JPhys Photonics

Ultrafast laser filamentation classification and analysis via neural networks

To cite this article: James A Grant-Jacob and Ben Mills 2026 *J. Phys. Photonics* **8** 015058

IOP Publishing *J. Phys. Photonics* **8** (2026) 015058 https://doi.org/10.1088/2515-7647/ae42a6

Journal of Physics: Photonics

OPEN ACCESS

RECEIVED
12 December 2025

REVISED
21 January 2026

ACCEPTED FOR PUBLICATION
5 February 2026

PUBLISHED
16 February 2026

PAPER

Ultrafast laser filamentation classification and analysis via neural networks

James A Grant-Jacob* **and Ben Mills**

Optoelectronics Research Centre, University of Southampton, Southampton SO17 1BJ, United Kingdom
* Author to whom any correspondence should be addressed.

E-mail: J.A.Grant-Jacob@soton.ac.uk and bm602@orc.soton.ac.uk

Keywords: lasers, ultrafast, filamentation, plasma, neural networks, deep learning

Abstract

Laser filamentation is a nonlinear propagation regime in which high-power ultrashort pulses self-guide through air, generating plasma emission and broadband spectra. It underpins applications ranging from remote sensing to lightning control, yet detecting the precise onset of filamentation remains challenging. Here, we demonstrate a deep learning approach for identification and analysis of filament formation. A convolutional neural network (CNN) achieved a higher accuracy in predicting onset directly from plasma emission images than a linear regression model. A complementary non-negative matrix factorisation–CNN (NMF–CNN) regression revealed that spatial emission structure encodes sufficient information to reconstruct broadband spectra with strong fidelity (median $R^2 = 0.953$), linking image features to underlying physical processes. This methodology establishes a route toward real-time detection and analysis of ultrafast nonlinear light–matter interactions, with implications for laser diagnostics, high-power beam control, and photonic sensing.

1. Introduction

When an intense ultrashort pulse propagates through air, the optical Kerr effect drives self-focusing, which is counterbalanced by plasma generation and higher-order defocusing mechanisms [1, 2]. This dynamic interplay gives rise to a self-guided filament, which is a narrow, high-intensity plasma channel (or channels) that can propagate over distances far exceeding the Rayleigh range [3]. Laser filaments are sustained at peak laser intensities around $\sim 5 \times 10^{13}$ W cm^{-2} and generate broadband radiation including supercontinuum, fluorescence, and ionisation signals [4, 5], making them valuable for applications in remote sensing, atmospheric diagnostics, and nonlinear spectroscopy [6–8].

Despite extensive research, many aspects of filament physics remain difficult to predict or quantify. Control over filament position, brightness, and stability is hampered by the strong dependence on beam profile, pulse energy, and focusing geometry [9]. Traditional approaches rely on threshold-based diagnostics or detailed plasma modelling, including breakdown thresholds and emission processes [10–13], but these methods are poorly suited for real-time detection or adaptive control. At the same time, transitions between collapse, stable filamentation, and post-filament breakup encode information about higher-order nonlinearities and clamping mechanisms, motivating new approaches to identify filament onset that are both rapid and physically interpretable.

In parallel, deep learning has emerged as a powerful framework for analysing complex scientific data, owing to increase in capability of graphics processing units [14], which enables large-scale training of neural networks on high-dimensional datasets. Convolutional neural networks (CNNs) are a type of deep learning neural network can classify images with high accuracy [15–17], and have been used in photonics for calibration-free polarimetry using disordered optics [18], image classification [19], super-resolution imaging [20], light scattering classification [21], and the design of photonic structures [22] and photonic devices [23].

Deep learning neural networks such as CNNs, therefore, have proven highly effective for extracting complex patterns directly from images, offering the potential, as demonstrated here, for accurate detection of filament onset across a range of focusing conditions. To connect image-based features to underlying physical phenomena, we trained a CNN to predict spectral representations derived from non-negative matrix factorisation (NMF) of measured spectra. NMF, a dimensionality reduction technique conceptually related to principal component analysis (PCA) [24, 25] that enforces non-negativity and produces additive modes that are physically interpretable, unlike the signed components of standard PCA. In this NMF–CNN framework, the neural network learns to map spatial features in plasma emission images to dominant spectral modes, providing a data-driven decomposition of filamentation dynamics whilst preserving interpretability. In the long run, such interpretable representations could enable neural-network-driven feedback control, enabling potential manipulation and stabilisation of filamentation dynamics.

2. Experimental methods

2.1. Laser filament generation and optical setup

Femtosecond laser pulses (190 fs, 1030 nm, 6 W, 6 kHz repetition rate, Pharos, light conversion) were focused into ambient air in a class 4 laser laboratory using a variety of plano-convex fused silica lenses with focal lengths ranging from 75 mm, 100 mm, 125 mm, 150 mm, 200 mm, 250 mm to 300 mm. The input beam diameter was 8.89 mm ($1/e^2$) and the maximum pulse energy was ∼1 mJ. Beam power was controlled electronically from 20% in steps of 1%–100% of maximum output using the internal attenuator. The linear-focus peak intensities at 100% laser power, range from $\sim 3.5 \times 10^{13}$ W cm^{-2} (300 mm lens) to $\sim 5.6 \times 10^{14}$ W cm^{-2} (75 mm lens), which is consistent with filament onset in air ($\sim 5 \times 10^{13}$ W cm^{-2}).

The resulting plasma emission was imaged orthogonally to the beam axis using a Basler a2A3088-57uc (3088 × 20 654 pixels, RGB) complementary metal oxide semiconductor camera fitted with a Tamron 16 mm adjustable lens. This configuration allowed sufficient imaging of the full filament length for all lens configurations, and these images henceforth are called plasma images. Filament morphology, including plasma channel length, brightness distribution, and visible self-channelling, was thus recorded for each power step and focal length by taking an image of the generated plasma. All imaging data of the filament were acquired using an integration time of 80 ms. To assess shot-to-shot stability, we analysed 100-frame sequences (80 ms per frame, ∼480 pulses per frame) at 20% and 100% power. The coefficient of variation (CV) of integrated filament emission was ∼2.4% at 20% power and ∼1.7% at 100% power, whilst RMS radius (the square root of the mean squared radial distance weighted by intensity) CV was below 2.5% and centroid drift remained under 3 pixel. These results confirm highly stable filament formation under our experimental conditions.

Spectral measurements were acquired using an Andor W Tek Glacier X TE-cooled CCD (charged coupled device) array spectrometer. The thermoelectric cooling significantly reduced dark current and thermal noise, which, combined with background subtraction and dark-frame correction, ensured high signal-to-noise ratio (SNR) during long integration times (up to 4.5 s). A Tamron 1:2.1 35 mm adjustable lens was used to collect the plasma emitted light and collimate to a 25 mm focal length N-BK7 plano-convex lens mounted in an adapter (CVH100-COL, Thorlabs) so that light could be focused into a fibre connected spectrometer. The spectrometer and another imaging camera (Basler daA1920-160uc, 1920 × 1200 pixels, RGB) were used interchangeably via a magnetic mounting system to ensure alignment of the spectrometer with the filament. Spectral data was recorded in the range 400–1100 nm with integration times ranging from 5 s for 75 mm focal length lens to 60 s for 300 mm focal length lens, to avoid saturation. A schematic of the optical setup, representative plasma images, and corresponding spectra for selected configurations are shown in figure 1. For each condition, one image and corresponding spectra were acquired.

The dataset comprised of 567 image–spectrum pairs (1% power increments giving 81 per lens) collected across multiple focal lengths and power settings. Data pairs were split into 324 samples for training, 81 samples for validation (20% of the training set), and 162 samples for testing (200 mm and 250 mm focal length lens data). Testing was performed on lenses not included in training to assess generalisation. Each image was cropped to a fixed region of interest and resized to 128 × 128 pixels to ensure compatibility with the CNN input layer, with data augmentation being applied during training using random rotations (±2°), scaling (±2%), and translations (±15 pixels) to improve generalisation and via essentially expanding the dataset.

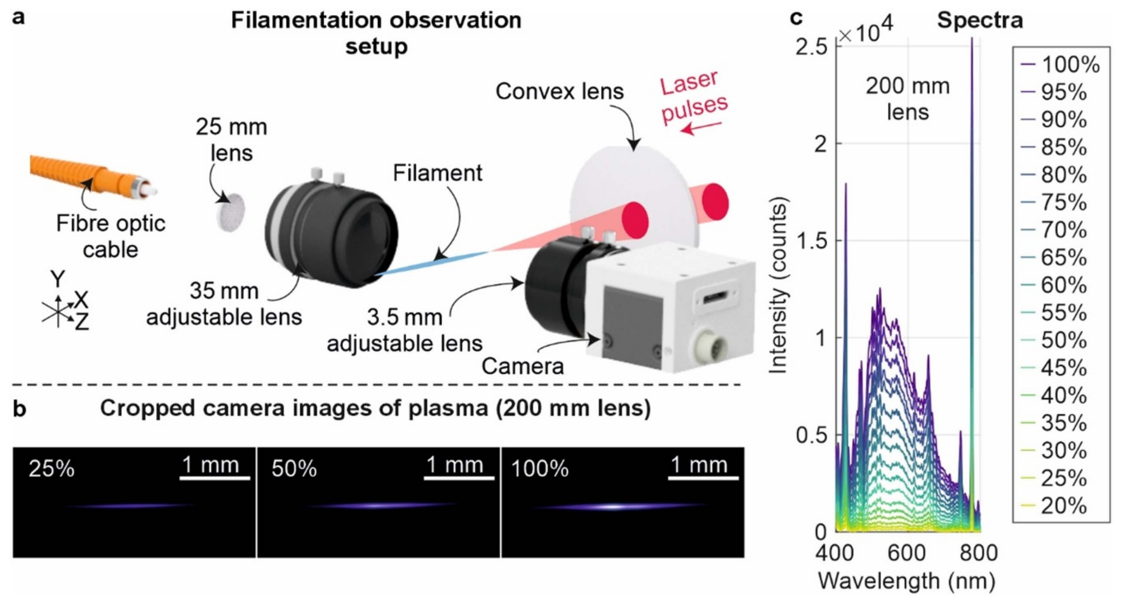

Figure 1. (a) Schematic of the experimental setup showing the femtosecond laser source (1030 nm, 190 fs), focusing lens (200 mm focal length in this instance, but changeable), side-imaging camera, and fibre-coupled spectrometer. The camera and spectrometer were positioned orthogonally to the filament axis. (b) Cropped side-view images of plasma generated using different powers. (c) Corresponding emission spectra for different laser power percentages of the maximum using a 200 mm focal length lens, showing broadband continuum features and relative intensity variation.

2.2. Neural networks

To assess the spectral fingerprints of filamentation and their diagnostic value, we employed two complementary neural networks, namely 1) direct CNN classification and 2) an NMF–CNN regression approach. The CNN classifier provides categorical predictions of filamentation onset, whilst the NMF–CNN regression yielded spectral components from input plasma images. The neural networks were trained and analysed using MATLAB R2024a on Windows 11 Enterprise Dell 7865 Workstation, that had 2 × NVIDIA RTX A6000 (each with 48 GB GDDR6 VRAM), 128 GB system memory and AMD Ryzen Threadripper PRO 5975WX 32-Cores (3.60 GHz) CPU.

2.3. Image classification CNN

To establish a robust diagnostic for filament onset, we trained a two-dimensional CNN directly on plasma emission images. As seen in figure 1(c), broadband plasma emission in the 500–700 nm region captures the key light-matter interactions during filament formation and was therefore used as a diagnostic to label CNN training images. For each lens and input power, the integrated spectral intensity in the 500–700 nm band of each spectrum was then computed as a function of input power. To objectively determine the filament onset threshold, we calculated the first and second derivatives of the broadband intensity with respect to power. In this work the inflection point, identified by a zero-crossing in the curvature (d^2I/dP^2), marked the onset of filamentation. All image frames acquired below this threshold were labelled as 'no filament', whilst those at and above the threshold were labelled as 'filament'. This derivative-based approach provided a reproducible and physically interpretable criterion for CNN supervision, directly linking plasma emission images to the nonlinear growth of broadband light. The broadband integrated spectral intensity curves and their corresponding inflection-point onsets are shown in figure 2. The filament onset thresholds determined from the inflection point in broadband growth were 65%, 80%, 100%, 85%, 75%, 80% and 100% for lenses of 75-, 100-, 125-, 150-, 200-, 250- and 300 mm focal length, respectively. These values quantify the onset of nonlinear plasma-driven emission rather than the subsequent clamped regime and provided the ground-truth labels for CNN training. These absolute onset thresholds are internally consistent within the present dataset, but they may vary between data-acquisition sessions on different days, as modest changes in overall laser alignment, coupling or delivered power can shift the input level at which super-linear broadband growth first appears, even though the underlying physical transition remains the same. The 75-, 100-, 125-, 150-, and 300 mm focal length lens data were used to train the neural network, and the 200 mm and 250 mm lens data were left out and used for testing the neural network.

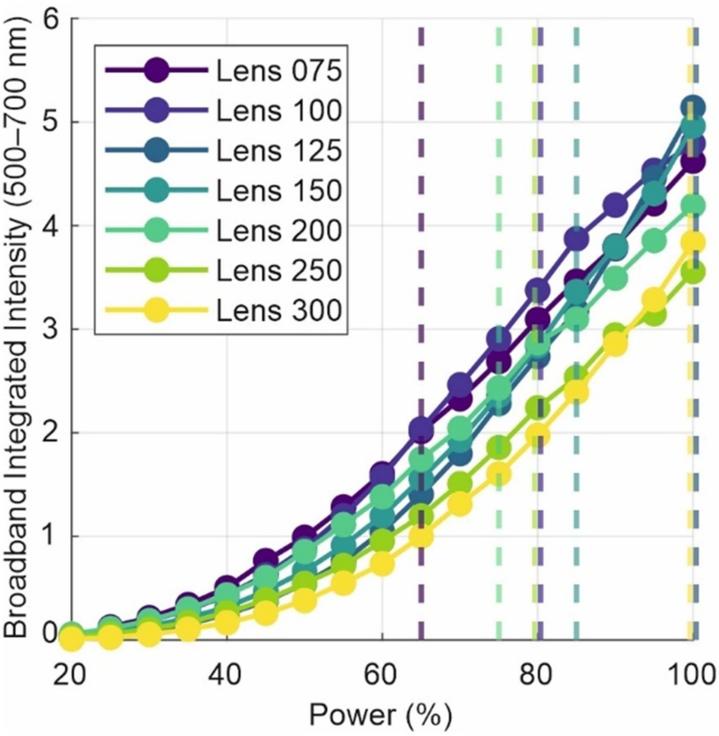

Figure 2. Broadband integrated spectral intensity (500–700 nm) as a function of input power for different focusing lenses. The onset of filamentation for each lens is determined by the inflection point in the energy growth curve, identified from a zero-crossing in the second derivative (d^2I/dP^2). Vertical dashed lines indicate the filamentation onset power, which defines the ground-truth labels for training the CNN.

The CNN comprised four convolutional blocks with increasing filter depth (16, 32, 64, and 128 filters; 3 × 3 kernels), each followed by batch normalisation, rectified linear unit activation [26], and max-pooling. After the convolutional stages, the network included a fully connected layer with 256 units, a dropout layer (rate of 0.5) to reduce overfitting, and a final fully connected layer with two outputs, followed by a softmax activation for binary classification. Training used the Adam optimiser [27] with an initial learning rate of 1×10^{-5}, L2 regularisation of 1×10^{-4}, a mini-batch size of 32, and up to 1000 epochs with learning-rate reduction on plateau. Validation was performed every 100 iterations using a held-out dataset. This architecture and training strategy were selected to evaluate the network's ability to generalise filament detection to previously unseen focusing geometries.

2.4. NMF–CNN hybrid

Whilst image-based classification of plasma enables robust detection of filament onset, spectra contain richer information about nonlinear processes such as supercontinuum generation, ionisation, and fluorescence. To couple these two diagnostic channels, we developed a hybrid framework that combines dimensionality reduction of measured spectra with CNN-based regression from plasma images. Spectra were interpolated onto a fixed 400–800 nm wavelength (since signal was minimal after 800 nm) grid (512 points) and averaged across measurements, yielding 1 × 512 input vectors for subsequent NMF analysis.

Specifically, measured spectra $\mathbf{s}_i \in \mathbb{R}^M$ (with $M = 512$ wavelength bins) were decomposed using NMF, yielding a reduced representation:

$$\mathrm{s}_i = \sum_{k=1}^{K} c_{ik} b_k,\ c_{ik} \geqslant 0$$

where $\mathbf{b}_k$ are the NMF basis spectra, and c_{ik} are the corresponding coefficients. Here $K \ll M$, with K $\sim$5 components found sufficient to capture the majority of spectral variance. Each experimental spectrum was thus represented as a coefficient vector $\boldsymbol{c}_i = (c_{i1},\ldots,c_{iK})$. To address potential correlations from a shared experimental setup, we normalised spectra to counts per second and interpolated to a common wavelength grid. We trained on NMF coefficients standardised across the dataset.

A 2D CNN was trained to map cropped plasma images directly to these NMF coefficient vectors. Note that the CNN does not predict spectra directly but predicts the latent representation in NMF space. The CNN was trained with mean squared error (MSE) loss function on standardised NMF coefficient targets and consisted of four convolutional stages (16–128 filters, batch normalisation, pooling, dropout), a global average pooling layer and a 256-unit fully connected layer, was trained to map plasma images directly to the NMF coefficient vectors. Training employed the Adam optimiser with a learning rate of 1×10^{-4}, L2 regularisation of 1×10^{-4}, minibatch size of 32, and a total number of 500 epochs. The same data was used as for the CNN, along with the same lens-based training and testing split.

Following training, test images were passed through the CNN to predict standardized NMF coefficient vectors $\hat{c}_i \in \mathbb{R}^K$. These were converted back to physical coefficients by inverting the z-scoring applied during training:

$$c_i = \widehat{c_i} \odot \sigma + \mu$$

where $\mu, \sigma \in \mathbb{R}^K$ are the training-set mean and standard deviation, and $\odot$ denotes elementwise multiplication. The full spectrum was then reconstructed by multiplying the coefficient vector with the fixed NMF basis matrix $H \in \mathbb{R}^{K \times M}$ (with $M = 512$ wavelength bins):

$$\hat{s}_i = c_i^{\top} H.$$

Performance was quantified using per-sample MSE across wavelengths and the coefficient of determination:

$$R^2 = 1 - \frac{\sum (s_i - \hat{s}_i)^2}{\sum (s_i - \bar{s}_i)^2}.$$

Note that R^2 here refers to the global goodness-of-fit for each reconstructed spectrum.

This hybrid NMF–CNN model effectively links spatial emission patterns to spectral content, revealing latent variables underlying filament formation. By learning a forward mapping from plasma images to spectral signatures, the neural network provides a data-driven diagnostic that unifies spatial and spectral information, enabling potential physical interpretation of nonlinear propagation regimes from captured plasma images. Because NMF optimises the basis H to best represent the spectral distribution present in a dataset, the components derived from lens measurements acquired during a separate experimental run (where spectra were collected across all lenses and power levels) may differ slightly in shape or ordering, for example due to small variations in beam size, alignment, or power stability. This behaviour is expected, as NMF solutions depend on the underlying spectral structure and are inherently non-unique under permutation and scaling, whilst remaining fixed and self-consistent within each dataset. To enable cross-lens comparison, components were aligned across the 200 mm and 250 mm datasets using spectral correlation so that components 1–5 correspond to the same underlying spectral features in all figure.

3. Results and discussion

To determine the performance of the CNN used to classify the onset of filamentation, we compared its predictions against ground truth labels for filament onset across a range of laser powers using a 200 mm and 250 mm focal length lens. The binary classification results are shown in figure 3. In addition, to evaluate filamentation onset detection, we compared a CNN classifier with a baseline linear regression approach across two lenses. Filamentation onset was defined as the power threshold beyond which plasma emission becomes significant, determined using spectral curvature analysis. The baseline model was implemented as a least-squares linear regression using MATLAB's 'fitlm' function. A single scalar predictor was derived from the total integrated spectral intensity for each laser power, representing overall brightness. This predictor was normalised for stability and mapped to the binary ground-truth label (No filament = 0, filament = 1). The linear regression output was thresholded at 0.5 to generate binary predictions, producing a simple 'linear classification' of filament onset versus power. This baseline assumes filament onset is driven solely by an overall increase in generated light intensity, without accounting for spatial structure or morphological changes in the plasma pattern. In contrast, the CNN uses full spatial image data that includes filament morphology and localised brightness variations.

As seen in the figure, for the 200 mm focal length lens, the CNN predicted onset at 72% power versus a ground truth of 75%, yielding an error of 3%, whilst linear regression predicted 87% with an error of 12%. For the 250 mm lens, CNN predicted 76% versus a ground truth of 80% with an error of 4%, compared to regression at 87% with an error of 7%. Across both lenses, CNN achieved a mean

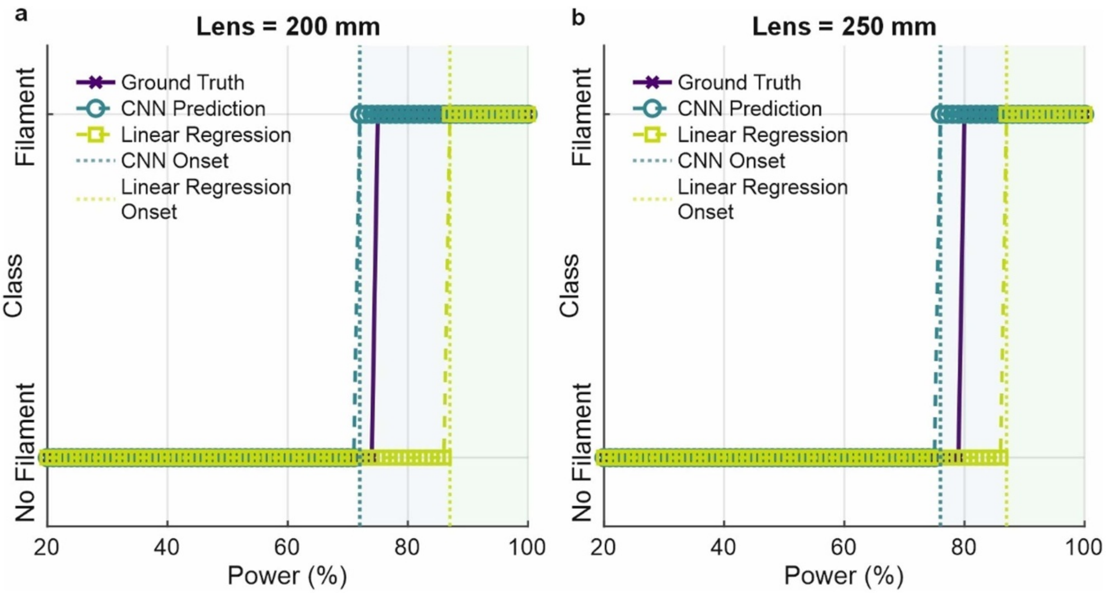

Figure 3. Filament classification results for data from using a lens of focal length (a) 200 mm and (b) 250 mm. Ground truth labels (blue, '×'), CNN predictions (teal, '○'), and linear regression predictions (lime green, '□') are plotted against laser power (%). The predicted onset of filamentation from the CNN and baseline regression models are indicated by dotted teal and dotted green vertical lines, respectively, with shaded regions highlighting the corresponding post-onset regimes.

absolute error of 3.5% compared to 9.5% for regression, representing an improvement of approximately 63%, with the linear regression tending to overestimate onset thresholds. These results demonstrate that incorporating spatial features from plasma images significantly improves onset detection accuracy compared to intensity-only models.

Following CNN classification, we tested the trained NMF–CNN on data from 200 mm and 250 mm lenses. In addition, to help quantify the NMF–CNN, we compared the NMF–CNN with a simple linear regression model that predicted the full spectrum from a single scalar image feature (the summed pixel values within the plasma image). The linear regression model was trained across all powers and fitted independently for each wavelength, producing slope and intercept values per spectral channel.

Predictions for the test dataset were obtained by applying these per-wavelength coefficients to the corresponding image intensities. Figures 4(a) and (b) show representative plasma images together with reconstructed spectra from the linear model and the NMF–CNN for the (a) 200 mm and (b) 250 mm lenses at input powers of 25%, 50%, 75% and 100%.

Whilst the linear regression captures the overall trend of increasing intensity with power, it systematically overpredicts spectral intensity at higher powers and fails to reproduce filament-enhanced line features. In contrast, the NMF–CNN provides substantially better agreement with the true spectra across all powers, with particularly strong improvements at filament onset and above.

To disentangle the complex spectral features generated during filamentation, figure 5 presents the relationship between NMF-derived spectral components and their spatial relevance using Grad-CAM [28] spatial attention maps that highlighted regions of the plasma image driving the predictions. For each condition, we plotted the first five NMF components alongside their Grad-CAM intensity maps (128 × 128 pixels), which are overlaid with 50% transparency on the plasma images. Figure 5 shows all five NMF components at 25% and 100% power for data from the (a) 200 mm and (b) 250 mm focal length lenses. For both lenses, at 25%, all components are weak, consistent with pre-filament conditions, and component 1 is nearly absent for the 250 mm focal length lens. At 100% for both lenses, component 1 dominates and becomes highly localised in the Grad-CAM overlays, reflecting broadband continuum emission from the filament core, whilst components 2 and 3 strengthen as they capture broad spectral regions around the halo of the filament. Components 4 and 5 feature line rich spectra and are spatially extended, as shown in the grad-CAM overlays.

These qualitative trends are supported by the integrated intensity curves in figure 6. For the 200 mm lens, component 1 rises sharply near 60%–70% power and saturates around 90%, consistent with filament onset and intensity clamping. Components 2 and 3 increase later, indicating the emergence of strong line emission as plasma develops. For the 250 mm lens, filament onset is delayed, with

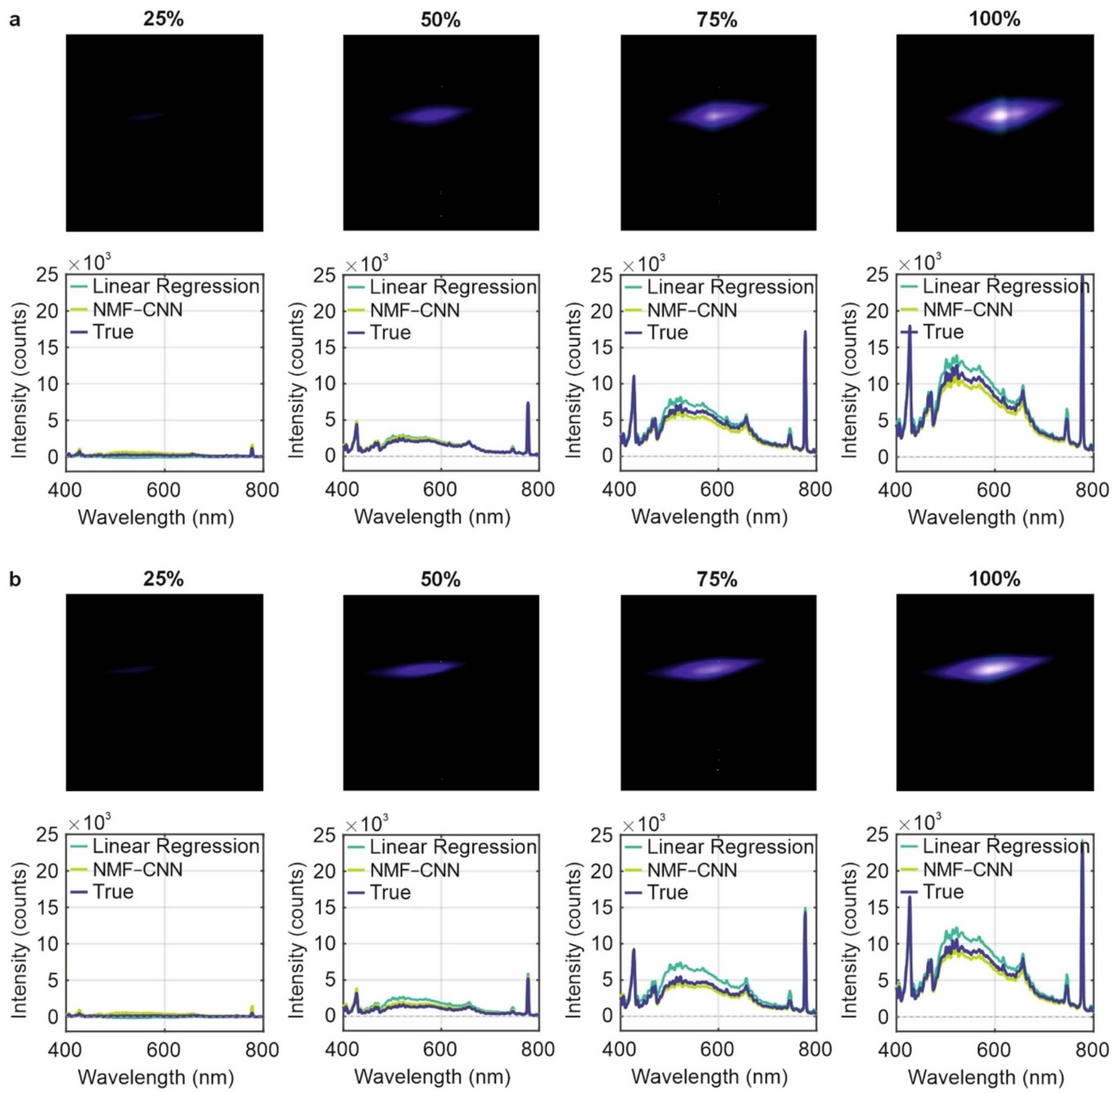

Figure 4. Representative results comparing the NMF–CNN regression model and a linear regression baseline trained to predict plasma emission spectra from image-derived features. Four test samples are shown for the (a) 200 mm focal length lens and (b) 250 mm focal length lens at input powers of 25%, 50%, 75%, and 100%. The top row shows the resized plasma images used for training (resized to 128 × 128 pixels, which stretches the vertical axis and makes the filament appear wider than in raw frames). The bottom row compares the ground-truth spectra (blue) with predictions from the linear regression model (green) and the CNN-based model (yellow).

component 1 beginning to rise at closer to 75%–80% power and saturating near 100%, reflecting the influence of focal length on filament formation. Components 2 and 3 increase later, reflecting the emergence of line emission, whilst components 4 and 5 remain weak across all powers. Together, figures 6 and 7 show that NMF decomposition isolates continuum and line contributions and provides a clear signature of filament onset and saturation, demonstrating that the NMF–CNN leverages spatial morphology beyond simple intensity scaling.

To further quantify these spatial–spectral relationships, we analysed Grad-CAM attribution trends for all NMF components across power levels in 5% steps. As shown in figure 7, we computed the normalised distance between the brightest Grad-CAM activation and the beam intensity peak, capturing spatial migration of spectral features. Component 1, associated with broadband continuum, initially shifts outward with increasing power for both lenses. For the 200 mm lens, this offset peaks near 70% power and then returns closer to the beam core at 100%, consistent with clamping and core re-intensification. In contrast, for the 250 mm lens, the outward migration persists to high power, reflecting a more distributed filament and weaker core confinement. For both lenses, components 4 and 5, which exhibit spiky, line-rich spectra, show large outward shifts, consistent with peripheral plasma regions where discrete transitions dominate. Components 2 and 3 remain comparatively close to the beam centre throughout, reflecting features that are less sensitive to filament widening. Whilst some fluctuations are present, these trends demonstrate that the CNN captures progressive morphological changes beyond simple power

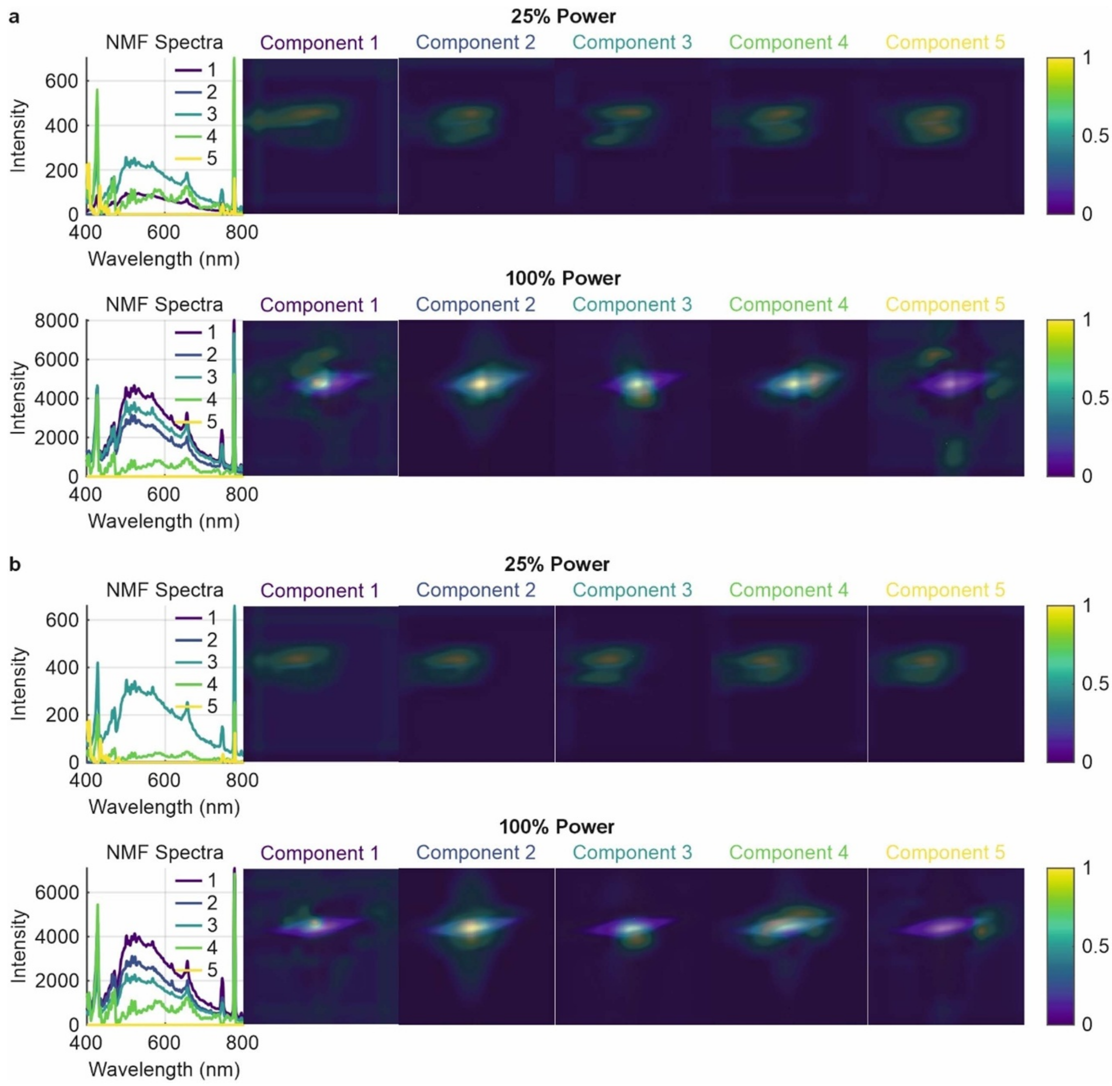

Figure 5. NMF component spectra (first column) and Grad-CAM activations map (columns 2–6, 50% transparency) overlaid on test images for components 1–5 at 25% laser power (top row) and 100% power (bottom rom) for (a) 200 mm focal length lens and (b) 250 mm focal length lens.

scaling, providing a quantitative link between spectral fingerprints and filament dynamics. Importantly, this spatial–spectral attribution is lens-dependent, revealing how focusing conditions influence the origin and localisation of continuum and line emissions. These observations align with known physics, yet our approach introduces a data-driven, interpretable metric that is not typically accessible from conventional diagnostics and could enable real-time control strategies for nonlinear propagation using reinforcement learning agents.

To quantify spectral prediction performance, absolute RMSE was computed across the full wavelength range, the 500–700 nm band, and equal 100 nm bins between 400 nm and 800 nm. Figure 8 shows median RMSE as a function of power and wavelength band. Errors are minimal at low powers across all bands but increase near filament onset, with the largest growth in the 500–600 nm region. The 500–700 nm band, which captures key filament features, remains relatively stable until onset and then rises moderately at full power. The 400–500 nm and 700–800 nm bands consistently show the lowest RMSE.

To evaluate prediction stability across the full power range, we analysed power-dependent performance across the entire wavelength domain. Figures 9(a) and (b) show median R^2 (coefficient of determination, value closer to 1 is better) and RMSE (root mean square error, lower value is better) versus power for both the NMF–CNN model and a linear regression model using the 200 mm and 250 mm lenses, alongside image and spectral SNR trends. For both lenses, the NMF–CNN exhibits poor performance at low powers (20%–40%), where image and spectral SNR are small, making NMF coefficient estimation unstable and driving R^2 strongly negative (performing worse than a mean-only baseline). As power increases, SNR improves steadily, and the NMF–CNN transitions sharply to high accuracy near

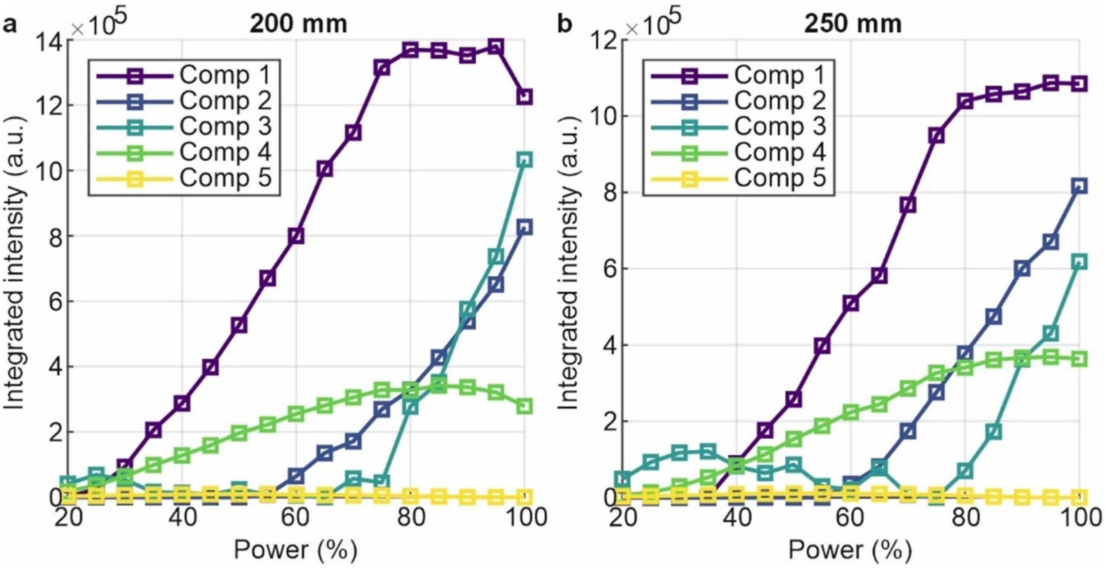

Figure 6. Integrated intensity of NMF components as a function of power for (a) 200 mm and (b) 250 mm focal length lenses. Component 1 rises sharply near filament onset and saturates at high power, consistent with intensity clamping.

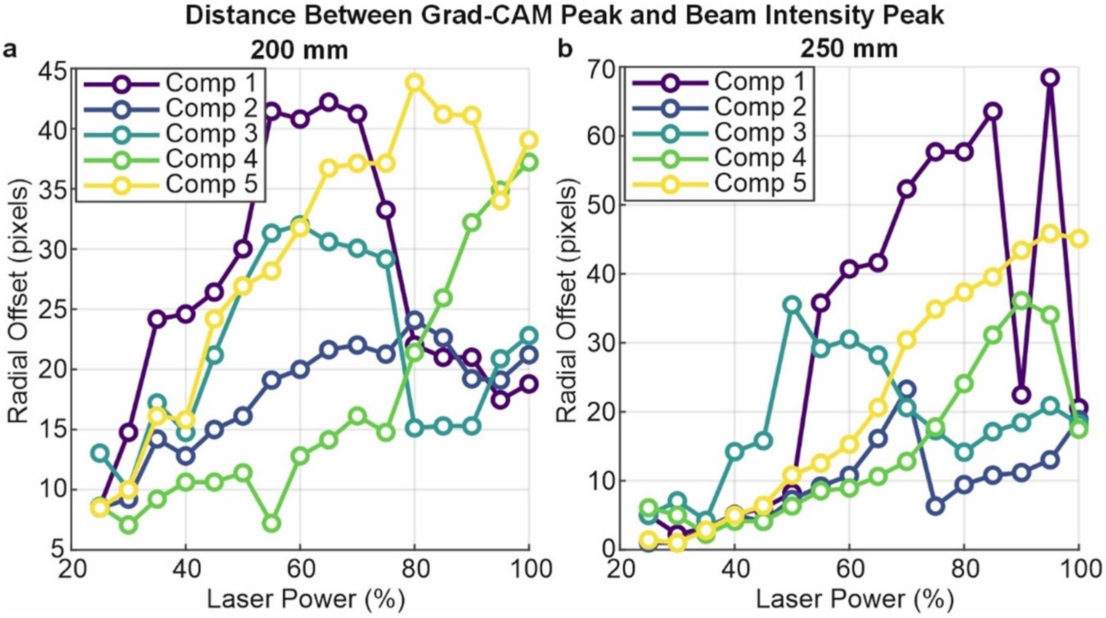

Figure 7. Radial offset between the Grad-CAM activation peak and the beam intensity peak for NMF components across laser power for (a) 200 mm lens and (b) 250 mm lens. Each curve tracks the radial displacement of the component-specific Grad-CAM maximum as power increases. These trends provide quantitative evidence of spatial–spectral coupling beyond simple power scaling.

filament onset (approximately 75% for the 200 mm lens and 80% for the 250 mm lens), with median R^2 approaching unity and RMSE stabilizing despite rising absolute intensity.

Although the absolute RMSE at high powers (e.g. ∼800 counts) appears large, this corresponds to only 10%–15% relative error given the 5000–8000-count dynamic range of the spectra in this regime. At low powers (<40%), the 250 mm lens exhibits higher RMSE than the 200 mm case because its plasma emission is weaker, leading to lower image and spectral SNR and noisier NMF coefficient estimates. Once power increases and SNR improves, both lenses show stable and accurate NMF–CNN predictions across the full spectral range.

The linear baseline performs slightly better than the NMF–CNN at very low powers but generally shows poorer agreement overall. Specifically, the R^2 remains negative at low powers, improves modestly mid-range, and then degrades again at high powers for the 250 mm lens. RMSE for the baseline is

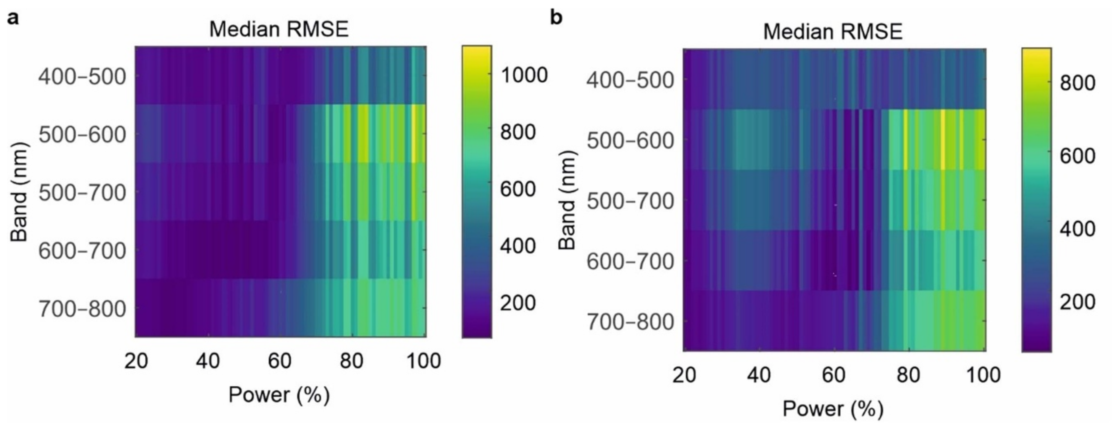

Figure 8. Median RMSE versus power for five wavelength bands for (a) 200 mm focal length lens and (b) 250 mm focal length lens. The colour bar indicates RMSE magnitude (blue is low error, yellow is high error).

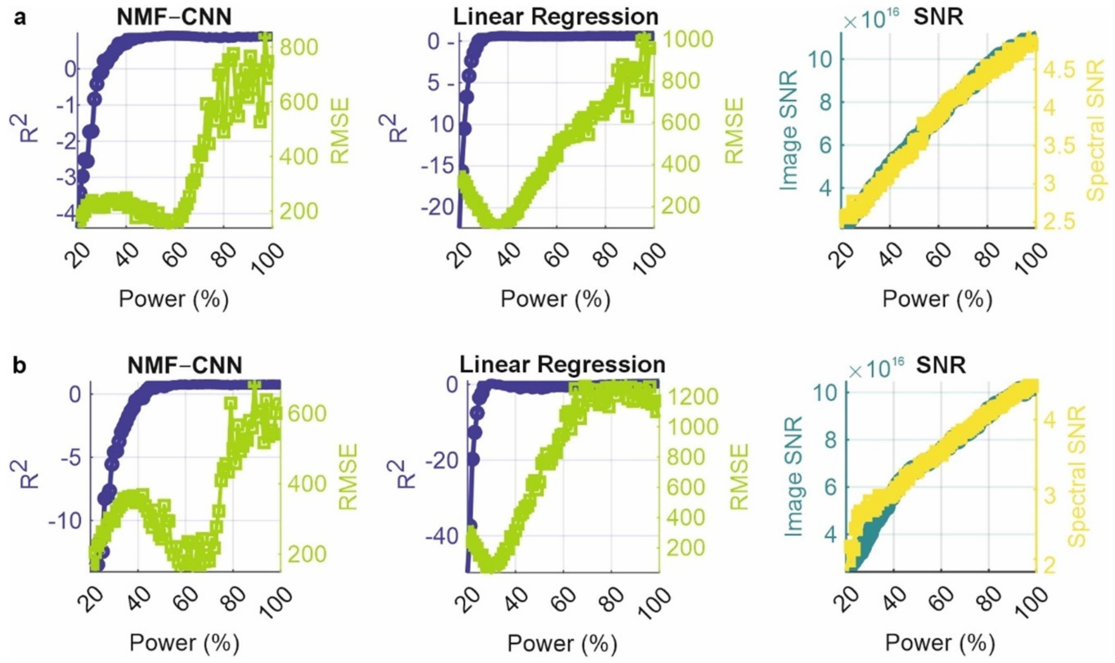

Figure 9. Power-dependent performance and SNR analysis for (a) 200 mm focal length lens and (b) 250 mm focal length lens. (Left) NMF–CNN: median R^2 (blue, left axis) and RMSE (green, right axis) versus power. (Middle) linear regression: median R^2 and RMSE versus power. (Right) image and spectral SNR versus power.

generally higher than NMF–CNN beyond filament onset, reflecting its inability to capture nonlinear spectral changes associated with filamentation. These trends are consistent with aggregate metrics, where for the 200 mm lens, the NMF–CNN achieves a median R^2 of 0.960 versus 0.895 for the linear regression, whilst for the 250 mm lens, NMF–CNN reaches 0.945 compared with 0.208 for the linear regression. Overall, the increased errors at low powers arise from the very low signal-to-noise ratio rather than model instability, and the predictions stabilise as soon as the plasma emission becomes measurable.

To assess generalisation beyond the dataset used for training, we applied the NMF–CNN model to a set of measurements acquired on different days under slightly different laser alignment. This independent dataset includes data from lenses 75, 100, 125, 150, 200, 250 and 300 mm, none of which were used during training the neural network. Figure 10 shows the measured and reconstructed spectra at 100% power. For clarity of comparison, the true spectra are normalised to a peak value of one and the NMF–CNN spectra are scaled using the same respective normalisation, ensuring that relative spectral shape differences are visually emphasised. The NMF–CNN recovers the broadband continuum and dominant spectral features for most lenses, with an average R^2 of 0.662 ± 0.232 across all focal lengths. Performance remains strong across all lenses. These results suggest that the model retains predictive

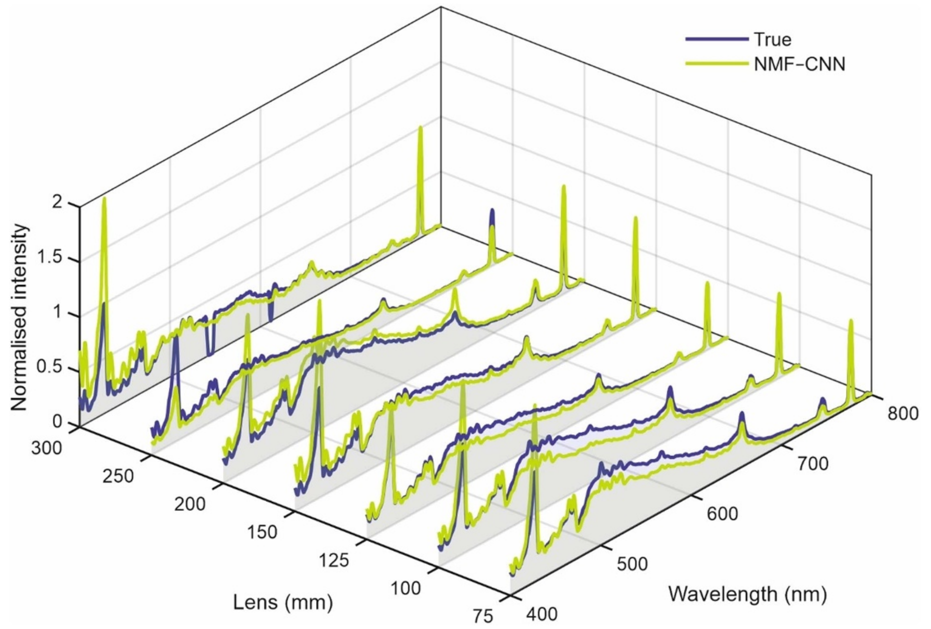

Figure 10. 3D normalised waterfall comparison between measured spectra (blue) and NMF–CNN-reconstructed spectra (green) at 100% power for lenses 75–300 mm. All data were acquired on different days and under slightly different laser alignment conditions than those used to train the model. For clarity, each true spectrum is normalised to its peak value, and the corresponding NMF–CNN prediction is scaled using the same per-lens normalisation factor. This cross-day evaluation demonstrates that the NMF–CNN retains predictive capability on independent measurements with modest alignment and brightness differences.

capability on independently acquired data, providing preliminary evidence for generalisation across measurement conditions.

Building on this foundation, the regression NMF–CNN extends the diagnostic power by reconstructing the full emission spectrum from single plasma images. The high predictive accuracy across the visible range demonstrates that the spatial features of observable plasma emission encode detailed spectral information, including the emergence of broadband continua and spectral peaks that evolve with increasing power. This dual approach bridges onset detection and spectral characterisation, showing how image-based deep learning can provide both discrete and quantitative insight into ultrafast nonlinear light–matter interactions.

Prediction accuracies are likely affected by experimental uncertainties such as spectral lens alignment, camera calibration settings, laser fluctuations, and air-current driven filament instabilities. Data and training limitations, including the diversity of the dataset, preprocessing choices, and the risk of overfitting, also contribute to residual differences between predicted and measured spectra. Despite these factors, the two complementary CNN approaches establish a methodology for filamentation diagnostics from plasma emission images. The classification CNN, trained using derivative-based onset labels, provides a robust method for identifying the threshold at which filamentation begins. This links image morphology directly to the nonlinear growth of broadband emission. Building on this foundation, the regression NMF–CNN extends the diagnostic power by reconstructing the full emission spectrum from single plasma images. The high predictive accuracy across the visible range demonstrates that the spatial features of observable plasma emission encode detailed spectral information, including the emergence of broadband continua and spectral peaks that evolve with increasing power. This dual approach bridges onset detection and spectral characterisation, showing how image-based deep learning can provide both discrete and quantitative insight into ultrafast nonlinear light–matter interactions. By coupling interpretable attribution with spectral reconstruction, the framework offers a pathway to physics-informed diagnostics, enabling future studies to probe filamentation dynamics and optimize experimental conditions using image-based feedback.

Although generalisation was evaluated by testing on lenses excluded from training, the limited diversity of experimental conditions may restrict applicability to other setups or environmental factors. To improve reconstruction of spectral features, future work should focus on expanding the dataset by orders of magnitude to support advanced architectures that enable greater generalisation and improved peak resolution. This expansion will include increased diversity across focal lengths, pulse regimes, and environmental conditions (i.e., turbulence), providing the foundation for deeper networks and capturing rare spectral features. With this richer dataset, architectural enhancements such as attention mechanisms [29] for fine-feature recovery and multi-scale CNNs [30] for hierarchical representation can be explored. These approaches specifically target the challenge of reconstructing sharp spectral peaks noted in this work. Whilst resizing to 128 × 128 pixels stretched the vertical axis, consistent preprocessing and labelling across data for all lenses and powers ensured that the CNN learned consistent spatial features. Nonetheless, future work will maintain aspect ratios to preserve physical fidelity when integrating physics-based information into neural networks for further control and understanding of filamentation.

Physics-informed constraints such as known emission line positions could be integrated to improve physical plausibility, along with adaptive NMF decomposition (where component structure and resolution adjust dynamically to better separate broadband continua from narrow spectral peaks) for better handling of broadband and narrow peaks. Additionally, capturing the 3D nature of filaments through multi-view or volumetric imaging (e.g., stereoscopic cameras or tomographic reconstruction) will likely enhance and refine existing spatial–spectral correlations, improving resolution and enabling architectures such as physics-informed neural networks [31] for developing generalised models.

4. Conclusion

In conclusion, we have introduced a data-driven framework for detecting and characterising femtosecond laser filamentation using optical imaging combined with deep learning. For two different lens setups, a neural network achieved higher accuracy in identifying the onset of filamentation directly from plasma emission images, compared with a linear regression model, with an onset prediction only 3.5% absolute error from the ground-truth determined threshold for the CNN and 9.5% absolute error for the linear regression model. This capability suggests that CNN-based diagnostics may enable predictive control of nonlinear propagation in future high-power laser systems. Additionally, a regression NMF–CNN demonstrated that plasma images encode sufficient spatial information to reconstruct broadband emission spectra with higher fidelity than a linear regression model, yielding a median coefficient of determination of 0.953 and 0.552 across the test data, respectively for the two types of models. Deviations between the ground-truth and predicted spectra likely occurred due to weak signal or noisy plasma images cases prior to onset, underscoring the physical link between spatial emission structure and spectral broadening.

These results highlight how combining dimensionality reduction with deep learning enables both accurate prediction and interpretable analysis of filament formation. Beyond serving as a diagnostic, this work demonstrates a methodology toward adaptive control of ultrafast nonlinear light–matter interactions. Such capabilities have important implications for controlled energy delivery, probing of strong-field physics, and the development of laser-based sensing platforms.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: https://doi.org/10.5258/SOTON/D3714 [32].

Funding sources

This work was funded by the Engineering & Physical Sciences Research Council (Grants EP/W028786/1, EP/T026197/1, EP/Z002567/1).

Conflict of interest

The authors declare no competing financial interest.

Author contributions

James A Grant-Jacob 0000-0002-4270-4247
Conceptualization (lead), Data curation (lead), Formal analysis (lead), Investigation (lead), Methodology (lead), Project administration (lead), Software (lead), Validation (lead), Visualization (lead), Writing – original draft (lead), Writing – review & editing (equal)

Ben Mills 0000-0002-1784-1012
Funding acquisition (lead), Resources (lead), Writing – review & editing (equal)

References

[1] Bergé L, Skupin S, Nuter R, Kasparian J and Wolf J-P 2007 Ultrashort filaments of light in weakly ionized, optically transparent media *Rep. Prog. Phys.* **70** 1633

[2] Schmitt-Sody A, Kurz H G, Bergé L, Skupin S and Polynkin P 2016 Picosecond laser filamentation in air *New J. Phys.* **18** 093005

[3] Couairon A and Mysyrowicz A 2007 Femtosecond filamentation in transparent media *Phys. Rep.* **441** 47–189

[4] Polynkin P and Moloney J V 2011 Optical breakdown of air triggered by femtosecond laser filaments *Appl. Phys. Lett.* **99** 151103

[5] Zhao J *et al* 2016 Strong spatial confinement of terahertz wave inside femtosecond laser filament *ACS Photonics* **3** 2338–43

[6] Chu W, Zeng B, Li Z, Yao J, Xie H, Li G, Wang Z and Cheng Y 2017 Range extension in laser-induced breakdown spectroscopy using femtosecond–nanosecond dual-beam laser system *Appl. Phys. B* **123** 173

[7] Dicaire I, Jukna V, Praz C and Milián C 2016 Spaceborne laser filamentation for atmospheric remote sensing *Laser Photon. Rev.* **10** 481–93

[8] Eisenmann S, Louzon E, Katzir Y, Palchan T, Zigler A, Sivan Y and Fibich G 2007 Control of the filamentation distance and pattern in long-range atmospheric propagation *Opt. Express* **15** 2779–84

[9] Reyes D, Baudelet M, Richardson M and Rostami Fairchild S 2018 Transition from linear- to nonlinear-focusing regime of laser filament plasma dynamics *J. Appl. Phys.* **124** 053103

[10] Noack J and Vogel A 1999 Laser-induced plasma formation in water at nanosecond to femtosecond time scales: calculation of thresholds, absorption coefficients, and energy density *IEEE J. Quantum Electron.* **35** 1156–67

[11] Vogel A, Noack J, Nahen K, Theisen D, Busch S, Parlitz U, Hammer D X, Noojin G D, Rockwell B A and Birngruber R 1999 Energy balance of optical breakdown in water at nanosecond to femtosecond time scales *Appl. Phys. B* **68** 271–80

[12] De Giacomo A, Gaudiuso R, Dell'Aglio M and Santagata A 2010 The role of continuum radiation in laser induced plasma spectroscopy *Spectrochim. Acta B* **65** 385–94

[13] Hossain A M, Ehrhardt M, Rudolph M, Lorenz P, Kalanov D, Zimmer K and Anders A 2023 Time- and position-dependent breakdown volume calculations to explain experimentally observed femtosecond laser-induced plasma properties *ACS Photonics* **10** 1232–9

[14] HajiRassouliha A, Taberner A J, Nash M P and Nielsen P M F 2018 Suitability of recent hardware accelerators (DSPs, FPGAs, and GPUs) for computer vision and image processing algorithms *Signal Process. Image Commun.* **68** 101–19

[15] Zeiler M D and Fergus R 2014 Visualizing And Understanding Convolutional Networks *Computer Vision—ECCV 2014* ed D Fleet, T Pajdla, B Schiele and T Tuytelaars (Springer International Publishing) pp 818–33

[16] Krizhevsky A, Sutskever I and Hinton G E 2012 ImageNet classification with deep convolutional neural networks. *Proc. of the 25th Int. Conf. on Neural Information Processing Systems* vol **1** (Curran Associates Inc.) pp 1097–105

[17] Karpathy A, Toderici G, Shetty S, Leung T, Sukthankar R and Fei-Fei L 2014 Large-scale video classification with convolutional neural networks *2014 IEEE Conf. on Computer Vision and Pattern Recognition* (IEEE) pp 1725–32

[18] Juhl M and Leosson K 2020 Polarimetry with disordered photonic structures *ACS Photonics* **7** 203–11

[19] Zhang J, Zhang S, Shi W, Hu Y, Dong Z-G, Li J and Lu W 2024 Advanced image classification using a differential diffractive network with "learned" structured illumination *ACS Photonics* **11** 5289–98

[20] Palounek D, Vala M, Bujak Ł, Kopal I, Jiříková K, Shaidiuk Y and Piliarik M 2024 Surpassing the diffraction limit in label-free optical microscopy *ACS Photonics* **11** 3907–21

[21] Grant-Jacob J A, Zervas M and Mills B 2024 Imaging pollen using a Raspberry Pi, LED and deep learning *Sci. Total Environ.* **955** 177084

[22] Ma W, Liu Z, Kudyshev Z A, Boltasseva A, Cai W and Liu Y 2021 Deep learning for the design of photonic structures *Nat. Photonics* **15** 77–90

[23] Jiang J, Chen M and Fan J A 2020 Deep neural networks for the evaluation and design of photonic devices *Nat. Rev. Mater.* **6** 679–700

[24] Jolliffe I T and Cadima J 2016 Principal component analysis: a review and recent developments *Philos. Trans. R. Soc. A* **374** 20150202

[25] 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

[26] Nair V and Hinton G E 2010 Rectified linear units improve restricted Boltzmann machines *Proc. 27th Int. Conf. on Machine Learning (ICML-10)* pp 807–14

[27] Kingma D P and Ba J 2014 Adam: a method for stochastic optimization. (arXiv:1412.6980)

[28] Selvaraju R R, Cogswell M, Das A, Vedantam R, Parikh D and Batra D 2020 Grad-CAM: visual explanations from deep networks via gradient-based localization *Int. J. Comput. Vis.* **128** 336–59

[29] Soydaner D 2022 Attention mechanism in neural networks: where it comes and where it goes *Neural Comput. Appl.* **34** 13371–85

[30] Han X-H, Zheng Y and Chen Y-W 2019 Multi-level and multi-scale spatial and spectral fusion CNN for hyperspectral image super-resolution *Proc. IEEE/CVF Int. Conf. on Computer Vision Workshops*

[31] Karniadakis G E, Kevrekidis I G, Lu L, Perdikaris P, Wang S and Yang L 2021 Physics-informed machine learning *Nat. Rev. Phys.* **3** 422–40

[32] Grant-Jacob J A and Mills B 2026 "Dataset - Ultrafast laser filamentation classification and analysis via neural networks" *University of Southampton Institutional Repository* (https://doi.org/10.5258/SOTON/D3714)