HIGH POWER LASER
SCIENCE AND ENGINEERING

RESEARCH ARTICLE

Applications of object detection networks in high-power laser systems and experiments

Jinpu Lin[ID], Florian Haberstroh[ID], Stefan Karsch, and Andreas Döpp[ID]

Ludwig-Maximilians-Universität München, Garching, Germany

Abstract

The recent advent of deep artificial neural networks has resulted in a dramatic increase in performance for object classification and detection. While pre-trained with everyday objects, we find that a state-of-the-art object detection architecture can very efficiently be fine-tuned to work on a variety of object detection tasks in a high-power laser laboratory. In this paper, three exemplary applications are presented. We show that the plasma waves in a laser–plasma accelerator can be detected and located on the optical shadowgrams. The plasma wavelength and plasma density are estimated accordingly. Furthermore, we present the detection of all the peaks in an electron energy spectrum of the accelerated electron beam, and the beam charge of each peak is estimated accordingly. Lastly, we demonstrate the detection of optical damage in a high-power laser system. The reliability of the object detector is demonstrated over 1000 laser shots in each application. Our study shows that deep object detection networks are suitable to assist online and offline experimental analysis, even with small training sets. We believe that the presented methodology is adaptable yet robust, and we encourage further applications in Hz-level or kHz-level high-power laser facilities regarding the control and diagnostic tools, especially for those involving image data.

Keywords: high repetition rate; laser–plasma accelerators; machine learning; object detection; optical diagnostics

1. Introduction

High-power laser systems with power reaching the petawatt level and repetition rate at a fraction of a hertz have emerged worldwide in the past few years[1–5]. With the fast development of high-repetition-rate operation capabilities in plasma targetry, high-power laser–plasma experiments can employ statistical methods that require a large number of shots. Studies for real-time optimization using evolutionary algorithms have been reported in recent years[6–11]. As the size of data to process has continued to increase, more advanced machine learning models have attracted increasing attention. By constructing predictive models, machine learning methods are employed to model the nonlinear, high-dimensional processes in high-power laser experiments. Various methods, including neural networks, Bayesian inference and decision trees, have been introduced for optimization tasks and physics interpretation[12–17]. Meanwhile, as the measurement and diagnostic tools evolve, digital imaging is playing an increasingly important role in experiments and, with it, machine learning methods to process image data.

In the case of a laser–plasma accelerator, image-based diagnostics can take a variety of forms, from the optical elements in the high-power laser facility, over shadowgraphy and interferometry of plasma dynamics, to scintillator signals generated by energetic electron or X-ray beams from the accelerator. In particular, the evolving structure of a plasma accelerator is challenging to visualize because of its microscopic size ($\sim 10^{-5}$ m) and its high velocity (approaching the speed of light). With the latest techniques, such as few-cycle shadowgraphy, taking snapshots of the plasma wake structure is enabled in femtosecond resolution over a range of picoseconds[18–20]. The latest generation of laboratory diagnostics for plasma structures is reviewed by Downer *et al.*[21].

In this paper, we demonstrate exemplary applications of an object detection network in the diagnostics in a high-power laser laboratory. We apply the object detector to few-cycle shadowgraphy of plasma waves, to an electron energy spectrometer and to detect optical damages in a high-power laser beamline. The results show that object detection enables possibilities in diagnostics and data analysis that have not yet been achieved using conventional methods. Moreover, due to the fast inference speed of the object detector, it paves the road towards real-time demonstration of such diagnostics during experiments.

Correspondence to: Jinpu Lin, Ludwig-Maximilians-Universität München, Am Coulombwall 1, 85748 Garching, Germany. Email: Lin.Jinpu@physik.uni-muenchen.de

1

2. Object detection algorithms

Since the development of convolutional neural networks (CNNs), computer vision has drawn attention from across various disciplines[22–24]. As a huge breakthrough in image recognition, a CNN allows categorizing images into certain classes. When a CNN classifies an input image, it learns a model that detects the specific patterns on that image. A pattern is detected by a 'filter' matrix, which has a pre-defined size relatively smaller than the size of the input image. It then takes the dot product of the filter matrix with a sub-matrix of the input image (in pixel values) that has the same size. The filter is 'convolved' with the input image as it slides across the entire input image matrix for all sub-matrices of its size, resulting in an output matrix of cross-products. Intuitively, a filter in a CNN is analogous to a neuron in a regular feed-forward neural network, and several filter matrices form a convolutional layer. A complex CNN can have multiple convolutional layers, and the final output matrix is compared with the input image to adjust the values of the filter matrices. This process is repeated over and over until the output matrix is close enough to the input.

An extension to classification tasks in computer vision is object detection. Unlike classification tasks such as image recognition, which assign one single label to the image, object detection aims to identify all the objects of interest in an image, classify each object and assign a label to it, and then locate them by drawing a bounding box around each object. For images with a fixed number of objects, the objects can be found using a standard CNN followed by a fully connected output network layer with a pre-specified length. However, the task becomes much more challenging when the number of interesting objects is not fixed in an image, leading to a varying length of the output layer of the neural network. This happens to be the case for most applications in high-power laser experiments, especially when scanning parameter spaces across various laser and plasma conditions.

Theoretically, the problem can be solved by splitting the image into many regions of interest and coupling a CNN to each region. However, the number of regions could be significant and easily exceed the computational limit. To make it computationally efficient, there are two families of methods to locate and label objects without determining the number of objects in advance. The region-based convolutional neural network (R-CNN)[25] and its later iterations (faster R-CNN, mask R-CNN) use a selective algorithm to propose a reasonable number of regions that may contain bounding boxes. It then applies a CNN to extract features from each candidate region and classify the feature into the known classes using a linear classifier. While the R-CNN is very accurate in locating the objects, its computational cost can be heavy.

The 'you only look once' (YOLO) family of algorithms[26] takes a different approach; a simplified methodology is illustrated in Figures 1(a)–1(c). YOLO algorithms split the image into a pre-determined number of grids, and define multiple bounding boxes for each grid. Unlike the R-CNN, which applies a network to each region, YOLO algorithms apply a single neural network to the full image. The network then predicts a probability for each class for each bounding box. Post-processing is performed to select the best bounding boxes based on the probability and the overlapping conditions regarding their neighbouring boxes. The greatest advantage of YOLO algorithms, as the name suggests, is that they make predictions with a single network evaluation instead of thousands in other methods, such as the R-CNN. Therefore, YOLO algorithms can be two or three orders of magnitude faster than the R-CNN, making it possible
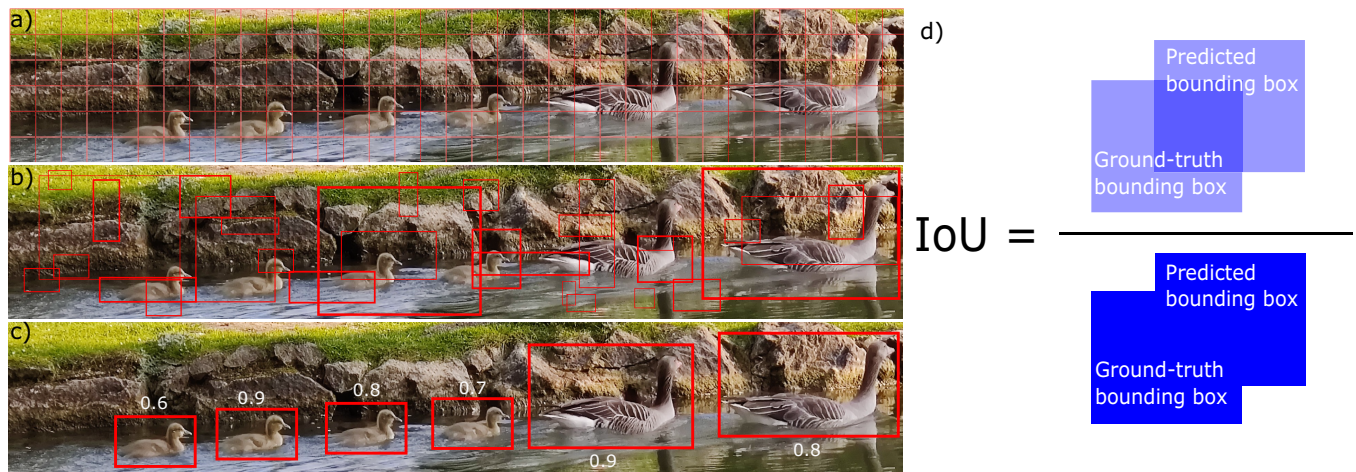
Figure 1. Step-wise illustration of the object detection method. The example image presents ducks creating and surfing on wakefields. (a) Split the image into small grid cells; (b) predict bounding boxes and confidences for each class; (c) final detected objects with confidences; (d) bounding box predicted by the object detector versus the ground-truth bounding box labelled manually. IoU is defined as their area of intersection divided by their area of union, where an ideal object detector would have IoU = 1.

for real-time object detection tasks. This is of particular interest to the community in high-power laser experiments, especially with the development of high-repetition-rate capabilities. However, it has to be pointed out that the YOLO algorithms' superiority in efficiency comes at the cost of prediction accuracy, such as in locating the bounding boxes.

While the original YOLO method was designed to work with square images, it is worth noting that the recent versions allow efficient training on non-square images using rectangular inference. Having such adaptability can benefit our practical applications in laser–plasma experiments that usually have multiple diagnostics with various image resolutions.

As a supervised learning task, validation is needed after training an object detection model. The commonly used evaluation metric in object detection is the intersection over union (IoU). To evaluate the model accuracy on a predicted bounding box, we manually label a ground-truth bounding box and the IoU calculates the area of the intersection as well as the area of the union; see Figure 1(d). The ratio of these two areas is defined as the IoU value between 0 and 1, where 0 means no intersection and 1 means completely overlapping. For a set of images, the performance of the object detection model is evaluated using the mean average precision (mAP), which is obtained by averaging over different IoU thresholds on each bounding box on each image. The box confidence score C is then defined as follows:

$$C = P_{\text{object}} \times \text{IoU},$$

where P_{object} is the probability that the box contains an object. The model considers the prediction to be a true prediction only if the box confidence score is higher than a minimum score. This minimum score is called a 'threshold confidence' and is set manually. For the dataset we use here, the threshold confidence is set to 10%–40% to find most objects of interest while excluding unwanted objects.

The algorithm we use in this project is the state-of-the-art object detector YOLOv5[27], which compared with its predecessors included a new PyTorch training and deployment framework. As a result, YOLOv5 is significantly faster and user-friendly while maintaining good prediction accuracy. Therefore, YOLOv5 is regarded as one of the standard test models when developing specific algorithms in the field of fast object detection.

3. Applications

In this section we are going to present three exemplary applications for object detection in the context of high-power laser experiments.

The experiment was performed at the Center for Advanced Laser Applications at the Ludwig-Maximilians-University of Munich using the ATLAS laser system. The on-target energy of the experiment is 6 ± 1 J. The pulses are centred at 800 nm and compressed to a length of 30 fs (full width at half maximum (FWHM)). Focused with a $f/33$ parabolic mirror, the peak intensity reaches from 1.7×10^{19} to 5×10^{19} W/cm^2, resulting in a normalized vector potential a_0, ranging from 2.8 to 4.8. Laval nozzles fed with hydrogen are used as the gas target and have a diameter of 5 or 7 mm.

3.1. Few-cycle shadowgraphy of plasma waves

Plasma waves excited by a laser-driven electron beam in a hybrid plasma accelerator are diagnosed. A hybrid plasma accelerator utilizes the dense, high-current electron bunch produced by a laser-wakefield accelerator to drive the plasma wave for a plasma-wakefield acceleration (PWFA)[19,28,29]. Unlike in PWFA driven by electron bunches from conventional radio frequency (RF) accelerators, the plasma density in a hybrid accelerator is higher, typically approximately 10^{18} cm^{-3}, which makes it possible for shadowgraphy using few-cycle optical probes[18–20]. The plasma evolution can be observed in detail in femtosecond resolution using a few-cycle probe beam. It is derived from the main laser driver, undergoes spectral broadening in a gas-filled fibre and is compressed to sub-10 fs by a set of chirped mirrors. Thus, the probe and driver are inherently synchronized. A practical problem in experiments is the variation of the plasma waves in the shadowgrams. This especially occurs when the laser–plasma parameters are being tuned, for instance, scanning the plasma target with respect to the laser focus. To locate the plasma waves regardless of the varying laser–plasma condition, an object detector is used.

3.1.1. Labelling and training

The object detector is applied to up to 200 manually labelled shadowgrams taken from various days of experimentation. Datasets of varying sizes are used for training, and a benchmark is listed in Table 1. While most of the shadowgrams have observable plasma waves, about 10% of the images do not. The labelled classes on the shadowgrams include the plasma waves, a shock front caused by a deliberate obstacle in the target's gas flow and the diffraction pattern caused by dust in the imaging beam path. The dataset is randomly split into a training set, a validation set, and a testing set by 70%, 20% and 10%, respectively. To further increase the size of the dataset, augmentations are applied to the labelled images, as shown in the second column in Table 1. In the augmentation process, copies of the original image are made and then the brightness and exposure are slightly changed. Note that augmentation is only applied to the training set and not the validation set or the testing set.

The training process utilizes the concept of transfer learning, where the knowledge from a pre-trained model for general object detection tasks is transferred to our model for a specific task. YOLOv5 provides a series of such

Table 1. Inference accuracy versus dataset size. The first column reports the size of the ground-truth (manually labelled) datasets for training, validation and testing. The second column reports the size of the augmented dataset for training, validation and testing. The third column presents the run time of the training process associated with each dataset, using a Tesla T4 GPU. The last two columns report the prediction accuracy of these datasets on two inference datasets, where inference set 1 has 50 images and inference set 2 has 1000 images.

Labelled set size	Augmented set size	Time	Inference set of 50 accuracy	Inference set of 1000 accuracy
15	35	5 min	68%	52%
30	52	8 min	98%	85%
50	124	11 min	100%	97%
120	299	15 min	100%	92%
200	449	28 min	100%	89%

pre-trained models, and here we use the second-smallest model (YOLO5s.pt). The run time of the training process is listed in the third column in Table 1. It is worth pointing out that the run time can be further reduced by transfer learning from a learned model using a small training set.

3.1.2. Results

In addition to the test dataset, the trained models are applied to two inference datasets, as shown in the last two columns in Table 1. The images in the inference sets are not used in the training, validation or testing processes. The first inference set contains 50 shadowgrams with observable and labelled plasma waves. The second inference set consists of 1000 images from various experiment days, where 68 of them do not have an observable plasma wave. The model's performance on these two inference datasets further proves that the training data are not biased to any specific orientation or location of the objects.

Comparing the five trained networks in Table 1, the medium-sized dataset with 50 pre-labelled shadowgrams provides the most accurate model in this case. The trained model has an mAP of 0.941 for an IoU threshold of 0.5. The model is then used to detect the target features on a shadowgram. An example is presented in Figure 2, showing the detected plasma wave (red), the shock (green) and the diffraction pattern caused by dust in the beam path (blue). A threshold confidence of 10% is applied when drawing the bounding boxes.

The plasma wavelength can be estimated as the plasma wave is located by the object detector. This is achieved by taking the Fourier transform of the region of interest (ROI), which is the red bounding box containing the plasma wave oscillation feature. The result of the Fourier transform is demonstrated on the top-right in Figure 2, and the peak is at approximately 27.5 μm.

As the laser and plasma parameters (pressure, longitudinal position, etc.) are being tuned during an experiment, the plasma wavelength changes accordingly. Figure 3 presents further analysis of the plasma oscillation, given the region defined by the object detector. In Figure 3(a), the backing pressure is scanned from 2 to 7 bar (1 bar $= 10^5$ Pa). Each data point represents the mean value of 20 consecutive laser shots, and the error bar measures the mean absolute deviation. The plasma wavelength is calculated by taking the Fourier transform of the plasma wave ROI at each pressure, and is plotted on the left-hand vertical axis. The electron density is calculated from the plasma wavelength, and is labelled on the right-hand vertical axis. Note that the right-hand vertical axis for the plasma density profile is set to have linear tick labels, and therefore the left-hand vertical axis for the plasma wavelength has nonlinear tick labels. The electron density versus the backing pressure is fit to a linear relation, with an R^2 value as high as 0.98. The curve fitting is shown by the dashed line. The gas target has also been characterized via separate interferometry measurements, and the resulting density is approximately 1×10^{18} cm^{-3} at a backing pressure of 2 bar and approximately 3×10^{18} cm^{-3} at 7 bar.

A similar analysis is presented in Figure 3(b), where the few-cycle optical probe is scanned over 1.2 mm relative to the shock position, from the upstream end to the centre of the gas target. At 0 mm, the first plasma bubble of the plasma wave overlaps with the density shock. The plasma density versus position before the shock is fitted to an exponential function, and the density away from the shock approaching the target centre is almost constant, both shown with blue dashed lines. The middle section of the density profile, shown as shaded circles, is lower than expected. There are two reasons for this method to be less reliable in this area. Firstly, the shock is overlapping with the plasma wave and the width of the shock is longer than the length of a plasma bubble. Therefore, taking the Fourier transform in this area gives a wavelength longer than it should be, and thus the data points in the greyed-out region are lower than expected. Secondly, since the plasma wave ROI is a few hundred micrometres long and contains over 10 plasma bubbles, the calculated plasma wavelength or density is an averaged value instead of the value at the exact probe position. A scale bar showing the length of the ROI is attached at the bottom right-hand corner. Therefore, the peak at 0 mm near the shock is less profound than expected, as it has been averaged with lower densities. At a long distance from the shock, however, the supersonic density profile is nearly constant and averaging over distance still results in a density plateau. To better resolve the density near the shock, methods that do not average over distance could be helpful, such as performing a windowed Fourier transform, a continuous wavelet transform or even a nested object detector.

Another interesting observation is the position jitter of the plasma wave. With the object detector, the vertical position of the plasma wave can be accurately determined, leading to an estimation of the jitter of the driver. The vertical
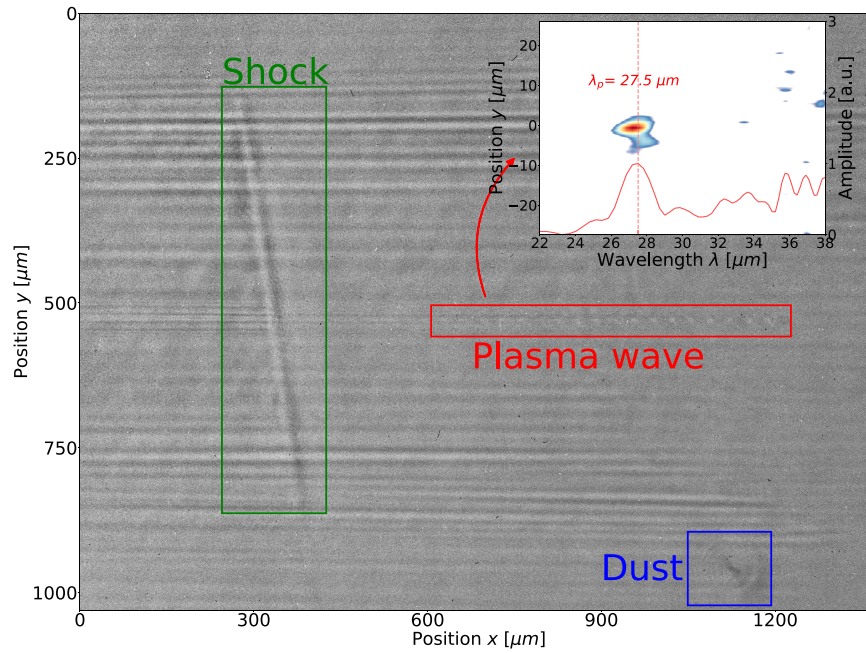
Figure 2. An example: the plasma wave, the shock and the diffraction pattern caused by dust are found by the object detector and located with bounding boxes. More shadowgrams with different shock positions, without shocks, with multiple dust patterns and with overlapping objects are attached in the supplementary material. The subplot on the top right is the Fourier transform of the region within the bounding box of the plasma wave (red). The plasma oscillation wavelength is estimated by integrating along the vertical axis, which peaks at 27.5 μm.
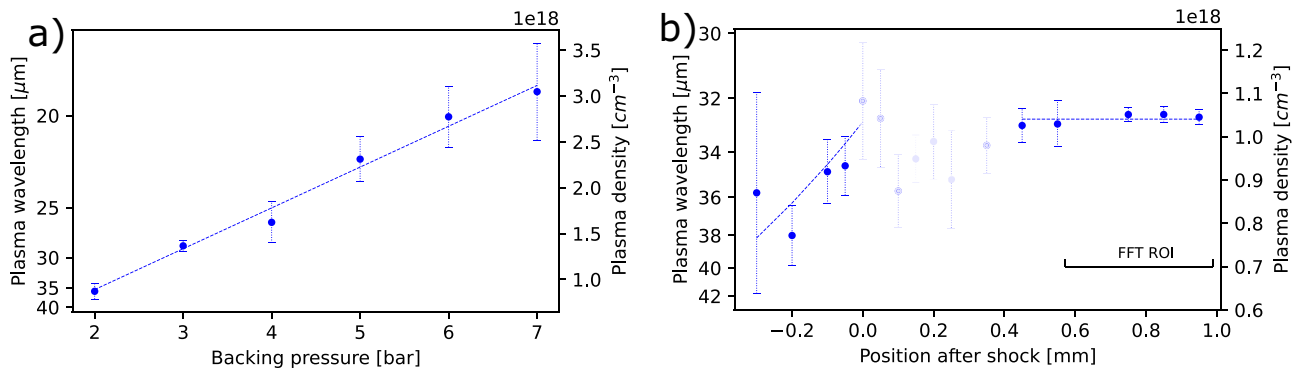

Figure 3. Plasma oscillation wavelengths (left-hand vertical axis) and plasma density (right-hand vertical axis) calculated from the Fourier transform results within the ROI defined by the object detector. (a) The backing pressure of the gas target is scanned from 1 to 6 bar. (b) The probe is moved from the upstream end to the centre of the gas target, and 0 mm is where the first plasma bubble of the plasma wave is at the density shock front. As mentioned in the main text, the region where the ROI includes the shock produces unreliable results and is thus greyed out.

position of the plasma waves from all shots is plotted in Figure 4(a), measured from the centre of the shadowgram, where a negative value means the plasma wave is on the top half of the shadowgram.

The plotted data were taken over approximately 5 hours in one experiment day, and are part of the inference set of 1000 shots mentioned in Table 1. The plasma wave's vertical positions in the camera's field of view of the first and last shot are manually compared to confirm that the positive drift is not artificial. Objects not perfectly centred in the bounding box account for an insignificant offset compared with the vertical jitter from shot to shot.

Note that the parameter scan performed during the experiment only affects the horizontal position of the plasma waves, and thus there is no intentional change in the vertical position of the plasma waves. A slight slope is observed in Figure 4(a), implying a linear drift of the plasma wave vertically over the day. The focusing optic is 6 m away from the gas jet, and the position angle is calculated accordingly. The shot-to-shot fluctuation is calculated from the difference in angle between two consecutive shots. As is shown in Figure 4(b), the majority of the shot-to-shot jitter falls below 4 μrad. This is in line with a separate measurement of the jitter in the laser driver, which is mostly within 3 μrad
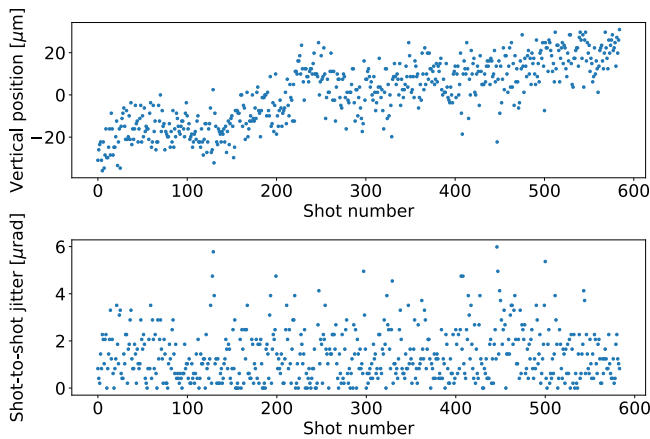
Figure 4. (a) Vertical position of the plasma wave moves over a day. (b) Jitter between every two consecutive shots, calibrated into a solid angle.

and, accordingly, the dominant source of the fluctuation in the plasma wave appears to originate from the pointing fluctuations of the main laser beam.

To summarize, the analysis in this application is enabled by the object detector, which tracks the position change of the plasma wave due to the parameter scan or even the beam jitter itself. While it is possible to select a fixed ROI for the Fourier transform without knowing the exact position of the plasma wave, such as the whole shadowgram, the estimated wavelength would not be reliable. This is due to the fact that such a maximum ROI includes too much unnecessary information, for example, the shock, the background noise or the tail of the plasma wave without observable oscillating structures. On the other hand, the object detector locates the obvious plasma wave structures of the first few bubbles, making it possible to exclude the interference of irrelevant information during the Fourier transform. Therefore, the use of the object detector effectively increases the signal-to-noise ratio during the Fourier transform calculation. Note that the shocks and diffraction patterns caused by dust are also detected on the shadowgrams, as shown in Figure 2. While these objects are not analysed in this study, they could find potential use, such as to relate the position of the shock to the position of the injection point in the accelerator[30].

3.2. Electron energy spectrometer

The second application regards the electron beams from a laser-wakefield accelerator. The energy spectra of the produced electron beams are measured by a magnetic electron spectrometer. The magnetic spectrometer consists of an 80 cm long, 0.85 T permanent magnet with a 4 cm gap. The electrons are deflected as they pass through the magnetic spectrometer, and intersect with the detector plane at different positions. The radius of the trajectory is determined by the energy of the electron, and thus the magnetic spectrometer is calibrated by particle tracking. Peaks on

the electron energy spectra are identified and the associated charge numbers are calculated, not only for mono-energetic beams but also for multi-energy and broadband beams. A calibrated tritium capsule is used as a constant absolute light source in order to calibrate the detected charge. The signal of the tritium capsule is also detected and labelled on the images, from which the charge value is calculated from the pixel intensity[29].

3.2.1. Labelling and training

The training and validation dataset consists of 50 images of electron energy spectra taken from various experiment days. The two labelled classes on the energy spectra are the peaks of the electron beam and the tritium capsule. The dataset is expanded to 82 images using rotational augmentation.

The model is transfer-learned from the YOLO5s.pt model and then fine-tuned with a lower learning rate. Both the training and the fine-tuning process take approximately 10 minutes on a Tesla T4 GPU.

3.2.2. Results

The trained model has an mAP of 0.904 for an IoU threshold of 0.5. After training, the algorithm can detect peaks in the electron energy spectra, as well as the location of the tritium capsule for charge calibration. The charge number of each peak is calculated and annotated alongside; see Figure 5. A threshold confidence of 20% is applied when drawing the bounding boxes.

Figure 5 presents six electron energy spectra with various shapes, positions and numbers of peaks. The spectra were taken on different experiment days. The peaks on the energy spectra are detected and labelled with the charge number, and the tritium capsule is also detected and highlighted in the small bounding box at the bottom of each image. The charge number in pico-coulomb is calibrated by the tritium capsule. Although it might be possible to estimate the charge of the peaks using pre-defined methods, applying an object detector has several advantages. Firstly, it allows for defining the region of a peak even if the peak is in an irregular shape. For instance, in Figures 5(b) and 5(f), it could be difficult to determine a peak using FWHM or other predetermined definitions. While traditional methods tend to implement human judgement via such definitions, object detection algorithms aim to imitate human decisions directly. Namely, the algorithm draws the bounding box the same way a human would draw it based on his/her knowledge and experience. Another advantage is that the object detector enables accurate recognition of multiple peaks, such as in Figures 5(d) and 5(e). While traditional methods may be able to select an area of interest and automize the charge estimation for electron energy spectra taken at similar energy featuring only a single peak, they cannot handle spectra with multiple peaks at changing positions. On the other hand, object detection algorithms can label spectra with
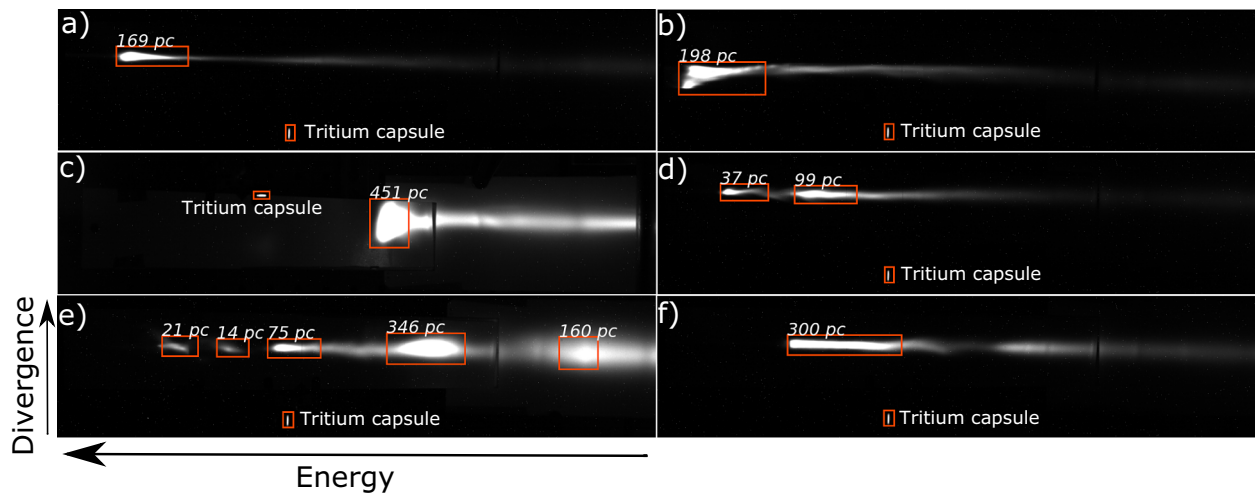
Figure 5. Labelled peaks with charge number on electron energy spectra. The charge numbers are calculated from the integral within each bounding box.

multiple peaks without knowing the number of peaks in advance due to the intrinsic nature of the algorithm. Lastly, live information on the electron charge would benefit the experimental logistics, giving extra feedback when tuning the laser–plasma parameters during experiments.

In addition to the mAP, the error in this method is further evaluated by comparing the estimated charge values of the predicted peaks with those of the manually labelled ground-truth peaks. The average prediction error is 9.6% for all the peaks shown in Figure 5. The charge value within the ground-truth boxes is calculated and labelled on the energy spectra in Figure 2 in the supplementary material. The charge value within the ground-truth boxes is calculated and labelled accordingly.

3.3. High-power laser damage on optics

The peak power of lasers has been ramping up since the invention of chirped-pulse amplification technology[31], entering the petawatt regime in several facilities world-wide[1,3,4,32–36]. While optics are carefully chosen for their high damage threshold, the large size of PW laser optics makes optical damage a main cost driver for operating high-power laser systems. In order to minimize damage propagation along the beam path, it is therefore crucial to detect the first occurrence of a damage spot on any optic and to use such an event to trigger a laser shutdown. In this section, we present the detection of laser damage on an optic in the laser chain (wedge) by imaging the stray light off the main compressor grating and analysing the imaging results using an object detector. In this proof-of-principle setup, the damages occurred on the wedges in the amplification chain; cameras looking at the compressor grating at the end of the chain saw diffraction patterns that can be recognized by an object detector. Of course, the same imaging/object detection algorithm could be applied for directly imaging the diffuse reflection off any laser optic.

3.3.1. Labelling and training
The training and validation dataset consists of 50 images of the grating surface taken from various experiment days. The only labelled class on the images is the damaged spot. The dataset is expanded to 114 images using augmentation in the image-level brightness and the exposure.

The model is transfer-learned from the YOLO5s.pt model and then fine-tuned with a lower learning rate. The training takes approximately 50 minutes and the fine-tuning takes approximately 30 minutes on a Tesla T4 GPU.

3.3.2. Results
The trained model has an mAP of 0.995 for an IoU threshold of 0.5. The model is applied to two inference sets. The first inference set has 1000 images when there is no optical damage in the beam path, and the second inference set has 1000 images when there is an optical damage in the beam path. The images were taken from different experiment days. The model detects no signal of damage in any of the images in the first set, while it detects the diffraction patterns from the optical damage in all 1000 images in the second set. The results prove the good consistency of the object detector, since it neither misses any damaged optics nor gives false labels. Figure 6 presents two examples of the prediction results in the inference sets, where Figure 6(a) is from the first set and Figure 6(b) is from the second set. A threshold confidence of 40% is applied when drawing the bounding boxes. This application demonstrates the potential to use object detection in any high-power laser system for immediate warnings on crucial optical elements, providing timely protection to the rest of the optics in the system. It has to be pointed out that this application, unlike the previous two, can also be treated as a classification problem to judge if there is an optical damage. A recent work by Tudor[37] finds an abnormal laser beam profile using a CNN. Amorin *et al.*[38] classified the microscopic damages on National
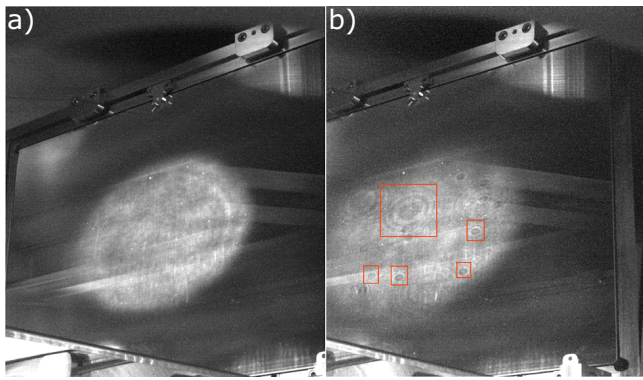
Figure 6. Detected interference pattern on a grating surface, originated from damages of previous optics in the amplification beam path: (a) is an image of the grating surface without damaged optics in the beam path, while (b) is an image of a grating surface with damaged optics in the beam path. The bounding boxes are drawn around the detected damage spots.

Ignition Facility (NIF) optics using a CNN connected by an ensemble of decision trees, which combines the feature extraction ability of CNNs and the decision-making strength of decision trees. It is also possible to trace where the damage occurs in the beam path. Li *et al.*[39] detected the laser damage using a CNN with the input being the diffraction ring feature itself instead of the image. Likewise, our future work is to narrow down the damage location to a range of longitudinal positions based on the size and periodicity of the ring structures, while utilizing the fast inference speed of this algorithm for real-time multi-damage positioning.

4. Summary and outlook

In this paper, we have provided three examples of how advanced computer vision techniques can be applied to assist online and offline experimental analysis in high-power laser facilities. We have shown satisfying predictions by fine-tuning a pre-trained network for general object detection tasks using only approximately 50 hand-labelled images. The learned model has been examined using not only the test dataset split from the input data, but also a separate inference set of 1000 images with various laser–plasma parameters. The model training is performed using accessible computational resources in GPU hours or below, while the inference time on an unseen image with the trained models takes only tens of milliseconds.

The main benefit of object detection is the possibility of real-time data analysis. For a human it is not possible to process all the important information of an experiment running at a 1-Hz shot rate. The object detector allows some degree of real-time, in-depth data analysis and visualization during an experimental campaign, which goes beyond that of what a human operator could achieve.

The presented methodology is widely adaptable and easy to implement. From a practical perspective, it does not require much expertise in machine learning to construct an architecture, and users only have to label a few dozen diagnostic images. For readers with slightly different laser parameters and diagnostics, we have shared our code in a GitHub repository linked in the supplementary material. The method can also be applied to kHz laser systems with little loss in the repetition rate. The algorithm can operate at approaching 100 Hz with lower prediction accuracy, while experimentalists usually have to average over approximately 10 laser shots to reduce fluctuations. Therefore, the presented methodology is not much slower than the 'effective' repetition rate of a kHz laser.

With the superiority in time and consistency, the prediction results are reliable and robust compared with manual human recognition. Thus, we anticipate that the concept of object detection will find wide applications in more image-related measurements and diagnostics in high-power laser experiments.

Acknowledgements

Supplementary Materials

To view supplementary material for this article, please visit http://doi.org/10.1017/hpl.2023.1.

References

1. J. H. Sung, H. W. Lee, J. Y. Yoo, J. W. Yoon, C. W. Lee, J. M. Yang, Y. J. Son, Y. H. Jang, S. K. Lee, and C. H. Nam, Opt. Lett. **42**, 2058 (2017).
2. https://www.cala-laser.de/.
3. J. Nees, A. Maksimchuk, G. Kalinchenko, B. Hou, Y. Ma, P. Campbell, A. McKelvey, L. Willingale, I. Jovanovic, C. Kuranz, A. Thomas, and K. Krushelnick, in *2020 Conference on Lasers and Electro-Optics (CLEO)* (IEEE, 2020), p. 1.
4. Z. Zhang, F. Wu, J. Hu, X. Yang, J. Gui, P. Ji, X. Liu, C. Wang, Y. Liu, X. Lu, Y. Xu, Y. Leng, R. Li, and Z. Xu, High Power Laser Sci. Eng. **8**, e4 (2020).
5. S. Borneis, T. Laštovička, M. Sokol, T.-M. Jeong, F. Condamine, O. Renner, V. Tikhonchuk, H. Bohlin, A. Fajstavr, J.-C. Hernandez, N. Jourdain, D. Kumar, D. Modřanský, A. Pokorný, A. Wolf, S. Zhai, G. Korn, and S. Weber, High Power Laser Sci. Eng. **9**, e30 (2021).
6. Z.-H. He, B. Hou, V. Lebailly, J. A. Nees, K. Krushelnick, and A. G. R. Thomas, Nat. Commun. **6**, 7156 (2015).
7. M. J. V. Streeter, S. J. D. Dann, J. D. E. Scott, C. D. Baird, C. D. Murphy, S. Eardley, R. A. Smith, S. Rozario, J.-N. Gruse, S. P. D. Mangles, Z. Najmudin, S. Tata, M. Krishnamurthy, S. V. Rahul, D. Hazra, P. Pourmoussavi, J. Osterhoff, J. Hah,

N. Bourgeois, C. Thornton, C. D. Gregory, C. J. Hooker, O. Chekhlov, S. J. Hawkes, B. Parry, V. A. Marshall, Y. Tang, E. Springate, P. P. Rajeev, A. G. R. Thomas, and D. R. Symes, Appl. Phys. Lett. **112**, 244101 (2018).

8. M. Noaman-ul Haq, T. Sokollik, H. Ahmed, J. Braenzel, L. Ehrentraut, M. Mirzaie, L.-L. Yu, Z. M. Sheng, L. M. Chen, M. Schnürer, and J. Zhang, Nucl. Instrum. Methods Phys. Res. Sect. A **883**, 191 (2018).

9. J. Lin, Y. Ma, R. Schwartz, D. Woodbury, J. A. Nees, M. Mathis, A. G. R. Thomas, K. Krushelnick, and H. Milchberg, Opt. Express **27**, 10912 (2019).

10. J. R. Smith, C. Orban, J. T. Morrison, K. M. George, G. K. Ngirmang, E. A. Chowdhury, and W. M. Roquemore, New J. Phys. **22**, 103067 (2020).

11. A. Englesbe, J. Lin, J. Nees, A. Lucero, K. Krushelnick, and A. Schmitt-Sody, Appl. Opt. **60**, G113 (2021).

12. K. D. Humbird, J. L. Peterson, and R. G. McClarren, IEEE Trans. Neural Networks Learn. Syst. **30**, 1286 (2018).

13. A. Gonoskov, E. Wallin, A. Polovinkin, and I. Meyerov, Sci. Rep. **9**, 7043 (2019).

14. K. D. Humbird, J. L. Peterson, B. K. Spears, and R. G. McClarren, IEEE Trans. Plasma Sci. **48**, 61 (2019).

15. A. Hsu, B. Cheng, and P. A. Bradley, Phys. Plasmas **27**, 012703 (2020).

16. R. J. Shalloo, S. J. D. Dann, J.-N. Gruse, C. I. D. Underwood, A. F. Antoine, C. Arran, M. Backhouse, C. D. Baird, M. D. Balcazar, N. Bourgeois, J. A. Cardarelli, P. Hatfield, J. Kang, K. Krushelnick, S. P. D. Mangles, C. D. Murphy, N. Lu, J. Osterhoff, K. Põder, P. P. Rajeev, C. P. Ridgers, S. Rozario, M. P. Selwood, A. J. Shahani, D. R. Symes, A. G. R. Thomas, C. Thornton, Z. Najmudin, and M. J. V. Streeter, Nat. Commun. **11**, 6355 (2020).

17. J. Lin, Q. Qian, J. Murphy, A. Hsu, A. Hero, Y. Ma, A. G. R. Thomas, and K. Krushelnick, Phys. Plasmas **28**, 083102 (2021).

18. A. Sävert, S. P. D. Mangles, M. Schnell, E. Siminos, J. M. Cole, M. Leier, M. Reuter, M. B. Schwab, M. Möller, K. Poder, O. Jäckel, G. G. Paulus, C. Spielmann, S. Skupin, Z. Najmudin, and M. C. Kaluza, Phys. Rev. Lett. **115**, 055002 (2015).

19. M. F. Gilljohann, H. Ding, A. Döpp, J. Götzfried, S. Schindler, G. Schilling, S. Corde, A. Debus, T. Heinemann, B. Hidding, S. M. Hooker, A. Irman, O. Kononenko, T. Kurz, A. M. de la Ossa, U. Schramm, and S. Karsch, Phys. Rev. X **9**, 011046 (2019).

20. H. Ding, A. Döpp, M. Gilljohann, J. Götzfried, S. Schindler, L. Wildgruber, G. Cheung, S. M. Hooker, and S. Karsch, Phys. Rev. E **101**, 023209 (2020).

21. M. C. Downer, R. Zgadzaj, A. Debus, U. Schramm, and M. C. Kaluza, Rev. Mod. Phys. **90**, 035002 (2018).

22. J. Gu, Z. Wang, J. Kuen, L. Ma, A. Shahroudy, B. Shuai, T. Liu, X. Wang, G. Wang, J. Cai, and T. Chen, Pattern Recogn. **77**, 354 (2018).

23. R. Yamashita, M. Nishio, R. K. G. Do, and K. Togashi, Insights imaging **9**, 611 (2018).

24. Z. Li, F. Liu, W. Yang, S. Peng, and J. Zhou, IEEE Trans. Neural Netw. Learn. Syst. **33**, 6999 (2021).

25. R. Girshick, J. Donahue, T. Darrell, and J. Malik, IEEE Trans. Pattern Anal. Machine Intell. **38**, 142 (2015).

26. J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, 2016), p. 779.

27. https://github.com/ultralytics/yolov5.

28. B. Hidding, A. Beaton, L. Boulton, S. Corde, A. Doepp, F. A. Habib, T. Heinemann, A. Irman, S. Karsch, G. Kirwan, A. Knetsch, G. G. Manahan, A. M. de la Ossa, A. Nutter, P. Scherkl, U. Schramm, and D. Ullmann, Appl. Sci. **9**, 2626 (2019).

29. T. Kurz, J. P. Couperus, J. M. Krämer, H. Ding, S. Kuschel, A. Köhler, O. Zarini, D. Hollatz, D. Schinkel, R. D'Arcy, J.-P. Schwinkendorf, J. Osterhoff, A. Irman, U. Schramm, and S. Karsch, Rev. Sci. Instrum. **89**, 093303 (2018).

30. J. Wenz, A. Döpp, K. Khrennikov, S. Schindler, M. F. Gilljohann, H. Ding, J. Götzfried, A. Buck, J. Xu, M. Heigoldt, W. Helml, L. Veisz, and S. Karsch, Nat. Photonics **13**, 263 (2019).

31. D. Strickland and G. Mourou, Opt. Commun. **55**, 447 (1985).

32. C. N. Danson, P. A. Brummitt, R. J. Clarke, J. L. Collier, B. Fell, A. J. Frackiewicz, S. Hancock, S. Hawkes, C. Hernandez-Gomez, P. Holligan, M. H. R. Hutchinson, A. Kidd, W. J. Lester, I. O. Musgrave, D. Neely, D. R. Neville, P. A. Norreys, D. A. Pepler, C. J. Reason, W. Shaikh, T. B. Winstone, R. W. W. Wyatt, and B. E. Wyborn, Nucl. Fusion **44**, S239 (2004).

33. E. W. Gaul, M. Martinez, J. Blakeney, A. Jochmann, M. Ringuette, D. Hammond, T. Borger, R. Escamilla, S. Douglas, W. Henderson, G. Dyer, A. Erlandson, R. Cross, J. Caird, C. Ebbers, and T. Ditmire, Appl. Opt. **49**, 1676 (2010).

34. S. Gales, K. A. Tanaka, D. L. Balabanski, F. Negoita, D. Stutman, O. Tesileanu, C. A. Ur, D. Ursescu, I. Andrei, S. Ataman, M. O. Cernaianu, L. D'Alessi, I. Dancus, B. Diaconescu, N. Djourelov, D. Filipescu, P. Ghenuche, D. G. Ghita, C. Matei, K. Seto, M. Zeng, and N. V. Zamfir, Rep. Prog. Phys. **81**, 094301 (2018).

35. W. Li, Z. Gan, L. Yu, C. Wang, Y. Liu, Z. Guo, L. Xu, M. Xu, Y. Hang, Y. Xu, J. Wang, P. Huang, H. Cao, B. Yao, X. Zhang, L. Chen, Y. Tang, S. Li, X. Liu, S. Li, M. He, D. Yin, X. Liang, Y. Leng, R. Li, and Z. Xu, Opt. Lett. **43**, 5681 (2018).

36. C. N. Danson, C. Haefner, J. Bromage, T. Butcher, J.-C. F. Chanteloup, E. A. Chowdhury, A. Galvanauskas, L. A. Gizzi, J. Hein, and D. I. Hillier High Power Laser Sci. Eng. **7**, e54 (2019).

37. P. Tudor, in *LPA Online Workshop on Control Systems and Machine Learning* (2022).

38. C. Amorin, L. M. Kegelmeyer, and W. P. Kegelmeyer, Stat. Anal. Data Mining **12**, 505 (2019).

39. Z. Li, L. Han, X. Ouyang, P. Zhang, Y. Guo, D. Liu, and J. Zhu, Opt. Express **28**, 10165 (2020).