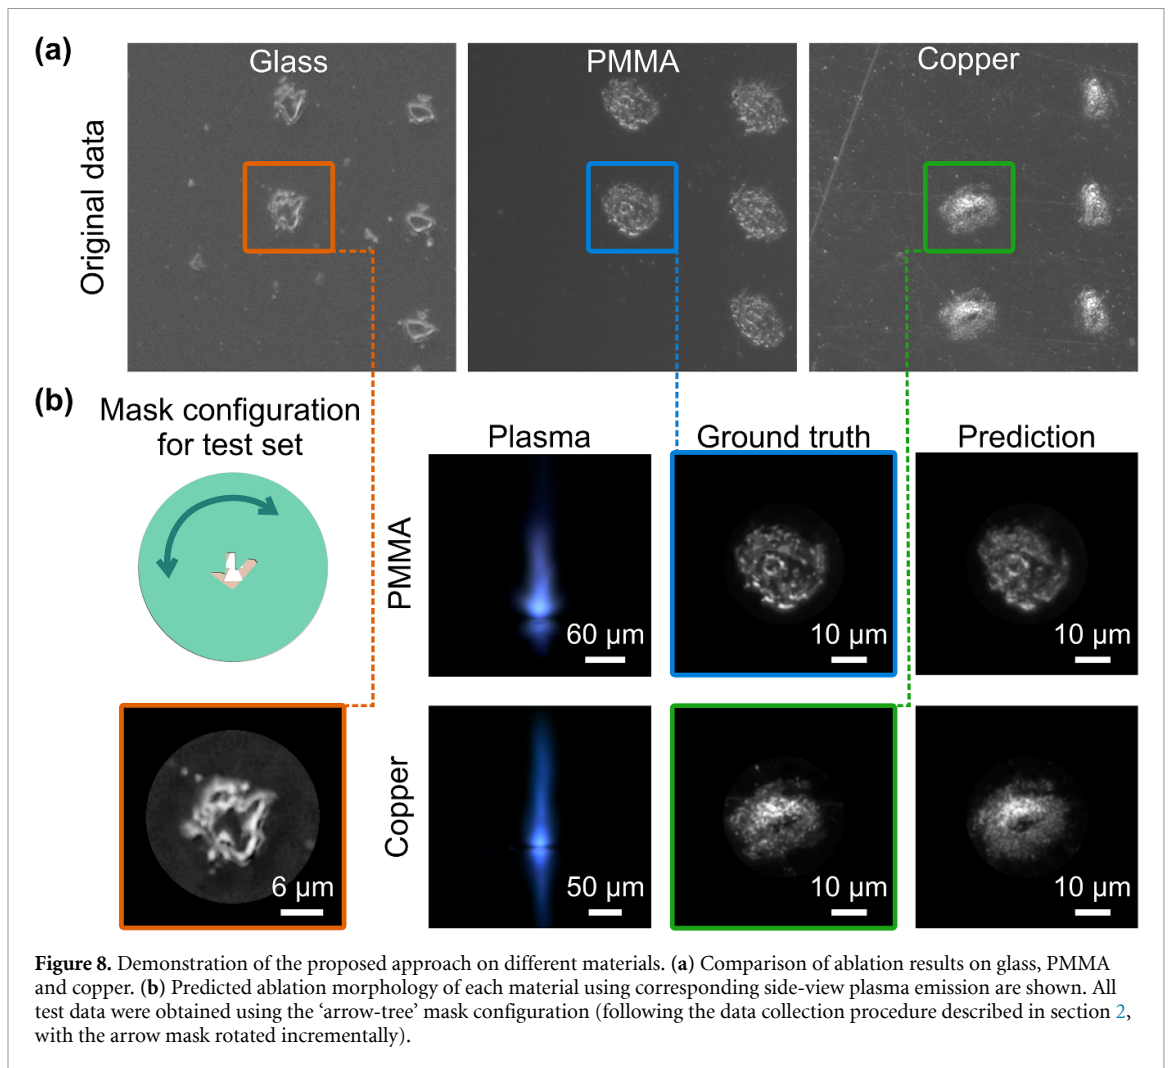
LPIPS is a perceptual image similarity metric designed to evaluate how closely two images resemble each other in a way that aligns with human visual perception (figure 6(a)). Two images are first passed through a CNN. AlexNet [36] was selected as it has been shown to provide strong performance. Feature activations are extracted from multiple layers of the network, and distances between corresponding features of the two images are computed. These distances ( $d_1-d_L$ ) are then linearly weighted (through  $\omega_i$ ) and aggregated into a single perceptual score. Unlike pixel-wise measurement, this metric uses hierarchical CNN features to evaluate perceptual differences in texture, structure, and semantics, offering a similarity assessment that aligns more closely with human visual perception. Figure 6(b) visualizes the feature maps across each layer from AlexNet. Early layers mainly capture local patterns and edges, whereas deeper layers emphasize broader structures and semantic features. The contributions from all layers are weighted and aggregated into a single scalar value, where lower LPIPS scores (closer to zero) indicate higher perceptual similarity. Figure 6(c) presents the distribution of LPIPS scores across the test set. The majority of samples distribute around 0.1–0.2, where the green dashed line indicates the perfect quality



(near 0), while the red dashed line indicates the poor quality (near 0.5), suggesting that most predictions maintain a reasonable degree of perceptual similarity to the ground truth.

## 5. Discussion

To further discuss the predicted ablation patterns along different orientations, we analyzed the images and separated Fourier components into vertical-dominant and horizontal-dominant groups for each test image and its corresponding ground truth. Using these subsets, we reconstructed orientation-specific images, one primarily capturing horizontal structures in the spatial domain and the other emphasizing vertical structures (figure 7(a)). The averaged intensity reconstructed from horizontal or vertical Fourier components is shown in figure 7(b). We then computed the LPIPS between the predicted and ground-truth images for each orientation (figure 7(c)). Both vertical and horizontal reconstructions yielded scores mostly within the range of 0.15–0.20. Figure 7(c) shows that the vertical reconstructions (blue stars) generally achieve lower LPIPS scores, indicating better perceptual similarity, compared with the horizontal reconstructions (orange dots). Considering the experimental setup, where the side-view



**Figure 8.** Demonstration of the proposed approach on different materials. (a) Comparison of ablation results on glass, PMMA and copper. (b) Predicted ablation morphology of each material using corresponding side-view plasma emission are shown. All test data were obtained using the ‘arrow-tree’ mask configuration (following the data collection procedure described in section 2, with the arrow mask rotated incrementally).

plasma was observed from the right side of the ablation pattern, this result is consistent with the expectation that vertical information (top-bottom structures) is better captured. The horizontal reconstructions also show reasonably good agreement with the ground truth even though only a single plasma projection was available. These findings suggest that plasma provides not only the information aligned with its primary viewing axis but also partial cues that contribute to reconstructing structures in the orthogonal orientation.

To explore the applicability of the framework beyond the glass substrate used in the main dataset, preliminary tests applied on PMMA and copper targets using the same approach are demonstrated in figure 8. All camera images were recorded after a single white-balance calibration to maintain consistent color representation across different materials. Compared with glass, PMMA and copper were investigated using higher pulse energy in order to obtain ablation features of comparable spatial dimension and plasma observations under the present experimental conditions, and the resulting crater morphologies tend to exhibit less clearly defined edges, making the crater boundary more difficult to identify directly from the surface images (figure 8(a)). For each material, 628 and 631 additional plasma-ablation image pairs were collected for PMMA and copper, respectively. Following the same evaluation strategy, one mask combination was excluded from the training pool and used exclusively for testing. Despite the differences in material response and surface appearance, the model is still able to recover the dominant crater morphology from the plasma projection (figure 8(b)). These results suggest that useful morphological information remains encoded in the plasma emission even when the surface features appear less clearly defined. A systematic investigation across different materials will be pursued in future work.

It should be noted that the present framework assumes a fixed and calibrated side-view imaging geometry, under which both training and testing data are collected. Variations in camera perspective may affect reconstruction accuracy, as the model learns spatial correlations under a specific projection geometry. Extending the framework to incorporate data augmentation or multi-view training strategies

would improve robustness to viewpoint changes and represent a promising direction for future development. From an implementation perspective, the computationally intensive training stage is performed during model development, whereas inference after training requires only a single forward pass. Consequently, the proposed approach remains compatible with real-time monitoring applications without requiring specialized high-performance hardware during deployment.

## 6. Conclusion

In this work, we demonstrate that spatially resolved ablation morphology can be inferred from a single side-view image of laser-induced plasma using a data-driven learning framework. By combining a pix2pix-based deep learning method with an additional edge-aware loss, the model reliably reconstructs dominant crater geometries and captures irregular contour features despite the indirect nature of the plasma observation. To validate the relationship between single side-view plasma and ablation images, PCA confirmed that the two modalities share correlated modes of variation, indicating that relevant information about ablation morphology is encoded in the plasma plume.

Fourier component analysis further shows that the learned mapping preserves directional and anisotropic characteristics of the ablation process, highlighting that spatially structured plasma features contribute to the reconstruction. These results establish a quantitative link between side-view plasma emission and ablation morphology, supporting the feasibility of plasma-based morphological sensing. Preliminary tests on metal and polymer targets further indicate that the approach can be extended beyond the glass substrate used in the main dataset. This suggests that the proposed framework may provide a foundation for applying plasma-based diagnostics to a broader range of material systems, as well as for developing closed-loop or reinforcement learning strategies in which plasma observations guide and optimize laser-material interactions. Future work will investigate multi-view integration and physics-informed constraints to enhance robustness and improve the reconstruction of fine-scale features.

## Acknowledgments

The authors acknowledge the use of the IRIDIS High Performance Computing Facility, and associated support services at the University of Southampton, in the completion of this work. This work was supported by the Engineering and Physical Sciences Research Council (EP/T026197/1, EP/W028786/1, EP/Z002567/1).

## Data availability statement

All data that support the findings of this study are included within the article (and any supplementary files) [37].

Supplementary data available at <https://doi.org/10.1088/3049-4761/ae6708/data1>.


## Conflict of interest

The authors declare no conflict of interest.


## ORCID iDs

Yuchen Liu  0009-0008-3636-1779

James A Grant-Jacob  0000-0002-4270-4247

Yunhui Xie  0000-0002-8841-7235

Michalis N Zervas  0000-0002-0651-4059

Ben Mills  0000-0002-1784-1012

## References

- [1] Phillips K C, Gandhi H H, Mazur E and Sundaram S K 2015 Ultrafast laser processing of materials: a review *Adv. Opt. Photon.* **7** 684–712
- [2] Chiang S and Albright C E 1989 Light-material interactions in laser material processing *7th Intl Symp on Gas Flow and Chemical Lasers* p 522
- [3] Gattass R R and Mazur E 2008 Femtosecond laser micromachining in transparent materials *Nat. Photon.* **2** 219–25
- [4] Faisal N, Zindani D, Kumar K and Bhowmik S 2019 Laser micromachining of engineering materials—a review *Micro and Nano Machining of Engineering Materials: Recent Developments* ed K Kumar, D Zindani, N Kumari and P Davim (Springer) pp 121–36
- [5] Kumar V, Verma R, Kango S and Sharma V S 2021 Recent progresses and applications in laser-based surface texturing systems *Mater. Today Commun.* **26** 101736
- [6] Liu H, Lin W and Hong M 2021 Hybrid laser precision engineering of transparent hard materials: challenges, solutions and applications *Light Sci. Appl.* **10** 162
- [7] Wang C, Wang Q, Qian Q and Di B 2020 The development of laser drilling: a review *IOP Conf. Ser.: Mater. Sci. Eng.* **782** 022067
- [8] Rethfeld B, Ivanov D S, Garcia M E and Anisimov S I 2017 Modelling ultrafast laser ablation *J. Phys. D: Appl. Phys.* **50** 193001
- [9] Dribinski V, Ossadtchi A, Mandelshtam V A and Reisler H 2002 Reconstruction of Abel-transformable images: the Gaussian basis-set expansion Abel transform method *Rev. Sci. Instrum.* **73** 2634–42
- [10] Hu Q, Wei X, Liang X, Zhou L, He W, Chang Y, Zhang Q, Li C and Wu X 2023 In-process vision monitoring methods for aircraft coating laser cleaning based on deep learning *Opt. Lasers Eng.* **160** 107291
- [11] Hazzan K E and Pacella M 2022 Surface defect detection and prediction in carbide cutting tools treated by lasers *Proc. CIRP* **108** 851–6
- [12] Westphal E and Seitz H 2021 A machine learning method for defect detection and visualization in selective laser sintering based on convolutional neural networks *Addit. Manuf.* **41** 101965
- [13] Wang B, Wang P, Song J, Lam Y C, Song H, Wang Y and Liu S 2022 A hybrid machine learning approach to determine the optimal processing window in femtosecond laser-induced periodic nanostructures *J. Mater. Process. Technol.* **308** 117716
- [14] Goodfellow I J et al 2014 Generative adversarial networks (arXiv:1406.2661)
- [15] Mirza M and Osindero S 2014 Conditional generative adversarial nets (arXiv:1411.1784)
- [16] You G, Qin Y, Zhao C, Zhao Y, Zhu K, Yang X and Li Y L 2023 A cGAN-based tumor segmentation method for breast ultrasound images *Phys. Med. Biol.* **68** 135002
- [17] Liu Y, Grant-Jacob J A, Xie Y, Chernikov F, Zervas M N and Mills B 2025 Selective laser cleaning of microbeads using deep learning *Sci. Rep.* **15** 15160
- [18] Wang J, Gao K, Zhang Z, Ni C, Hu Z and Chen D 2021 Multisensor remote sensing imagery super-resolution with conditional GAN *J. Remote Sens.* **2021**
- [19] Grant-Jacob J A, Mills B and Zervas M N 2023 Live imaging of laser machining via plasma deep learning *Opt. Express* **31** 42581–94
- [20] Jiang L and Tsai H L 2004 Prediction of crater shape in femtosecond laser ablation of dielectrics *J. Phys. D: Appl. Phys.* **37** 1492
- [21] Song S, Jiang L and Ji P 2024 Phenomenological modeling for femtosecond laser processing of fused silica *J. Manuf. Process.* **120** 365–77
- [22] Tani S and Kobayashi Y 2022 Ultrafast laser ablation simulator using deep neural networks *Sci. Rep.* **12** 5837
- [23] Yin Z, Liu Q, Sun P, Zhou Y and Ning Z 2023 Surface feature prediction for laser ablated 40Cr13 stainless steel based on extreme learning machine *Materials* **16** 505
- [24] Grant-Jacob J A, Zervas M N and Mills B 2024 Spatial analysis of femtosecond laser generated plasma using principal component analysis *Sci. Rep.* **14** 30301
- [25] Cao G, Huang L, Tian H, Huang X, Wang Y and Zhi R 2018 Contrast enhancement of brightness-distorted images by improved adaptive gamma correction *Comput. Electr. Eng.* **66** 569–82
- [26] Pearson K L I I 1901 On lines and planes of closest fit to systems of points in space *Lond. Edinb. Dubl. Phil. Mag. J. Sci.* **2** 559–72
- [27] Hotelling H 1933 Analysis of a complex of statistical variables into principal components *J. Educ. Psychol.* **24** 417–41
- [28] Greenacre M, Groenen P J F, Hastie T, D’Enza A I, Markos A and Tuzhilina E 2022 Principal component analysis *Nat. Rev. Methods Primers* **2** 100
- [29] Isola P, Zhu J-Y, Zhou T and Efros A A 2016 Image-to-image translation with conditional adversarial networks (arXiv:1611.07004)
- [30] Grant-Jacob J A, Everitt C, Eason R W, King L J and Mills B 2021 Exploring sequence transformation in magnetic resonance imaging via deep learning using data from a single asymptomatic patient *J. Phys. Commun.* **5** 095015
- [31] Kalra D S and Barkeshli M 2024 Why warmup the learning rate? underlying mechanisms and improvements (arXiv:2406.09405)
- [32] Kingma D P and Ba J 2014 Adam: a method for stochastic optimization (arXiv:1412.6980)
- [33] Venkatanath N, Praneeth D, Maruthi Chandrasekhar B, Channappayya S S and Medasani S S 2015 Blind image quality evaluation using perception based features *2015 Twenty First National Conf. on Communications (NCC)* pp 1–6
- [34] Zhou W, Bovik A C, Sheikh H R and Simoncelli E P 2004 Image quality assessment: from error visibility to structural similarity *IEEE Trans. Image Process.* **13** 600–12
- [35] Zhang R, Isola P, Efros A A, Shechtman E and Wang O 2018 The unreasonable effectiveness of deep features as a perceptual metric. *2018 IEEE/CVF Conf. on Computer Vision and Pattern Recognition* pp 586–95
- [36] Krizhevsky A, Sutskever I and Hinton G E 2012 ImageNet classification with deep convolutional neural networks *Proc. 26th Int. Conf. on Neural Information Processing Systems* vol 1 pp 1097–105
- [37] Yuchen L, Fedor C, James G, Yunhui X, Michalis Z and Ben M 2025 Dataset for “Deep learning-based reconstruction of irregular laser ablation crater geometry from single plasma projection” *University of Southampton Institutional Repository* (<https://doi.org/10.5258/SOTON/D3680>)