Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Machine-learning assisted analysis of oil-film interferometry images and application to 3D boundary layers

André Weingärtner

A.Weingartner@soton.ac.uk

University of Southampton

Manuel Aguiar Ferreira

City St George's, University of London

Bharathram Ganapathisubramani

University of Southampton

Research Article

Keywords: wall-shear stress, oil-film interferometry, three-dimensional boundary layer, convolutional neural network, machine learning

Posted Date: November 25th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11639163/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Machine-learning assisted analysis of oil-film interferometry images and application to 3D boundary layers

André Weingärtner[1*], Manuel Aguiar Ferreira[1,2] and Bharathram Ganapathisubramani[1]

[1*]Department of Aeronautical and Astronautical Engineering, University of Southampton, Southampton, SO17 1BJ, United Kingdom.
[2]Department of Engineering, City St George's, University of London, London, EC1V 0HB, United Kingdom.

*Corresponding author(s). E-mail(s): a.weingartner@soton.ac.uk;
Contributing authors: M.Aguiar-Ferreira@soton.ac.uk;
G.Bharath@soton.ac.uk;

Abstract

Although wall-shear stress (WSS) is a central quantity in fluid mechanics, experimental determination remains challenging. Even more so in three-dimensional (3D) flows, where measuring both the direction and magnitude of the WSS vector becomes particularly demanding. An established method to directly measure the WSS under a boundary layer is oil-film interferometry (OFI), but the application in 3D flows is nontrivial, and results are easily corrupted. In this study, we present a new approach to analyze images from discrete OFI measurements that can be applied both in 2D and 3D flows. The result is a quantitative map of the vectorial WSS over a relatively large area. The method combines machine learning with classical approaches, where the directionality of the WSS is determined using a convolutional neural network (CNN) and its magnitude through a novel autocorrelation approach. We show that the CNN is a reliable method to find the direction of the shear stress and the autocorrelation is more robust than the classical Fourier transform method. Finally, we present results from applying the method to both 2D and 3D flow cases. Besides a canonical 2D zero-pressure gradient turbulent boundary layer, results of the 3D boundary layer around a simplified wing-body junction are shown.

Keywords: wall-shear stress, oil-film interferometry, three-dimensional boundary layer, convolutional neural network, machine learning

1 Introduction

The wall-shear stress (WSS), τ_w, characterizes the tangential stress exerted by a fluid on a solid boundary as a consequence of the no-slip condition at the wall. It is a fundamental quantity in fluid mechanics, governing a range of phenomena. The integral of τ_w over a surface determines the skin-friction drag, while in wall-bounded turbulent flows it defines the friction velocity, $u_\tau = \sqrt{\tau_w/\rho}$, which provides the characteristic velocity scale for inner (viscous) scaling of the near-wall region. Wall-shear stress is also closely associated with surface degradation and erosion, making its accurate characterization relevant to a broad range of engineering and geophysical applications.

However, experimentally τ_w is difficult to determine, even in stationary, two-dimensional (2D) flows. Typical approaches include fitting a universal velocity profile to boundary layer data and thus indirectly inferring the value of τ_w [1], direct measurements by means of a floating-element balance [2] or oil-film interferometry (OFI), which will be described in more detail. A comprehensive overview of WSS measurement techniques is given by Vinuesa and Örlü [3]. When studying three-dimensional (3D) flows, where the WSS is generally not aligned with the flow in the free stream, measuring τ_w becomes even more challenging. OFI has been applied in 3D flows previously, which will be briefly mentioned at the end of the following section.

1.1 Oil-film interferometry

The OFI method, as first proposed by Tanner and Blows [4], relates the deformation of a thin oil film to the local WSS. A droplet of oil is placed on a smooth surface and, as it is swept away by the flow, it leaves a thin film of oil behind. This thinning process continues and at some point, the height of the oil film reaches a similar order of magnitude as the wavelength of the incoming light, and interferometric effects become relevant. Figure 1(a) shows a typical setup of an OFI experiment. A camera images oil films that are placed on a transparent surface from below at an oblique angle α. Monochromatic light is generated through a sodium lamp with a wavelength of $\lambda_{\text{light}} = 589.3\,\text{nm}$. In the presented setup, a large piece of white paper is used as a light reflector, generating an evenly lit background. Figure 1(b) shows a zoomed-in view of a thin oil film on the surface. The incoming light is reflected twice, at the glass-oil interface and the oil-air interface. Consequently, these two reflected waves feature a path difference, which can either result in constructive or destructive interference, depending on the relative phase angle. Note that in some experimental setups, the camera and light source are placed above rather than below the oil. This has the advantage that the substrate is not required to be transparent anymore, but still reflective. However, optical access might be more challenging. The physical phenomena and equations are similarly valid, as the additional air-glass refraction in the setup only leads to a translation of the fringes [5].

Due to the local variation in height of the oil film, interference patterns emerge across the film, called Fitzeau fringes. These interference lines indicate regions of constant oil height. An example of a typical interference pattern is shown in Figure 2. The yellow lines are all circular arcs with equal diameter, placed by eye. The fact that the diameter is constant throughout the droplet is relevant for the later processing

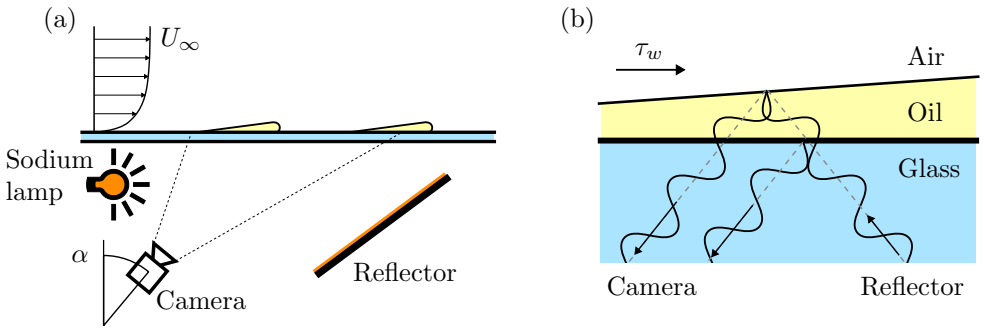

Fig. 1 Schematic of a typical experimental setup for OFI. The right-hand side shows a zoom into the oil film, detailing the emergence of interference patterns.

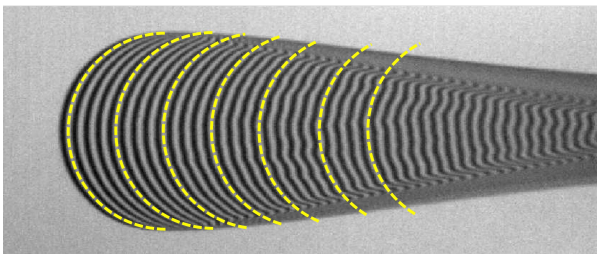

Fig. 2 Fitzeau-interference pattern as typically observed in an OFI experiment. The yellow lines are circular arcs of equal diameter.

methods. Based on the optical parameters, namely light wavelength λ_{light}, refractive indices of oil and air, n_{oil} and n_{air}, and the oblique camera viewing angle α, the difference in oil height between two fringes can be calculated as

$$\Delta h = \frac{\lambda_{\text{light}}}{2\sqrt{n_{\text{oil}}^2 - n_{\text{air}}^2 \sin^2 \alpha}}. \tag{1}$$

Now, we assume that the oil is only deformed by the shear stress acting on its surface, which is justified with the film height being around one viscous unit [5]. The equation describing the height h of a two-dimensional oil film under shear is

$$\frac{\partial h}{\partial t} + \frac{\partial}{\partial x}\left(\frac{\tau_{w,x} h^2}{2\mu}\right) + \frac{\partial}{\partial z}\left(\frac{\tau_{w,z} h^2}{2\mu}\right) = 0, \tag{2}$$

where x and z are the streamwise and spanwise direction, respectively, τ_w is the WSS (in the two directions) and μ the dynamic viscosity of the oil. In one dimension and assuming a constant τ_w, this can be written as

$$\frac{\partial h}{\partial t} + \frac{\tau h}{\mu}\frac{\partial h}{\partial x} = 0. \tag{3}$$

Using this equation, the WSS can be expressed in terms of the mean peak distance of the interferometric pattern, or the fringe spacing, λ_f,

$$\tau_w = \frac{\mu}{\Delta h} \frac{\partial \lambda_f}{\partial t}. \tag{4}$$

Now, the only challenge is to determine λ_f from the interference image. Usually, a region in the center of the droplet is selected manually, and a Fourier transform is employed to compute the mean fringe wavelength. Note that another, but found to be less robust method is to track individual fringes over time and determine their velocity, for example by building an x-t diagram [3]. More elaborate approaches, such as a method based on empirical mode decomposition and Hilbert transforms was proposed by Chauhan et al. [6]. They demonstrated that for particularly noisy images, this method can be advantageous. Regardless of the applied method, λ_f changes linearly in time for a constant τ_w so that a straight line can be fitted to the data to determine the constant value of $\mathrm{d}\lambda_f/\mathrm{d}t$, from which then τ_w can be calculated.

Gluzman et al. [7] developed a simplified approach to perform OFI measurements along a curved surface. Using photogrammetry, they were able to map the surface of a Gaussian bump from a set of calibration images and project the fringe images onto the surface. In 3D flows, the application of OFI becomes more challenging, as the direction of the WSS has to be assessed separately. Ruedi et al. [8] performed OFI measurements in a 3D turbulent boundary layer wind tunnel, comparing it to other techniques. By manual visual analysis, they determined the direction of the oil flow and then tracked fringes in that direction, calculating their velocity and thus the shear stress.

1.2 Convolutional neural networks

Convolutional neural networks (CNNs) are a well-established class of machine-learning models for image analysis. CNNs exploit the spatial structure of images through convolutional layers, in which learnable kernels are applied locally across the input to extract features. Successive convolutional layers can learn increasingly complex representations, with early layers typically capturing low-level features such as edges and textures and deeper layers capturing more abstract features. The use of local receptive fields and shared convolutional kernels also substantially reduces the number of parameters relative to fully connected architectures.

A key development in CNN architectures was the introduction of residual neural networks (ResNets) by He et al.[9]. ResNets address the difficulties associated with training very deep networks through the use of residual, or skip, connections. Rather than requiring a stack of layers to directly learn a desired mapping, these connections allow the layers to learn a residual mapping with respect to their input. This facilitates the optimization of substantially deeper networks while mitigating the degradation of performance that can occur as network depth increases. Here, we make use of a ResNet architecture for image feature extraction and a subsequent regression task. The regression is realized through a fully-connected neural network, or multilayer perceptron (MLP), where the layer width is gradually reduced up to the outputs.

Vinuesa et al. [10] give an overview of the capabilities of machine learning, specifically in the field of experimental fluid mechanics. They identify a significant potential in improving measurement techniques, enhancing experimental design and modeling. More relevant for the machine learning aspects of the present study, Schulte-Sasse et al. [11] presented a CNN-based method to automatically determine the direction of the wall streamlines from oil-flow visualizations. In their approach, they divide the complete oil-flow image into windows and separately infer the local WSS angle with a neural network, returning continuous wall streamlines. More recently, Mehdizadeh Youshanlouei et al.[12] investigated the use of machine-learning methods to quantitatively analyze OFI images and determine the mean fringe spacing. They present a machine-learning workflow to determine the WSS from the images. However, the reported approach of finding the slope $\mathrm{d}\lambda_f/\mathrm{d}t$ achieved an accuracy of approximately $\pm 5\%$ relative to the reference measurements, which does not represent a clear improvement over established approaches. This is particularly striking given that the uncertainty of carefully conducted OFI measurements is often cited as $\pm 1\%$ [3, 5]. Both previous papers employ classical Visual Geometry Group (VGG) neural networks. Rooted in the field of meteorology, Coney et al. [13] developed a related method using a more modern ResNet-type approach. They were able to identify location, direction and amplitude of trapped lee waves from vertical velocity data.

In the present study, we introduce a method to quantitatively determine the WSS under 3D turbulent boundary layers. It is based on the fundamental OFI measurement approach and combines machine-learning approaches with classical data analysis. In a 3D flow, the WSS features a direction and a magnitude. Hence, from here on it will be denoted with a bold symbol indicating a vector quantity ($\boldsymbol{\tau_w}$).

The paper is structured as follows. In Section 2, we introduce the developed methods of 3D WSS determination from OFI images. The corresponding subsections describe first the identification of droplets in a larger image, then finding their direction using a machine-learning algorithm and finally the fringe wavelength in that direction using an autocorrelation approach. Section 3 shows the application of the methods to experimental wind-tunnel data. Besides data from a canonical 2D boundary layer, we present continuous WSS fields around a simplified wing-body junction. Finally, conclusions are drawn in Section 4.

2 Methods

As in the standard OFI technique, in a 3D flow oil droplets placed on a substrate are gradually thinned by the flow. The underlying assumption here is that the one-dimensional method of Tanner and Blows [4] continues to be locally valid. While Equation 3 can be expressed in an arbitrary direction, deriving it from Equation 2 assumes a constant WSS in the local flow direction, which is invalid by default in 3D flows. However, we assume that the spatial gradients of $\boldsymbol{\tau_w}$ are sufficiently small across the measurement area (i.e. over one oil droplet), so that the theory remains valid. Although not shown here, when plotting the interference data in an x-t diagram (see e.g. [6]), the fringes developed linearly in space for most droplets, indicating that the assumption holds. In some extreme cases, where the droplet was placed in a region of

large shear gradients, this was not valid, but a straight line still represented a good first-order approximation and is expected to adequately represent the average WSS across the window. Another assumption is that the direction of the oil flow is constant in the analyzed window and that it aligns with the direction of the local WSS, which will be discussed later.

The implemented algorithm follows three steps, which are described separately in the following subsections. First, we identify the droplets in a larger image, then determine the direction and finally the magnitude of the WSS $\boldsymbol{\tau_w}$ for each droplet separately.

2.1 Droplet identification

The first step in the data processing pipeline is to automatically identify oil droplets in the dewarped image, rather than having to select them manually. For this, we implemented a machine-learning based "You Only Look Once" (YOLO) object detection algorithm, first introduced by Redmon et al. [14]. More specifically, we employed the more modern variant YOLOv8 [15], which made it possible to define key points (such as oil edges) in the training images, improving the overall performance. Furthermore, the SAHI (Slicing Aided Hyper Inference) algorithm was used, which scans the total image with interrogation windows and separately performs the object detection. For an overview of different YOLO architectures, see the review by Terven et al. [16].

The detection was performed for a single image, typically around the middle of the measurement time. Here, interference patterns are usually clearly visible, and the wavelength is still not too large, so that the fringes can be captured by an interrogation window. This approach provided a reliable and relatively fast detection of the droplets, despite large image sizes. A result of the detection will be shown later in the results section, see Figure 12. Note that while the algorithm generally works very reliably, it did not always detect exactly the center of the droplets. However, the developed method for finding the fringe spacing λ_f is agnostic to this and makes use of the whole fringe pattern (Section 2.3). Now, a box around each detected droplet was defined and analysis was performed for each image stack sequentially. The provided software package includes an interactive reviewer, where the detection threshold can be adjusted dynamically, and detections can be added or removed manually [17].

2.2 Wall-shear stress direction

To determine the direction of the local WSS for each measurement (i.e., each droplet), we employed a neural network, which was found to consistently predict the angle in all investigated cases. The relevant details, as well as some results, will be presented in the following subsections.

2.2.1 Training data

For the training data, droplet images from various previous experiments featuring 2D flows were used, where the direction of the WSS is known to be in the direction of the free stream. The images were then digitally rotated by random angles θ from 0° to 359.9° and the value of θ was used as a label for training. To avoid artifacts from the

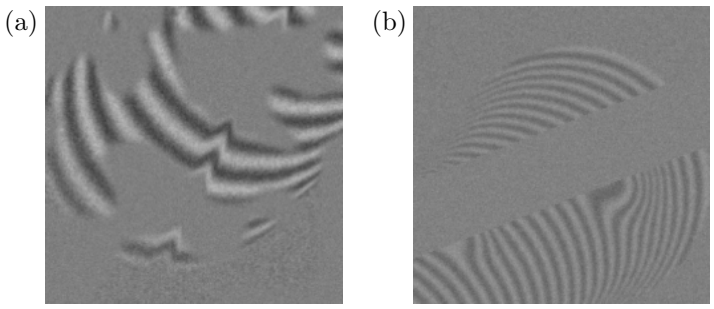

Fig. 3 Sample training images. (a): real and (b): artificial interference pattern, showcasing the applied augmentations.

digital rotation affecting the CNN performance, several additional augmentations were implemented, such as Gaussian noise, blurring, intensity degradation, resizing, deletion of random patches, and local phase dislocations. Square images of size 224×224 were found to be most practical due to camera resolution and common CNN architectures. However, in order to make the model as general as possible, the network was also trained on smaller images, where zero-padding was employed to bring them to the required input size.

In addition to the real experimental training images, synthetic droplet interference patterns were generated with the goal to generalize the model better. We observed that the circular interference patterns within one droplet feature a clear pattern. In all cases, the circular arcs are of equal diameter and merely shifted, as highlighted in Figure 2. Note that this diameter is defined by the initial size of the droplet and hence does not change with time either. The same pattern was consistently observed in the literature, e.g. [3]. This implies that the shape of the oil film assumes a specific geometrical shape in that region, which is a linearly slanted cylinder. Hence, the 3D shape could be recreated digitally and interference patterns were calculated applying the thin-film interference equation. Artificial noise could be added directly to the 3D shape, so that more realistic disturbances in the interference patterns were generated and the resulting patterns were tapered at various angles to better resemble realistic droplets. This process is described in more detail in the appendix (Section A) and the corresponding MATLAB code is provided [17].

Figure 3 shows examples of a real (a) and an artificial (b) training image. The final data set contained a total of almost 900,000 8-bit images, where 5% were reserved for validation and around 3% for testing. Before training, several preprocessing steps were performed. First, the time-mean of each image series was subtracted to avoid the effects of an irregular background. Then, each image was normalized by dividing it by its maximum and then subtracting its mean. The maximum normalization is beneficial since the intensity of each individual image is irrelevant to the directionality encoded in it.

2.2.2 Loss function

For the given task, the angle θ is represented through its sine and cosine and the network performs a circular regression. This avoids the discontinuity at $0°/360°$ and

places the angle onto the unit circle, as well as bounding the target values to the interval $[-1, 1]$,

$$\boldsymbol{y} = \begin{pmatrix} \sin\theta \\ \cos\theta \end{pmatrix}. \tag{5}$$

The neural network predicts a two-dimensional vector $\mathbf{z} = (z_{\sin}, z_{\cos})$, which is normalized to unit length to obtain the predicted representation $\hat{\boldsymbol{y}}$. The predicted angle $\hat{\theta}$ can then be recovered using

$$\hat{\theta} = \arctan2(\hat{y}_{\sin}, \hat{y}_{\cos}). \tag{6}$$

The angular error was quantified using the geodesic distance on the unit circle, analogous to the use of geodesic losses for neural-network-based orientation regression [18, 19]. The loss for each sample is

$$L_{\theta,i} = \arctan2(|\hat{\boldsymbol{y}} \times \boldsymbol{y}|, \hat{\boldsymbol{y}} \cdot \boldsymbol{y}). \tag{7}$$

Here, the $\times$ operator identifies the two-dimensional cross product of the two vectors, which can be achieved by adding a third dimension of zeros and performing a 3D cross product. Note that, since both vectors lie on the unit circle, this expression is equal to the smallest angular difference between $\hat{\theta}$ and θ.

The loss used for the optimization is the mean geodesic loss over a batch of N samples,

$$L_\theta = \frac{1}{N} \sum_{i=1}^{N} L_{\theta,i}, \tag{8}$$

with the loss expressed in radians.

The loss expression in Equation 7 was found to perform better than the commonly used cosine loss $L_{\cos} = 1 - \cos(\hat{\theta} - \theta)$, which behaves quadratically when approaching zero. The employed geodesic loss, on the other hand, behaves linearly with respect to the magnitude of the angular difference around zero. Consequently, it provides larger gradients for small errors, which was found to lead to a better performance given the same network architecture and training procedure. The first derivative of the loss function is discontinuous at zero, which, however, did not affect the training substantially.

2.2.3 Network architecture and training

The proposed model employs a compact residual convolutional neural network for feature extraction followed by a fully connected regression head, as schematically shown in Figure 4. The input is a single-channel 224×224 interference image. An initial convolutional stem comprising a 3×3 strided convolution, batch normalization, ReLU activation and max pooling performs initial feature extraction and spatial downsampling. This is followed by four residual stages with channel widths of 16, 32, 64 and 96, respectively. Each residual block comprises two 3×3 convolutional layers with batch normalization, with a ReLU activation following the first convolution and after the residual addition. A projection shortcut consisting of a 1×1 convolution and batch

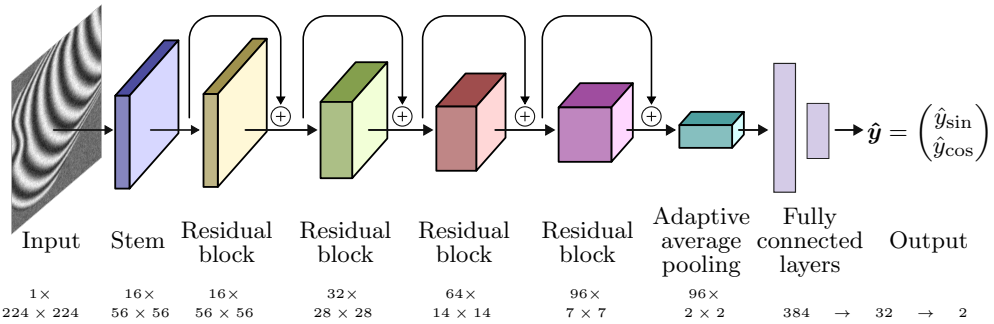

Fig. 4 Schematic of the architecture of the employed neural network.

normalization is used, where the spatial dimensions or number of feature channels change. Following the convolutional feature extractor, 2-D dropout with a probability of 0.15 is applied, followed by adaptive average pooling to 2×2. The resulting $96 \times 2 \times 2$ feature representation is flattened to form a 384-dimensional feature vector. The regression head consists of fully connected layers (MLP) mapping this representation from 384 to 32 and subsequently to the two output values. The network has a total of 234,450 trainable parameters. Note that due to the adaptive average pooling, the input image can have any size. However, we performed the training exclusively for 224×224 images, which is the recommended input size.

Slight variations in the network architecture, such as a larger number of channels, were found to only marginally improve the losses during training, and have even less influence on the result when applied to real data. Hence, the employed model appears to be a good compromise between a relatively small network that is still able to predict the angles suitably well.

The model was trained using the AdamW optimizer with a maximum learning rate of 3×10^{-4} and a weight decay of 10^{-2}. A OneCycle learning-rate schedule was employed over the complete training run, with a batch size of 196 images and a total of 100 training epochs. During training, the model was optimized using the geodesic angular loss described above, with automatic mixed-precision computation employed. After each epoch, the model was evaluated independently on both artificial and real validation datasets. The model featuring the lowest real-data validation loss was retained as the best-performing model.

2.2.4 Sample angle results

As described earlier, OFI measurements require a time series of images, nominally featuring the same angle. Now, each droplet image was passed to the network for inference separately. An example from a real interference pattern in a 3D flow is presented in Figure 5(a), showing one preprocessed image from the time series (here: image #80). Figure 5(b) shows the angle output from the CNN for each image in the series. The predicted angles lie well within $\pm 1°$, which gives confidence in the employed method. Now, the circular median across the angles is calculated, indicated by the red horizontal line. An arrow at this angle of $\theta = 218.25°$ is then drawn on top of the raw

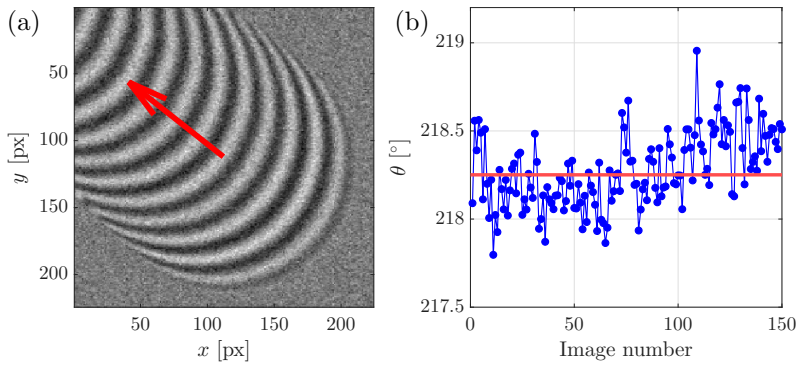

Fig. 5 Result of the CNN for angle prediction when applied to 3D flow data. (a): raw interference fringes on a single image and the identified direction. (b): CNN result for each image in the time series and the circular median value in red.

interference pattern. Note that the angle prediction does not vary significantly across the set of 150 images, although the film height decreases slightly over time, which was generally true for the observed interference patterns.

The supplied software contains the readily trained trained neural network weights including inference code implemented in Python that follows the outlined procedure [17].

2.2.5 Assumptions

The underlying assumption behind this method is that the direction of the oil flow is aligned with the direction of the WSS. This was already investigated by Ruedi et al. [8], who compared OFI to a turnable wall hot wire in a 3D turbulent boundary layer wind tunnel. Through visual judgment, they estimated the direction of the local WSS and found that the difference between the methods was below 2°. However, for this to be generally true, some assumptions must be made which have not been discussed in detail previously.

Once interference effects can be observed, the oil film is expected to be sufficiently thin so that only the WSS is acting on it. However, the initial height of the oil is significantly larger. It can be estimated using the volume formula of a spherical cap,

$$V = \frac{1}{6}\pi h(3a^2 + h^2), \tag{9}$$

with a being the base radius of the cap. We found that a typical droplet has a volume of around 8 μL and a base diameter $2a$ of roughly 8 to 10 mm prior to starting the wind tunnel. In our experiments, the droplets were applied from a syringe with a dispensing tip and allowed to drop naturally once they reached a certain size, which is expected to be mainly a function of oil density and surface tension. Using these values, the initial height of the oil droplet can be calculated to be around $0.2 - 0.3$ mm. In a 3D boundary layer, the flow direction is known to vary near the wall, so that this height might be relevant. To overcome this, we have to assume that the final OFI interference pattern is agnostic to a potential initial deflection in a slightly different direction. The

same assumption has to be made regarding the wind-tunnel startup period, where the flow and WSS directions might be different from those during the actual experiment.

Here, some comments can be made: first, the oil film thins within a few seconds due to the flow acting on it, so that very quickly the only acting factor is the local WSS. Furthermore, even if the oil is initially deflected in a slightly different direction, the acting WSS can still deform the footprint of the oil. This is dependent on the surface tension of the oil and is difficult to assess. However, we expect that there would be clear asymmetries in the fringe patterns if the local forcing due to the WSS and the oil flow direction were different. This was only observed in extreme cases, such as very close to obstacles, where the spatial gradients in WSS are also significant, so that the assumption of a constant magnitude and angle throughout the interrogation window likely does not hold. In all other cases, the fringes strongly resembled those observed in 2D flows. Nonetheless, this should be kept in mind for future applications of the technique, as it might become relevant in some cases.

2.3 Wall-shear stress magnitude

While the commonly used Fourier-transform-based approach of determining the fringe spacing λ_f (Section 1.1) works well in most cases, it has several shortcomings. For example, the observed interference patterns are not always clean and may be contaminated by dust or by imperfections in the substrate, leading to disturbances in the pattern. If the center region of the droplet is affected, the information may not be usable at all, while clear interference patterns in the lateral regions are still present. Furthermore, the manual selection of an interrogation window may cause some ambiguity, especially when applied to a large number of droplets. Therefore, an automatized method that by default makes use of the whole interference pattern was sought after.

The original approach was to use the same neural network introduced for the angle determination and have the fringe spacing λ_f as an output besides $\cos\theta$ and $\sin\theta$. While this approach produced good results, it was generally inferior to the Fourier-transform method. Even though the average gradient of $\mathrm{d}\lambda_f/\mathrm{d}t$ was mostly within $\pm 3\%$ of the Fourier prediction, the individual fringe spacing predictions did not follow a straight line (as in Figure 7(b)), but instead featured slight undulations around it. This is consistent with previous studies making similar attempts [12]. Variously sized neural networks with different architectures were tested, but none featured the same quality and robustness of the classical method for clean input data. Instead, a new approach based on an autocorrelation analysis was developed, which will be described in the following.

2.3.1 Autocorrelation approach

This method builds on the aforementioned observation that in the early region of the oil droplet, the interference pattern exhibits a set of evenly spaced circular arcs of equal diameter. As a result, a 2D autocorrelation map of the image contains a clear periodic signal in the direction of the WSS. This is demonstrated in Figure 6(a) which shows the normalized autocorrelation map R of a single (mean-subtracted) interference image, such as in Figure 5(a). It can be seen that there is a clear periodic pattern

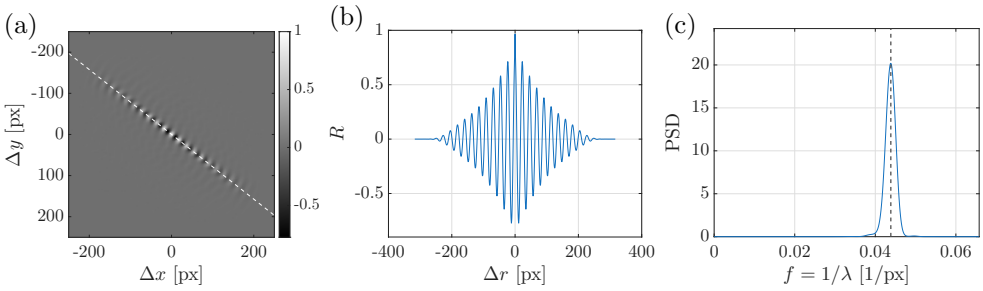

Fig. 6 Steps in the proposed method to extract λ_f from droplet images. (a): autocorrelation map of an interference pattern, (b): extracted signal along the white dashed line in (a), and (c): power spectrum of the signal plotted in (b), with the maximum location highlighted.

in the direction of the droplet oil flow, here $\theta = 218.25°$. A white dashed line at this specific angle, which was determined from the neural network (Section 2.2), is added to the figure. Now, the data along this line can be extracted by interpolation, as shown in Figure 6(b). The correlation coefficient R is plotted against the pixel shift Δr in the previously specified direction, where the first and last 20% are tapered using a Tukey window to reduce potential edge effects. Now, by calculating the power-spectral density (PSD) of that signal, the dominant frequency can be extracted – Figure 6(c) – which is the reciprocal of the sought after fringe spacing λ_f. Note that it was found that subtracting the mean spectrum across all images before identifying the maximum was generally beneficial and returned values closer to the reference. The associated errors of this method are investigated in more detail in Section 2.4.

The proposed method of determining the fringe spacing comes with several clear advantages. First, it is very easy and straightforward to implement and computationally cheap. The autocorrelation map is by default centered around the origin $(0, 0)$, so that the location of the droplet within the window is irrelevant. Also droplets of different sizes, given that the fringes are sufficiently visible, can be processed. Furthermore, the method makes use of the whole image rather than just a manually selected region, avoiding ambiguity while also making the result insensitive to small-scale noise in the interference pattern. Additionally, in cases where the central region is corrupted due to dust, scratches in the substrate or similar, valid results can still be achieved, as will be demonstrated later. Another advantage is that the algorithm returns clean wave patterns, even for relatively short interrogation windows, as shown in Figure 5.

The only input parameter to this method is the size of the frame that contains the droplet. While the CNN employed for the angle determination is limited to a fixed input size of 224×224 pixels, the correlation window can assume any rectangular size (symmetric around the center point identified by the YOLO algorithm). A larger window might contain more fringe periods, which will likely improve the accuracy of the wavelength prediction. On the other hand, the method assumes that the WSS magnitude remains constant within the window, which should be kept in mind if the flow features strong gradients. When applying the method to 2D data, a rather long window in the streamwise direction might be beneficial. It should be noted that when applying the method to real 3D flow data, the assumption of the interference pattern featuring arcs of equal diameter was not always valid. This is, for example,

the case in the image shown in Figure 5(a), which is located in a region of noteworthy WSS gradient. When investigated in detail, one can find that the diameter decreases slightly in the direction of the oil flow, which is likely due to a spatially changing WSS amplitude. However, this does not significantly affect the quality of the correlation and thus of the predicted fringe spacing and the result is expected to represent an average λ_f across the window.

Some comments regarding the accuracy of the WSS magnitude determination can be added here. Segallini et al. [5] performed an in-depth analysis of errors in OFI measurements and gave some recommendations to maximize the measurement accuracy of the method. For example, they recommend discarding the first fringe(s) in the interference pattern due to the pressure field associated with the viscous stagnation point. This cannot be implemented here, as by default the whole interference pattern in the frame is used. This error can be reduced by starting the image acquisition as late as possible, so that the film slope is small. Furthermore, it should be noted that this error is balanced by the fact that in the early times of running the experiment, the number of fringes is larger, so that the influence of a few "bad" fringes is naturally weighted less. At the same time, it should be kept in mind that there needs to be a sufficient number of wave periods visible in the frame to determine the wavelength accurately. For good accuracy, the mentioned paper recommends a minimum of around ten visible fringes. This cannot always be satisfied here due to the spatial WSS gradient and the connected requirements in interrogation window size. However, the $\lambda_f(t)$ data was generally found to still follow a straight line, even with slightly fewer fringes (see e.g. Figure 7). As will be shown in Section 2.4, this tends to long wavelengths being slightly underestimated. The found bias is below 0.5 px, but depending on the application case and desired accuracy, this should be kept in mind.

2.3.2 Sample magnitude results

Applying the described method sequentially to a time series of images, the relevant quantity for the WSS determination, $\mathrm{d}\lambda_f/\mathrm{d}t$, can be calculated. Sample fringe data from the previously investigated case is given in Figure 7(a). It shows the last image of the series, together with red lines that serve two purposes. First, they show the previously determined direction of the WSS. Second, they indicate the region in which the classical Fourier-transform-based method is applied, which is used as a reference here. Figure 7(b) compares the two methods of wavelength determination. The data follow the expected straight lines and lie virtually on top of each other. The calculated gradients of the linear fits $\mathrm{d}\lambda_f/\mathrm{d}t$ deviate by around 0.2%.

Now, we present some noisy data in Figure 8, where dust appears to have affected the central region of the interference pattern, which is used by the Fourier transform method. The displayed image is #150 in the time series, i.e. the last one. As a result, the classical method struggles to determine the wavelength of each image and the λ_f plot does not show a straight line. The autocorrelation, on the other hand, makes use of the whole image and the final λ_f data follows a straight line much better and is expected to return a gradient value closer to the truth.

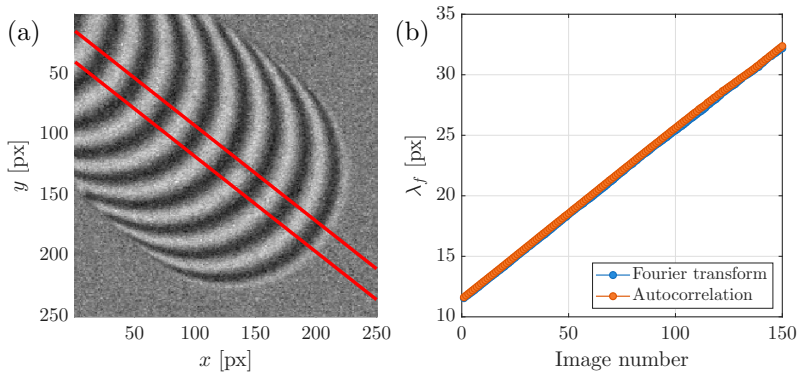

Fig. 7 Comparison of the wavelength calculation methods. (a): raw image with the red lines showing the determined flow angle and indicating the region, in which the Fourier transform is performed. (b): Corresponding results of wavelength for a series of images.

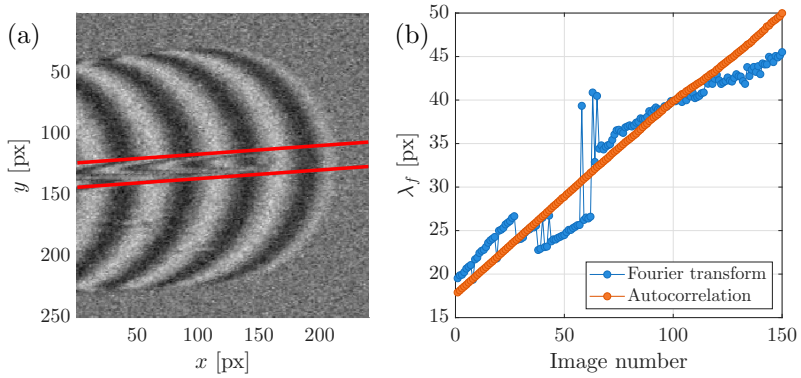

Fig. 8 Comparison of the wavelength calculation methods for noisy fringe images. For descriptions of the shown lines, see Figure 7.

The supplementary material contains the Python code to analyze fringe data using the described methods. It includes both, a wavelength prediction using the proposed autocorrelation approach as well as the classical Fourier transform method, and quantifies the deviation [17].

2.3.3 Viewing angle correction

When determining the dimensional WSS magnitude from the fringe spacing evolution $\mathrm{d}\lambda_f/\mathrm{d}t$ using Equations 1 and 4, further considerations must be made. As the cameras have an oblique view onto the surface on which the oil is placed, the viewing angle relative to the normal of the measurement surface $\tilde{\alpha}$ varies for each pixel in the dewarped image, which directly affects the WSS calculation. This can be seen conceptually in Figure 1(a). OFI experiments are not rarely performed at a camera angle around $\alpha = 45°$, which is the point where the sensitivity to the angle is largest. Hence, it is instrumental to implement an angle correction. Note that we here use $\tilde{\alpha}$ as the viewing angle and α as the camera angle.

For the following process, we assume that the camera image has already been dewarped/rectified through a camera calibration using a pinhole model. This model assumes that all light rays pass through a single point (the pinhole) before reaching the camera sensor. Besides a scaling factor "mm per pixel" – which is required for the dimensional determination of $\mathrm{d}\lambda_f/\mathrm{d}t$ – the calibration provides the extrinsic camera parameters, namely three rotation angles r_x, r_y and r_z and a translation vector $\boldsymbol{t}$. Three-dimensional rotation matrices $\boldsymbol{R_x}$, $\boldsymbol{R_y}$ and $\boldsymbol{R_z}$ can be constructed and the overall 3×3 rotation matrix is

$$\boldsymbol{R} = \boldsymbol{R_z}\boldsymbol{R_y}\boldsymbol{R_x}. \tag{10}$$

Now, the a coordinate transformation from a point in world coordinates $\boldsymbol{x_w}$ to camera coordinates $\boldsymbol{x_c}$ can be written as

$$\boldsymbol{x_c} = \boldsymbol{R}\boldsymbol{x_w} + \boldsymbol{t}. \tag{11}$$

At the camera center, $\boldsymbol{x_c} = \boldsymbol{0}$ and the corresponding position in world coordinates is $\boldsymbol{x_w} = \boldsymbol{c}$. Furthermore, the rotation matrix R is orthogonal, so that we can write

$$\boldsymbol{c} = -\boldsymbol{R}^T\boldsymbol{t}. \tag{12}$$

For a given image pixel with coordinates (u, v), the pinhole calibration provides a mapping from pixel coordinates to a corresponding physical point on the observed planar surface. In the present implementation, this mapping is given by

$$p_x = s_x u + b_x, \qquad p_y = s_y v + b_y, \tag{13}$$

where s_x and s_y are the physical length-per-pixel scale factors in the x- and y-directions, respectively, and b_x and b_y are the corresponding physical-coordinate offsets. The vector from the camera center $\boldsymbol{c}$ to the physical point $\boldsymbol{p} = [p_x, p_y, 0]$ is

$$\boldsymbol{v} = \boldsymbol{p} - \boldsymbol{c}, \tag{14}$$

where $||\boldsymbol{v}||$ is the length of the viewing ray. The normal vector of the imaged physical plane is $\boldsymbol{n_z} = [0, 0, 1]$, so that the absolute angle relative to $\boldsymbol{n_z}$ can be calculated as

$$\tilde{\alpha} = \arccos\left(\frac{|v_z|}{||\boldsymbol{v}||}\right). \tag{15}$$

This is performed for the whole camera field of view, so that for each pixel, the viewing angle $\tilde{\alpha}$ is determined and then used for the calculation of Δh in Equation 1 and hence $|\boldsymbol{\tau_w}|$ in Equation 4.

2.4 Accuracy

To estimate the accuracy of the newly developed method, we again make use of the algorithm to create artificial fringe images, generating new (unseen) data of size $224 \times$

Table 1 Overview of errors in the angle prediction. Within-run variation shows the discrepancy in angle prediction for a single droplet in time, and across-run accuracy compares the determined angle for one run to its ground truth.

	Within-run variation		**Across-run accuracy**	
Metric [°]	Clean	Augmented	Clean	Augmented
MAE	0.16	0.22	0.19	0.53
STD	0.23	0.33	0.23	0.76
P95 AE	0.41	0.65	0.46	1.53

224, as outlined in Section A in the appendix. Here, the ground truth of wavelength and direction is known and we first use the CNN to infer the flow angle and then compare the two methods of wavelength determination to the reference. Note that for this part, the images were not as strongly augmented as shown previously in Figure 3. To have more realistic data, no random patches or slices were deleted. However, the analysis was run twice, once for clean data (other than noise/blurring) and in the other one, the simulated "dust" leading to variations in the interference patterns was retained. A total of 2000 batches, each containing 50 images in "time" (with different time steps) was used to arrive at converged statistics.

The variations in the angle determined by the CNN are displayed in Table 1. Here, we separate the variation within one run relative to its circular median, as shown in Figure 5(b), from the across-run accuracy. While the first one quantifies the consistency of the prediction for a given (fixed) input angle, the latter compares the median prediction from one run to the labeled truth, which is the more relevant parameter for the actual WSS angle prediction. For both measures, we report the mean absolute error (MAE), the standard deviation of the signed angular errors (STD), and the 95th percentile of the absolute errors (P95 AE). Generally, the angle MAEs are well within 1°, even for deteriorated data, giving confidence in the method. As expected, the augmentations overall increase the across-run inaccuracies and especially lead to a larger 95th percentile AE, which indicates a larger tail of errors. This is likely a result of the applied randomness in the augmentations, where a few extreme cases exhibit a stronger deterioration of images and thus make the angle prediction harder. This is, however, still just above 1.5°.

Now, we can calculate the wavelength in the direction of the previously determined angle, using either the autocorrelation or the Fourier transform method, as shown in Table 2. Here, the first four rows show the pooled error for each individual image and the mean-absolute percentage error (MAPE) is added as well, which relates the error of each image to its individual wavelength. The last two rows compare the determined gradient of the wavelength $\mathrm{d}\lambda/\mathrm{d}t$ to the ground truth, which is again the more important quantity for the actual WSS magnitude determination. For each method, we report the mean absolute gradient error across runs together with its standard deviation ($\pm$), and the percentage errors are calculated analogously relative to the reference gradient. We can see that the new autocorrelation method performs better in all cases compared to the direct Fourier transform. This is especially true for augmented data, which is likely due to the fact that the autocorrelation makes use of

Table 2 Overview of wavelength prediction accuracies, comparing autocorrelation and Fourier transform for both clean and augmented data.

	Autocorrelation		**Fourier transform**	
Metric	Clean	Augmented	Clean	Augmented
MAE [px]	0.152	0.222	0.195	0.492
STD [px]	0.166	0.463	0.295	1.574
MAPE [%]	0.47	0.66	0.55	1.36
P95 AE [px]	0.486	0.743	0.891	1.611
Gradient AE [px/frame]	0.0063 ± 0.0032	0.0099 ± 0.0157	0.0096 ± 0.0067	0.0238 ± 0.0503
Gradient AE [%]	0.76 ± 0.33	1.11 ± 1.49	1.09 ± 0.63	2.57 ± 5.04

the whole image and therefore is less affected by local perturbations and missing data. Overall, the gradient estimation of the proposed method appears very robust, giving a mean absolute error of just above 1% even for the augmented data.

It should be noted that the pooled frame-level wavelength statistics were calculated only for images where the prediction deviates from the reference by a maximum of 100%. This filter is applied as there were very few small-wavelength images where a significantly larger wavelength was identified, which artificially inflated the error. This was especially prevalent in the Fourier transform approach, but the removed data was less than 0.2% of the total images and not more than three images in a single set. Hence, the linear gradients are computed from the whole data series using a Huber robust linear regression algorithm and it was found that the gradient results are virtually unaffected by a few outliers.

Now, we can assess the errors in the wavelength prediction as a function of the wavelength, plotted for the two methods in Figure 9. The lines show the systematic error (bias), and the errorbars indicate variability (STD) of the prediction in that wavelength bin, again plotted for clean (a) and augmented (b) data. It shows that both methods consistently underpredict the wavelength towards longer waves, with the STD increasing simultaneously. Given that the bias is approximately equal for clean and augmented data, it can be attributed to the limited number of visible fringes in the frame, which was mentioned previously (Section 2.3.1). Especially for wavelengths above 40 px, the autocorrelation consistently performs better in both regards, featuring a lower bias and a lower STD. For augmented data, the Fourier transform shows a significant amount of variation, increasing towards larger wavelengths, which is less pronounced for the autocorrelation method. It should be noted that the window size plays a significant role for the wavelength estimation, as it limits the number of visible fringes. Here, we tested the "worst case" of small windows of 224×224 pixels, where a larger window will always provide more information for both methods and is expected to reduce the error.

3 Application

After developing the method and showcasing the output for individual droplet images, we now apply it to an actual large-scale OFI experiment in a wind tunnel. Besides a

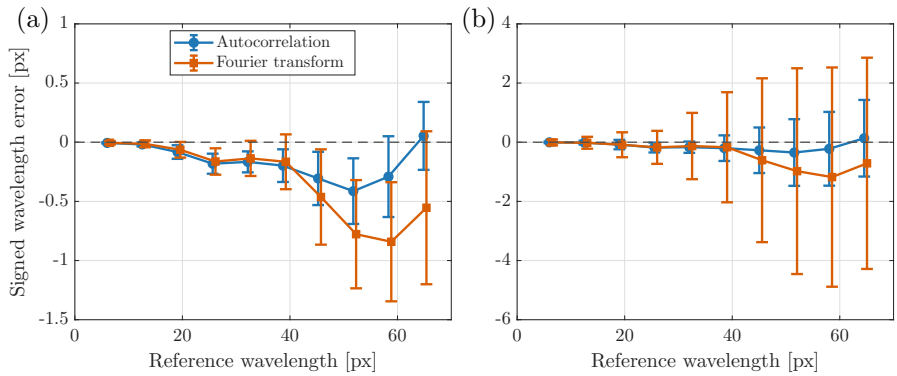

Fig. 9 Prediction error as a function of the wavelength. (a) shows clean and (b) augmented input images. The lines indicate a systematic wavelength bias and the error bars the standard deviation of the prediction in each wavelength bin. Note the difference in the vertical axis limits.

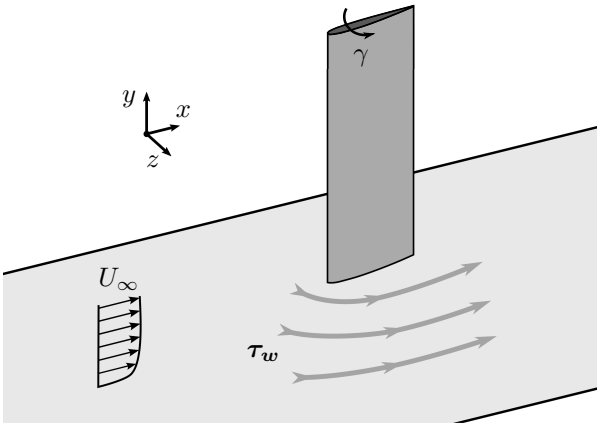

Fig. 10 Schematic of the employed experimental setup, featuring a simplified wing-body junction in the test section of the wind tunnel.

canonical 2D boundary layer, the flow case is a 3D turbulent boundary layer around a simplified wing-body junction and will be described in more detail in the following.

3.1 Experimental setup

Experiments were conducted in the Boundary-Layer Wind Tunnel (BLWT) at the University of Southampton, Department of Aeronautical and Astronautical Engineering. The BLWT is a closed-loop tunnel with a test-section of size $1 \times 1.2 \times 12\,\mathrm{m}$ (height $\times$ width $\times$ length). The tunnel speed was kept at a constant value of $U_\infty = 20\,\mathrm{m/s}$ for the current experiments and the temperature was controlled to a nominal value of $18\,^\circ\mathrm{C}$. The adjustable ceiling of the wind tunnel was adjusted to zero-pressure gradient (ZPG) conditions and the boundary layer along the smooth tunnel floor was left to develop naturally for around $7.5\,\mathrm{m}$.

To perform OFI, we used a setup similar to the one shown in Figure 1 with the cameras imaging the interference patterns from below the tunnel through a glass floor panel. To achieve a larger field of view (FOV), four LaVision CMOS CX3 25MP cameras were installed in a 2×2 grid setup, spanning a combined FOV of around 600×350 mm. Each camera was equipped with a 50 mm lens and a Scheimpflug adapter to allow to focus across the whole image. The nominal camera angles α were around $45°$, which was found to be a good compromise between FOV size and minimizing distortions in the image when dewarping. The monochromatic light source was a sodium lamp with a wavelength of 589.3 nm. As soon as interference fringes became visible on the cameras, the image acquisition was started. We recorded a total of 150 images with a time separation of two seconds. For the experiments, we used Polycraft Dow Corning 200/50 silicone oil, applied by hand from a syringe with a dispensing tip. It has a nominal dynamic viscosity of $\mu = 50\,\mathrm{cSt} = 5 \times 10^{-5}\,\mathrm{m^2/s}$. A temperature calibration of the viscosity was performed and for each run, the wind-tunnel temperature was recorded and the oil viscosity was determined. Note that prior to the experiment, the wind tunnel was allowed to run for a sufficiently long time so that we can assume the tunnel walls (and hence the oil) to be of the same temperature as the air.

For the 3D experiments, a wing was placed vertically in the test section after the ZPG distance. It spanned the entire height of the test section, acting as a simplified wing-body junction. The airfoil cross section is a NACA0012 with a chord length of 0.3 m and for the shown data, it was fixed at an angle of attack of $\gamma = 10°$. As a result, a 3D boundary layer develops around the airfoil and we employed the suggested method to study the 2D WSS distribution on the wind-tunnel floor around the wing. A schematic of the setup is shown in Figure 10.

3.2 Application to 2D boundary layers

As a first benchmark, the 2D boundary layer without the wing was investigated as a baseline case, using the developed techniques. Here, we present results from a total of 139 OFI measurements for the ZPG case, scattered across the measurement area. Figure 11(a) shows a histogram of all determined WSS angles. We can see that the data are roughly normally distributed around the mean angle close to $\theta = 0°$. The slight misalignment might be due to a rotation of the reference plate when calibrating the cameras. Otherwise, the vast majority of the angles is within $\pm 2°$.

WSS magnitude results obtained using the two methods are shown in Figure 11(b), where each data point represents one OFI droplet. It shows that for most measurements, the two methods agree to within $\pm 1\%$. The outliers can mostly be attributed to interference patterns that are distorted in the central region due to dust, where the Fourier transform struggles to pick up the dominant fringe wavelength (as for example shown in Figure 8). Note that in previous studies, OFI has been compared to a floating-element drag balance in this facility with good agreement [2]. Hence, this shows that the method can be successfully applied in 2D flows.

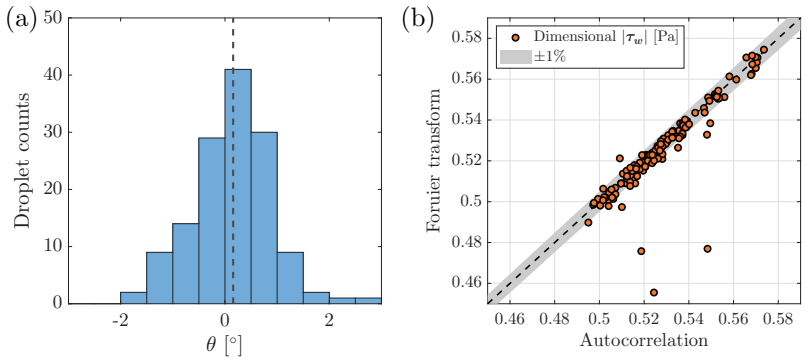

Fig. 11 Results from the OFI experiments in the 2D boundary layer. (a): histogram of all detected angles and (b): WSS amplitude $|\boldsymbol{\tau_w}|$ in Pa, determined with the two methods. The dashed diagonal lines indicates perfect agreement and the gray shaded area a variation of ±1%.

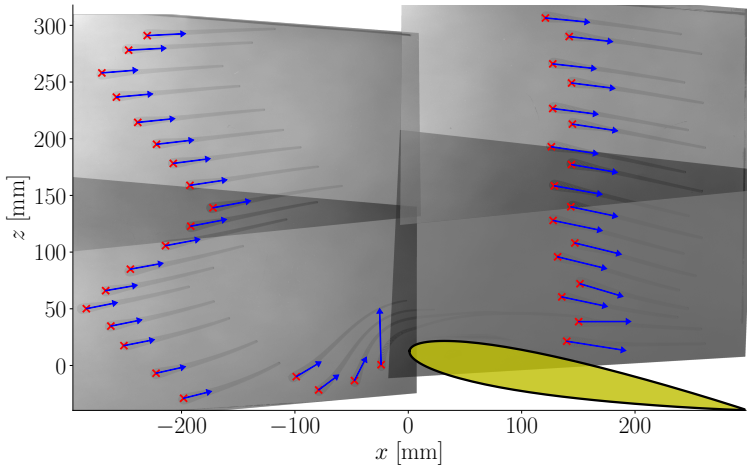

Fig. 12 Bottom-up view of the area investigated with OFI, dewarped and stitched from the four individual cameras. Shown data is the result of one measurement run with the described methods applied. The red crosses indicated droplets identified by the YOLO algorithm and blue arrows indicate direction and scaled magnitude of the local WSS.

3.3 Application to 3D boundary layers

Finally, we present results from the wing-body junction, featuring a 3D boundary layer. Figure 12 shows the output from one experimental run in a bottom-up view of the measurement area. The vertically placed wing is marked in yellow, and the droplets identified by the YOLO algorithm are shown with red markers. As before, for each droplet the employed algorithm determines first the direction and then the magnitude of the local WSS. This is shown in the figure, where the blue arrows indicate angle and scaled magnitude of $\boldsymbol{\tau_w}$.

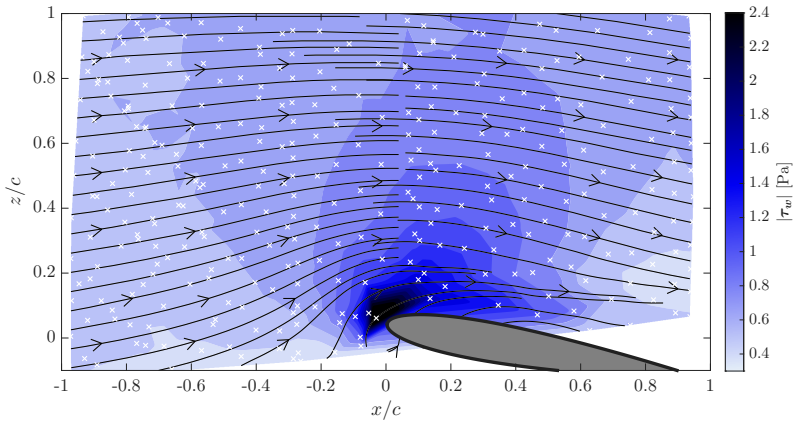

Fig. 13 Wall-shear stress $\boldsymbol{\tau_w}$ magnitude and direction in the vicinity of the wing-body junction. White markers indicate individual measurements and the data is interpolated to generate a contour plot. The streamlines indicate the direction of the local WSS (wall streamlines).

Similar runs were performed with different oil droplet locations, so that the complete WSS distribution can be determined. For a total of eight repeated runs, this data is shown in Figure 13. Here, the white markers indicate individual measurements (i.e. oil droplets) and the contour color plot is determined on the convex hull of the data points by linear interpolation. The colored contour plot indicates the absolute value of the WSS, and the overlaid streamlines represent the WSS direction, i.e. wall streamlines. The data was slightly smoothed to reduce the effect of scarce measurements near airfoil's nose. A clear maximum in $|\boldsymbol{\tau_w}|$ is visible near the leading edge of the airfoil, where the WSS also exhibits the most extreme deflection. Moving away laterally from the suction side of the airfoil, the contour lines somewhat resemble the expected shape of C_p contours around an airfoil. This is to be expected as an accelerated boundary layer features higher shear stresses by default. In fact, by calculating the friction coefficient C_f based on the local boundary-layer edge velocity, this effect can be suppressed, and regions where different flow physics arise can be highlighted, as shown in [20]. For example, indications of the corner vortex forming in the junction become visible more clearly.

4 Conclusions

In this study, we present and apply a new approach to analyze OFI interference patterns. Besides 2D flows, it allows us to quantitatively determine the vectorial WSS that a 3D boundary layer exerts on the substrate. The method combines machine learning and classical data analysis to separately find the angle and magnitude of one OFI measurement, namely, one oil droplet. The method builds on the underlying assumption that the two-dimensional similarity solution by Tanner & Blows [4] continues to be valid locally, even in 3D flows. The analyzed data indicate that this requirement is mostly fulfilled, and that even in regions of strong WSS gradients, it presents a decent first-order approximation.

The developed method follows a three-step approach. Firstly, we identify the oil droplets in a larger image using a YOLO algorithm. Thereafter, each droplet is processed separately, where first the direction of oil flow (and hence of the WSS) is identified using a CNN. Then, the dominant wavelength is determined for each time instant using a new autocorrelation approach, which is then used to find the temporal gradient of the fringe spacing. Further effects that have not been discussed previously are investigated, such as the influence of the changing viewing angle of the camera across a planar FOV.

Finally, we apply the method to experimental wind-tunnel data. First, we test it on a canonical 2D boundary layer and then in the 3D flow around a simplified wing-body junction. By performing several separate OFI experiments with oil droplets at varying locations and then interpolating, a complete map of WSS can be generated. From this, regions of maximum and minimum WSS can be identified, which can be relevant for future modeling approaches or for decisions in engineering design processes. The developed software is designed to be generally applicable, regardless of how the data was acquired. The provided package contains the source code, trained neural network weights, and sample datasets to analyze oil-film interferometry (OFI) images, complementary to the present publication [17].

The code for this study along with some sample data are openly available under the following DOI: https://doi.org/10.5258/SOTON/D4049.

Acknowledgements. We gratefully acknowledge the financial support from EPSRC (Grant Ref no: EP/W026090/1).

Appendix A Generation of artificial interference patterns

Based on the observation that in the early regions of the droplet, its shape resembles a linearly slanted cylinder, we generated this shape digitally. This is shown in Figure A1(a), where a small offset of 10% of the maximum height is added to simulate the oil edge. Note that these calculations were made in dimensional units, although for the later processing, only pixel units are used. Now, the shape is tapered slightly, smoothened using a Gaussian kernel and some large-scale noise is randomly added to simulate dust particles. The latter is achieved by low-pass filtering random noise and applying a threshold, and then also scaling the noise level with the maximum film height at each time step. Furthermore, the 3D shape is rotated by a random angle, which is set to zero for the plotted data. The result of these steps is shown in Figure A1(b).

Now, the interference pattern can be calculated. As described in the Introduction (Figure 1b), the path difference Δx between the two light waves is twice the local height of the oil. Assuming a light wavelength of $\lambda_{\text{light}} = 589\,\text{nm}$, the phase difference between these is given as

$$\Delta\phi = 2\pi \frac{\Delta x}{\lambda_{\text{light}}}. \tag{A1}$$

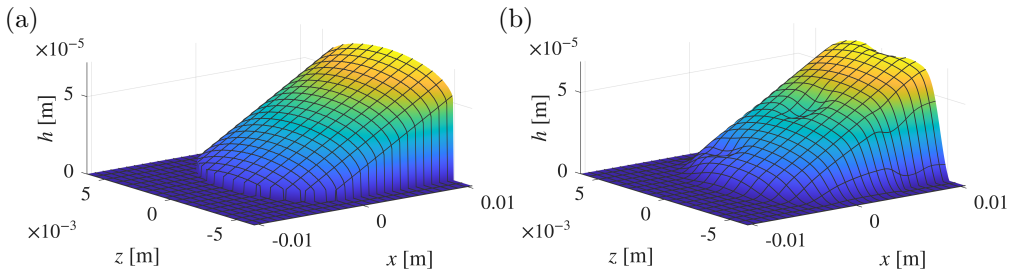

Fig. A1 Sample of a generated 3D surface resembling an oil film. (a): linearly slanted cylinder and (b): the same data, now tapered, smoothed and with some large-scale noise applied.

Now, the two phase-shifted waves are summed up to simulate interference. The amplitude of the resulting wave is

$$A = 2\cos\frac{\Delta\phi}{2}, \tag{A2}$$

where the incoming wave amplitude was assumed to be unity. The light intensity is the square of the amplitude, and using a trigonometric identity and normalizing to the interval $[0, 1]$, it can be expressed as

$$I = \frac{1}{2}(\cos\Delta\phi + 1). \tag{A3}$$

Using the relation between the gradient in oil film height $\mathrm{d}h/\mathrm{d}x$ and the fringe spacing λ_f,

$$\frac{\mathrm{d}h}{\mathrm{d}x} = \frac{1}{2}\frac{\lambda_{\text{light}}}{\lambda_f}, \tag{A4}$$

images with a desired fringe wavelength can be generated. Finally, some artificial image degradation was applied to the interference image, such as Gaussian blurring and white noise. Results are shown in Figure A2, where (a) shows the first and (b) the last interference pattern in the time series.

To make the data set as general as possible, several parameters are set randomly, but within reasonable intervals. This includes the droplet radius, the amount of tapering, the application of the artificial dust,the level of background brightness, the rotation angle, and the offset of the droplet center. Furthermore, time series of varying starting wavelengths and time steps were generated, where the latter was simulated by varying the droplet height linearly in time. Note that the angle of incoming light, which is inherently oblique (see Figure 1), was not used as a parameter here. This is justified as it can be absorbed in the path difference Δx, which was already randomized.

The corresponding MATLAB code for generating fringe patterns is openly available, along with the instructions and some sample data [17].

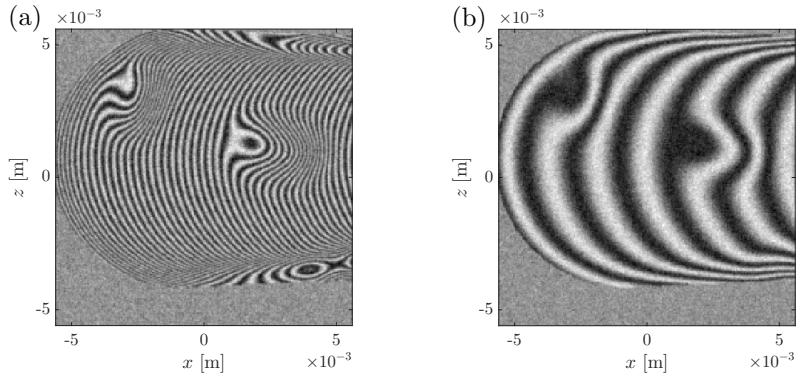

Fig. A2 Interference patterns corresponding to the 3D surfaces shown in Figure A1. (a) shows the first and (b) the last image in the time series.

References

[1] Musker, A.J.: Explicit expression for the smooth wall velocity distribution in a turbulent boundary layer. AIAA journal **17**(6), 655–657 (1979)

[2] Aguiar Ferreira, M., Costa, P., Ganapathisubramani, B.: Wall shear stress measurement using a zero-displacement floating-element balance. Experiments in Fluids **65**(4), 56 (2024)

[3] Vinuesa, R., Örlü, R.: Measurement of wall-shear stress. In: Experimental Aerodynamics, pp. 393–428. CRC Press, Boca Raton (2017)

[4] Tanner, L., Blows, L.: A study of the motion of oil films on surfaces in air flow, with application to the measurement of skin friction. Journal of Physics E: Scientific Instruments **9**(3), 194–202 (1976)

[5] Segalini, A., Rüedi, J.-D., Monkewitz, P.A.: Systematic errors of skin-friction measurements by oil-film interferometry. Journal of Fluid Mechanics **773**, 298–326 (2015)

[6] Chauhan, K., Ng, H.C., Marusic, I.: Empirical mode decomposition and hilbert transforms for analysis of oil-film interferograms. Measurement Science and Technology **21**(10), 105405 (2010)

[7] Gluzman, I., Gray, P., Mejia, K., Corke, T.C., Thomas, F.O.: A simplified photogrammetry procedure in oil-film interferometry for accurate skin-friction measurement over arbitrary geometries. Experiments in Fluids **63**(7), 118 (2022)

[8] Ruedi, J., Nagib, H., Österlund, J., Monkewitz, P.: Evaluation of three techniques for wall-shear measurements in three-dimensional flows. Experiments in Fluids **35**(5), 389–396 (2003)

[9] He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016)

[10] Vinuesa, R., Brunton, S.L., McKeon, B.J.: The transformative potential of machine learning for experiments in fluid mechanics. Nature Reviews Physics **5**(9), 536–545 (2023)

[11] Schulte-Sasse, J., Steinfurth, B., Weiss, J.: Automatic extraction of wall streamlines from oil-flow visualizations using a convolutional neural network. Experiments in Fluids **66**(5), 86 (2025)

[12] Mehdizadeh Youshanlouei, M., Lazzarini, L., Talamelli, A., Bellani, G., Rossi, M.: Automated local measurement of wall shear stress with ai-assisted oil film interferometry. Sensors **26**(2), 701 (2026)

[13] Coney, J., Denby, L., Ross, A.N., Wang, H., Vosper, S., Niekerk, A., Dunstan, T., Hindley, N.: Identifying and characterising trapped lee waves using deep learning techniques. Quarterly Journal of the Royal Meteorological Society **150**(758), 213–231 (2024)

[14] Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 779–788 (2016)

[15] Jocher, G., Chaurasia, A., Qiu, J.: Ultralytics yolo. Zenodo (2023)

[16] Terven, J., Córdova-Esparza, D.-M., Romero-González, J.-A.: A comprehensive review of yolo architectures in computer vision: From yolov1 to yolov8 and yolo-nas. Machine learning and knowledge extraction **5**(4), 1680–1716 (2023)

[17] Weingärtner, A., Aguiar Ferreira, M., Ganapathisubramani, B.: Code and Sample Data for "Machine-learning Assisted Analysis of Oil-film Interferometry Images and Application to 3D Boundary Layers". https://doi.org/10.5258/SOTON/D4049

[18] Mahendran, S., Ali, H., Vidal, R.: 3d pose regression using convolutional neural networks. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) Workshops, pp. 2174–2182 (2017)

[19] Salehi, S.S.M., Khan, S., Erdogmus, D., Gholipour, A.: Real-time deep pose estimation with geodesic loss for image-to-template rigid registration. IEEE Transactions on Medical Imaging **38**(2), 470–481 (2019)

[20] Weingärtner, A., Ganapathisubramani, B.: Experimental study of three-dimensional turbulent boundary layers in a simplified wing-body junction. In: 22nd LISBON Symposium (2026)