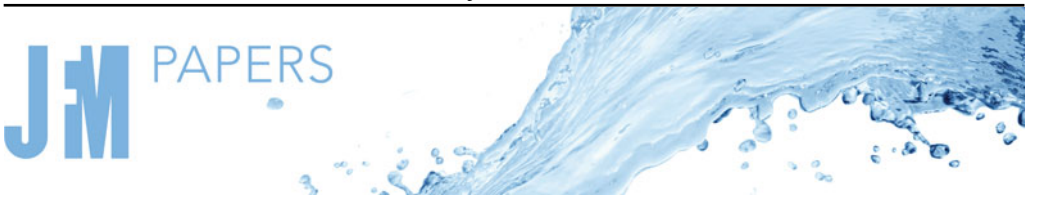
J. Fluid Mech. (2025), *vol.* 1024, A35, doi:10.1017/jfm.2025.10929



Geometry of large-scale uniform momentum zone interfaces

Guangyao Cui[1] **and Ian Jacobi**[1]

[1]Faculty of Aerospace Engineering, Technion Israel Institute of Technology, Haifa 32000, Israel
Corresponding author: Ian Jacobi, ijacobi@technion.ac.il

(Received 8 August 2025; revised 10 November 2025; accepted 10 November 2025)

Uniform momentum zones (UMZs) are widely used to describe and model the coherent structure of wall-bounded turbulent flows, but their detection has traditionally relied on relatively narrow fields of view which preclude fully resolving features at the scale of large-scale motions (LSMs). We refine and extend recent proposals to detect UMZs with moving-window fields of view by including physically motivated coherency criteria. Using synthetic data, we show how this updated moving-window approach can eliminate noise contamination that is likely responsible for the previously reported, high fractal dimension of UMZ interfaces. By applying the approach to channel flow direct numerical simulation (DNS), we identify a significant number of previously undetected, large-scale UMZ interfaces, including a small fraction of highly linear interfaces with well-defined streamwise inclination angles. We show that the inclination angles vary inversely with the size of the UMZ interfaces and that this relationship can be modelled by the opposing effects of shear-induced inclination and vortex-induced lift-up on hairpin packets. These geometric properties of large-scale UMZ interfaces play an important role in the development of improved stochastic models of wall-bounded turbulence.

Key words: turbulent boundary layers, turbulence modelling, boundary layer structure

1. Introduction

The analysis of coherent structures remains the dominant paradigm for classifying and modelling the physical processes that generate and sustain turbulent flows (Jiménez 2018). Coherent structures have traditionally been defined in terms of the intensity of highly correlated flow variables (momentum, vorticity, invariants of the velocity-gradient tensor) following Robinson (1991) and more recently, they have been defined in terms of

© The Author(s), 2025. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

their predictive or causal/explanatory power (Encinar & Jiménez 2023; Cremades *et al.* 2024).

For wall-bounded turbulent flows, the uniform momentum zone (UMZ) has become a classic category of coherent structure that represents regions of nominally uniform, instantaneous streamwise velocity. The boundaries separating distinct UMZs have been associated with thin, internal-shear layers (Meinhart & Adrian 1995); the individual vortex cores assembled along the 'backbone' of hairpin packets (Adrian, Meinhart & Tomkins 2000); the sweep and ejection quadrant events responsible for the Reynolds stresses (Laskari *et al.* 2018); as well as the more general ramp-like structures inclined downstream along the wall. However, UMZ boundaries have never been defined in terms of these various physical phenomena. Since their introduction, UMZ edges have been defined only in terms of the boundaries between regions of relatively uniform momentum as determined by a specific calculation procedure, leaving the physical interpretation slightly ambiguous.

Adrian *et al.* (2000) first established the general procedure for detecting UMZs based on calculating the probability density function (p.d.f.) of streamwise velocities, which was later refined by de Silva, Hutchins & Marusic (2016). They established best-practices for identifying the modal velocities that characterise each UMZ and showed how the number of UMZ layers increases with Reynolds number. They also modelled the zonal structure of UMZs by a superposition of self-similar attached eddies. de Silva *et al.* (2017) focused on the boundaries between UMZs, and showed how the thickness of the interfaces scales with the Taylor microscale, λ_T, and the velocity jump across the interfaces scales with the friction velocity, u_τ. They also reported that the iso-contours used to identify the UMZ interfaces exhibited fractal behaviour. This fractal behaviour was not technically a result of the UMZ itself, but was rather a consequence of the iso-velocity edge-detection procedure traditionally used to demarcate UMZs. In fact, Heisel *et al.* (2022) showed that all velocity iso-contours in wall-bounded flows (not just those associated with UMZs) exhibit fractal behaviour and they related that self-similarity to the power-law decay of streamwise velocity structure functions. Heisel *et al.* (2020) examined the wall-normal spacing of UMZ interfaces and showed that the thickness of UMZs within the log-layer can be explained by Prandtl's mixing length theory. UMZs have also been examined in the presence of free stream turbulence (Hearst *et al.* 2021) and in the atmospheric surface layer (ASL), where Heisel *et al.* (2018) showed that the UMZ interfaces follow the same trends reported in canonical laboratory studies.

Although most studies have focused on the ensemble statistics of UMZs collected over spatial snapshots, a few recent studies have considered their temporal evolution. Laskari *et al.* (2018) measured the residence time of UMZs and showed that most UMZs tend to persist for time scales much smaller than those associated with the large-scale motions (LSMs) and very-large-scale motions (VLSMs) that are energetically important to wall-bounded flows (Guala, Hommema & Adrian 2006). Laskari *et al.* (2022) tracked the spatio-temporal evolution of a simplified proxy for the UMZ interfaces (that did not correspond to physical UMZs) and showed how the temporal persistence of this proxy scaled in viscous units.

The intuitive interpretation of UMZs and their connection with so many other physical features of wall-bounded flows make them an ideal building block for turbulence modelling. de Silva *et al.* (2017) used individual step functions to represent the velocity jumps associated with specific UMZ boundaries and calculated a superposition of these step functions based on empirically obtained probability density functions (p.d.f.s) to reconstruct the mean velocity profile. Ehsani *et al.* (2024*a*) extended this stochastic approach to include full piecewise-constant velocity profiles with multiple UMZ boundaries to reconstruct the mean velocity and Reynolds stress profiles of wall-bounded

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

flows. Ehsani *et al.* (2024*b*) then further extended the piecewise-constant stochastic model to reconstruct the full spatial velocity field, including the structure of the internal layers.

A key problem with the use of UMZs as a building block for turbulence analysis and modelling is the difficulty in identifying UMZs objectively and robustly. Adrian *et al.* (2000) developed the basic identification methodology that involves: (i) calculating a p.d.f. of instantaneous streamwise velocities over a finite spatial extent, $\mathcal{L}$; (ii) identifying the anti-modal velocities of that p.d.f. which correspond to the UMZ interfaces (the modal velocities characterise the UMZs themselves); and (iii) identifying the spatial locations of the UMZ interfaces based on the iso-contours of the anti-modal velocities. de Silva *et al.* (2016) studied the effect of the field of view (FOV), $\mathcal{L}$, on the UMZ detection. If the window is too small, the p.d.f.s are poorly converged; too large and the p.d.f.s become overly smoothed, obscuring the structure of the UMZs. They reported an optimal size, $\mathcal{L}^+ \approx 2000$. The identification of the modal and anti-modal velocities is typically performed by automatic peak-finding in a histogram approximation of the p.d.f., although kernel estimation methods have also been advocated (Fan *et al.* 2019) and criticised (Heisel *et al.* 2022).

The real difficulty in UMZ eduction emerges in the application of the iso-velocity contours to identify UMZ interfaces. Because the iso-velocity contour is defined based on a p.d.f. computed from a limited spatial extent, the contour can never extend beyond that extent and thus identification of the full extent of large-scale UMZs is precluded by construction. However, this limitation is problematic, not least because LSMs and VLSMs contain at least half of the turbulent kinetic energy (TKE) in wall-bounded flows (Guala *et al.* 2006). Heisel *et al.* (2022) recently attempted to overcome this limitation, while continuing to employ a fixed interrogation window. They segmented a larger FOV into independent, non-overlapping interrogation windows and then employed a fuzzy clustering algorithm to connect UMZ edges across the different windows. Although they were able to detect edges up to twice the boundary layer thickness, δ, they were not able to extend the FOV to connect edges the size of LSMs and VLSMs. Moreover, their fuzzy algorithm still generated iso-velocity levels globally, from the entire FOV, which did not necessarily align with local structures inhabiting particular sub-regions of that field.

To detect UMZs over the full range of length scales relevant to wall-bounded flows, including the LSMs associated with the ramp-like structures near the wall, it is necessary to include local information about the anti-modal velocities over regions that extend beyond the usual fixed FOV, but without expanding the window size and without relying on globally defined iso-velocity levels. The solution to this problem is to slide the FOV in space (or time) and collect anti-modal velocities locally at each spatial (temporal) position. However, this means that the UMZ edges must be constructed from a collection of locally defined anti-modal velocity points and not the continuous iso-velocity contours that had previously been used. So far, two studies have employed this local, moving window approach – Laskari *et al.* (2018) in a water channel and Heisel *et al.* (2018) in the atmospheric surface layer – although they did not focus on differences between their UMZ edges and those detected by the traditional approach, even though this change in approach can alter the properties of UMZs detected, particularly their size, shape and inclination.

In this study, we examine the impact of a local, moving window UMZ detection scheme on the identification of large-scale UMZ interfaces and their size-dependent properties. In § 2, we refine previous efforts at the moving window procedure and introduce physically motivated coherence criteria for the local, anti-modal points. Using this updated moving window approach, we show how the resulting UMZ interfaces detected from channel flow simulations do not exhibit the fractal behaviour associated with the iso-velocity contour approach. We then develop a synthetic UMZ model to explain how the fractal behaviour

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

can emerge from iso-velocity contours due to non-physical noise contamination under the traditional fixed window approach, and thus why the smoother UMZ interfaces found via moving windows are more likely to reflect physically meaningful boundaries. We also show evidence for a significant presence of UMZs on the scale of LSMs that had not been observed previously, along with a range of UMZ edges smaller than the fixed window sizes which were also impossible to detect without a moving window. The large-scale UMZ edges also exhibit a strongly linear and well-defined downstream inclination, not previously reported in edges detected via fixed windows. In § 3, we discuss the size, shape and inclination properties of the resulting UMZ interfaces. The inclination angle appears to vary with the streamwise extent of the interfaces, and we model this behaviour in terms of a balance between mean shear and vortex lift-up effects. We also validate the trends in UMZ edge size distribution and inclination angle using a higher Reynolds number channel direct numerical simulation (DNS), confirming that they are robust features of wall-bounded flows. Finally, we provide additional evidence in support of a previously proposed hypothesis that the distribution of UMZ thicknesses in the log-layer can be explained in terms of Prandtl's mixing length theory.

2. Local UMZ detection

2.1. *Moving-window local p.d.f.*

Although most UMZ analysis has been performed on experimental velocity fields obtained from particle image velocimetry (PIV), we used velocity fields from a DNS of turbulent channel flow with half-height, h, at $Re_\tau = 2003$ by Hoyas & Jiménez (2006). The DNS provided a very long streamwise (x) extent, without compromising on the small-scale resolution, which made it well suited for a spatial moving window in the streamwise direction. The computational domain in streamwise (x), wall-normal (y) and spanwise (z) directions was $L_x \times L_y \times L_z = 8\pi h \times 2h \times 3\pi h$. The UMZ analysis was performed on streamwise/wall-normal instantaneous velocity fields that were extracted from five independent volumetric snapshots of the DNS, leading to a total number of 23 040 fields. For consistency with previous reports of UMZs from PIV data (de Silva *et al.* 2016, 2017), we interpolated the DNS data onto a uniform grid with resolution $\Delta s^+ = 50$. The influence of spatial resolution on UMZ detection is discussed by de Silva *et al.* (2014).

Figure 1(*a*) shows an instantaneous streamwise velocity field from the channel where a large-scale, ramp-like structure is clearly discernible extending for more than 2–3h. The traditional UMZ procedure identifies a single FOV of length $\mathcal{L}^+ = 2000$, marked by the dashed black line, in which the p.d.f. is calculated and the UMZs are detected. Clearly, the large ramp-like structure is bifurcated by this fixed FOV, making it impossible to capture the full extent.

The p.d.f., $p(u)$, of the instantaneous streamwise velocities, u, within the FOV is shown above the snapshot in figure 1(*a*). The p.d.f.s are calculated by kernel density estimation, following Fan *et al.* (2019), although histogram binning has widely been used in the past. The kernel estimate avoids the arbitrariness of selecting the number of bins and returns a smooth p.d.f. estimate with finite support (Silverman 2018). However, an arbitrary bandwidth, B, for the kernel smoothing must still be selected – here we choose $B = 0.012$ based on Scott (2015). The choice of kernel estimation versus histogram and its effect on the number of UMZs detected is discussed in Appendix A. The velocity p.d.f. indicates two broad modes, although the taller of the two is itself divided into two less prominent modes. The division points between modes depend on the prominence, ξ, of the anti-modes (valleys). Here, we choose a prominence of $\xi = 0.15$ that is higher than those

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

Δs^+	$\mathcal{L}^+$	B	ξ	$d_x/\Delta s$	d_y/λ_T	d_u/u_τ
50	2000	0.012	0.15	2	1	0.5

Table 1. Parameters for the UMZ interface detection: spatial resolution of the velocity field, Δs^+; FOV, $\mathcal{L}^+$; bandwidth in the p.d.f. kernel estimation, B; prominence of anti-mode, ξ; and the maximum anti-modal point separations in streamwise, d_x, wall-normal, d_y, and velocity, d_u, dimensions.

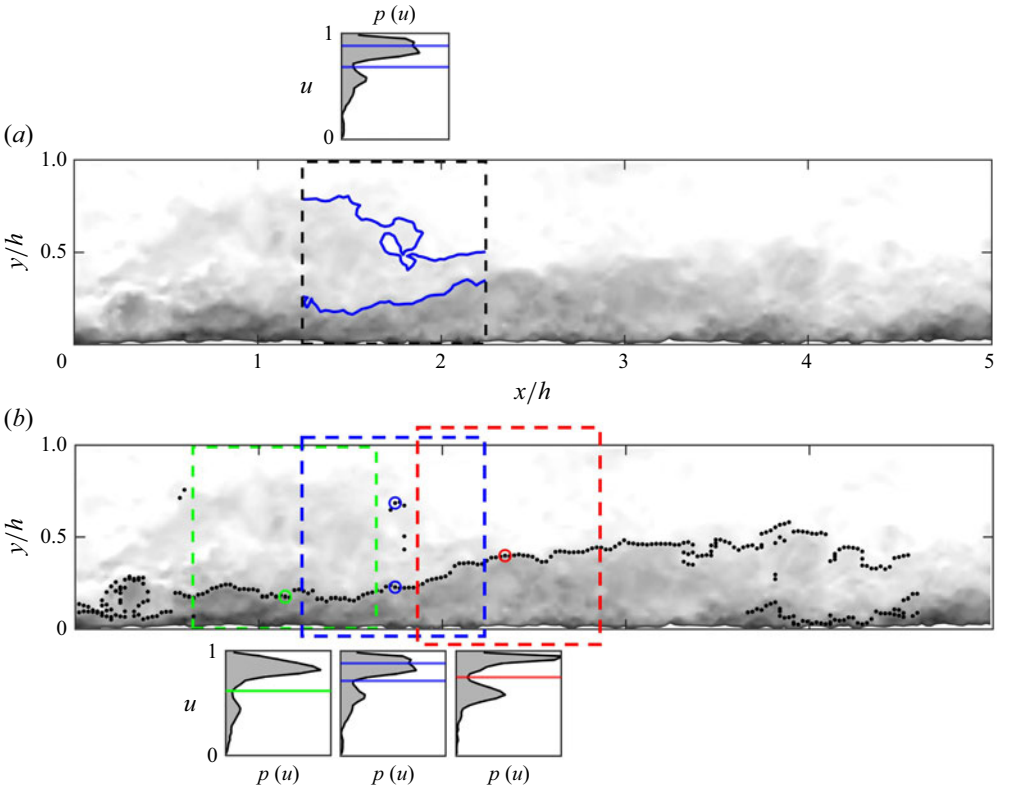


Figure 1. An illustrative snapshot of the instantaneous streamwise velocity field, u, from the DNS channel flow at $Re_\tau = 2003$ (Hoyas & Jiménez 2006). (*a*) The fixed interrogation window of length $\mathcal{L}^+ = 2000$ is shown in dashed lines with its corresponding p.d.f. above. The iso-velocity contours corresponding to the anti-modal points in the p.d.f. are shown in blue. (*b*) A sequence of three sliding interrogation windows with their corresponding p.d.f.s below, and the individual anti-modal velocity points marked at the centre of each interrogation window. The black points are the anti-modal points for all remaining moving windows.

used in previous studies. We examine the effect of this choice by percolation analysis in Appendix B. The value of the anti-modal velocity, u_{am}, was then extracted from the p.d.f. using a three-point parabolic fit of the minima between peaks, following Heisel *et al.* (2018). The parameters for the anti-mode detection are listed in table 1.

The traditional UMZ procedure locates the iso-velocity contours of the anti-modal velocity, u_{am}, within the fixed field of view, shown as the two blue lines in figure 1(*a*). Small, closed contours are typically ignored, following de Silva *et al.* (2017) and Laskari *et al.* (2018) who included only edges that extended across the whole field of view. We note that the anti-modal velocity with the smaller prominence results in a much more

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

complicated, multi-valued contour line than the anti-modal velocity corresponding to the higher prominence.

In the current UMZ procedure, we replace the fixed FOV with a moving FOV and calculate a local p.d.f. for each new FOV position in the streamwise direction, similar to what was done by Laskari *et al.* (2018) and Heisel *et al.* (2018). Three fields of view are shown in the coloured dashed boxes in figure 1(*b*), with the corresponding local p.d.f.s marked below them. (The centre blue FOV is the same as the fixed FOV above.) The anti-modal points of the p.d.f.s were identified following the same procedure, along with the iso-velocity contours within each FOV, but instead of marking the entire contour line, only the contour points located at the midpoint of the FOV are recorded as coloured circles. Therefore, each field of view contributes points at only a single x-location. Repeating this process for all streamwise locations results in the set of black points that are superposed over the velocity snapshot.

Under this procedure, it is clear that the individual anti-modal velocity points trace along the edge of the large-scale, ramp-like structure, extending far beyond the streamwise extent of any single FOV. We note that the blue point corresponding to the low-prominence anti-mode has very few anti-modal neighbours from neighbouring fields of view, indicating that its identification in the fixed window case in figure 1(*a*) was likely spurious. If a slight shift of the window results in the disappearance of the anti-mode, then the anti-mode cannot be considered part of a robust UMZ interface. However, the black anti-modal points typically have many neighbouring anti-modes and the only question is under what criteria should they be connected into coherent structures.

2.2. *Coherence criteria for connecting UMZ anti-modal points*

Each anti-modal point is defined by a spatial location (x, y) and its anti-modal velocity, u_{am}, and therefore, we construct coherence criteria for both space and velocity in terms of maximum separations between neighbouring anti-modal points: d_u for velocity, d_x for streamwise distance and d_y for wall-normal distance. These maximum separations were then used in a recursive clustering algorithm known as DBSCAN (Ester *et al.* 1996).

The choice for maximum velocity separation, d_u, was based on the scaling of velocity jumps across UMZ interfaces that was studied by de Silva *et al.* (2017) for the traditional UMZ detection scheme. They showed that the velocity jump across the shear layer scales on the friction velocity and is approximately 1–2 u_τ. Therefore, if two anti-modal points are to be considered part of a single coherent UMZ edge, we required that their anti-modal velocity separation, d_u, be bounded by the friction velocity, $d_u \lesssim 0.5u_\tau$. It is important to note that this criterion allows for some variation in anti-modal velocity within a single coherent UMZ edge, which is a fundamental difference from the iso-velocity contour-based approach to detecting UMZ interfaces that enforces a constant velocity for the edge (at least within the resolution of the contour algorithm).

A similar maximum velocity separation criterion was also used by Laskari *et al.* (2018) and Heisel *et al.* (2018), but they applied it to the modal-velocity, and they did not apply any additional criteria on the spatial coherence of the UMZ edge points. The study by Heisel *et al.* (2018) highlights a problem with connecting UMZ edges based only on modal velocities: it can result in corresponding anti-modal structures that are spatially discontinuous. (See, for example, their figure 9*d* in which the upper and lower edges of a UMZ exhibit a spatial discontinuity as a result of intervening modal structures in figure 9*c*.) The coherence criteria used by Laskari *et al.* (2018) and Heisel *et al.* (2018) are therefore useful for tracking modal velocity coherence, but are problematic from the perspective of identifying physical UMZ interfaces, which correspond to anti-modal

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

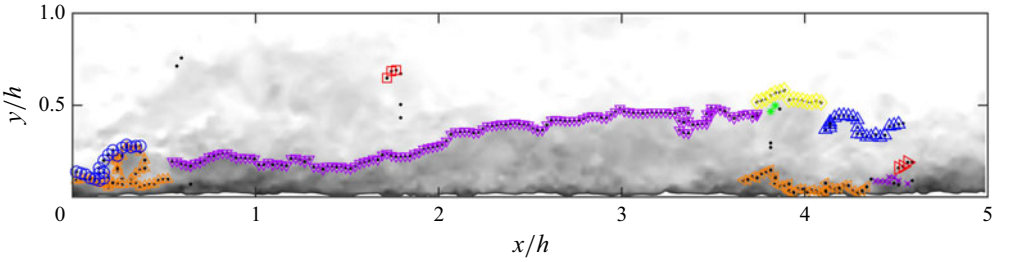


Figure 2. Clusters of the anti-modal points from figure 1(*b*) obtained using the density-based clustering algorithm. Each cluster is marked with a unique colour and symbol style.

velocity locations in space, not modal velocity locations. We therefore apply the velocity coherence to anti-modal velocities and we explicitly enforce spatial coherence as well.

The choice for maximum wall-normal spatial separation, d_y, was based on the thickness of the internal shear layers associated with UMZ interfaces, which de Silva *et al.* (2017) found to be approximately half of the Taylor microscale, $0.5\,\lambda_T$. This length scale is agnostic to the fractal dimension of the interface shape since it was calculated based on conditionally averaged velocity jumps across the interface and not the specific interface amplitude. Therefore, for edge coherence, we also required that the vertical separation between two anti-modal points, d_y, be bounded by that Taylor microscale, $d_y \lesssim \lambda_T$. The Taylor microscale, λ_T, varies across the channel; we used a constant value of λ_T evaluated at $y/h = 0.15$, following Heisel *et al.* (2021), and this choice appeared robust even at higher Reynolds number, as discussed in § 3.6.

Finally, the maximum streamwise spatial separation, d_x, was chosen to be less than twice the spatial resolution, $d_x \lesssim 2\,\Delta s$, where the factor of two allows for enough separation that a single missing anti-modal point does not interrupt an entire UMZ interface. The specific values for the coherence criteria are reported in table 1 and studied by percolation analysis in Appendix B.

2.3. *Clustering UMZ edges from anti-modal points*

Having defined the physical criteria for coherence of UMZ interfaces in terms of the spatial and velocity proximity of individual, anti-modal points, we clustered the anti-modal points together using a density-based clustering algorithm, referred to as DBSCAN (Ester *et al.* 1996), that begins with a single, isolated anti-modal point, identifies all of the points within a fixed neighbourhood of that point, measured by d_{max}, and joins them to the cluster. This process then repeats recursively for each new member of the cluster until the cluster is fully populated, before then moving to the next, isolated anti-modal point and starting again. In our case, we extended the classical DBSCAN algorithm to define the clustering neighbourhood in terms of the three independent separations for the three coherence criteria noted previously: d_x, d_y and d_u, instead of a single d_{max}. Only pairs of anti-modal points that are simultaneously inside all three neighbourhoods are joined to a common cluster. We also enforce a minimum cluster size of three points.

Figure 2 labels the clusters obtained from all of the anti-modal points illustrated by the black dots in figure 1(*b*). Each cluster is marked with a unique colour and symbol style. Note that some points are not members of any cluster. The very large scale ramp structure is now clearly identified (purple) as a single UMZ edge, something which is not possible using fixed interrogation windows. In addition to the longer streamwise extent, this UMZ interface also looks different from the edges obtained by the classical iso-velocity contour approach.

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

2.4. *Local UMZ edges versus iso-velocity contour edges*

The UMZ interfaces detected via the moving interrogation window were smoother than edges obtained via iso-velocity contours. In this section, we will explain how this smoothness is a consequence of the additional information provided by using local velocity p.d.f.s instead of fixed fields of view with iso-velocity contours. In particular, we will use synthetic data to illustrate how the iso-velocity contours in fixed windows can mistake noise for UMZ edges, whereas the moving window ignores the noise, resulting in lower fractal dimension UMZ edges.

The original UMZ technique used contours of the anti-modal velocities to identify UMZ interfaces, but UMZ interfaces and iso-velocity contours are not necessarily the same thing. A UMZ edge is defined physically to always divide between two distinct wall-normal regions of uniform velocity, whereas the iso-contour of velocity simply identifies the location of a particular anti-modal velocity value calculated across a full FOV. It does not necessarily divide two regions of uniform velocities at every local streamwise position, because the contour is based on an anti-modal velocity value that is calculated from the entire FOV, not each local streamwise point. For instance, if we examine a point of an anti-modal contour in the middle of the FOV, all of the data used to calculate that contour are located within $0.5\mathcal{L}$ of the point itself. Whereas if we look at a point of the contour on the far left (or far right) side of the FOV, then only half of the data used to calculate that contour was located within $0.5\mathcal{L}$; the other half of the data was farther away and thus less relevant to describing the local velocity distribution and UMZs.

Unfortunately, there is no way to avoid the fact that contours of anti-modal velocities will always be contaminated with non-local velocity information in their p.d.f.s, since the FOV for the p.d.f. must have a finite width. Nevertheless, the amount of this non-local contamination has a direct influence on the extent to which the resulting UMZ edges represent physically meaningful UMZs or just random velocity fluctuations with the same nominal value as the anti-modal velocity.

To illustrate this contamination effect, we developed a synthetic data set with two neighbouring UMZs separated by a UMZ edge and compared the resulting analysis with a UMZ obtained from the channel flow DNS, as shown in the three columns of figure 3.

In figure 3(*a*), we constructed a background mean velocity profile in piecewise form, using the classic Prandtl $1/7$th power-law profile near the wall, and then two linear velocity segments farther from the wall, up to the half-height of the channel. In the linear region centred on $y = 0.5$, we inserted two adjacent UMZs, each of thickness $0.18h$ and length $0.6h$, where the upper UMZ has a uniform velocity matched to the background profile evaluated at its upper edge and the lower UMZ has a uniform velocity matched to the background profile evaluated at its lower edge, thus resulting in a velocity jump at the UMZ interface. The UMZs are situated slightly downstream of the middle of the velocity field of length $2h$. In the first scenario, illustrated in the first column of the figure, we assumed that the UMZ interface is perfectly flat. Outside the two UMZs, we added noise to the linear velocity profile, with amplitude of approximately 1 % of the local mean velocity, to represent the turbulence intensity in the real flow (which is approximately 6 % at the same wall-normal location).

In figure 3(*b*,*c*,*d*), we applied the fixed window UMZ technique for windows of length $\mathcal{L} = h$ (equivalent to $\mathcal{L}^+ \approx 2000$ for the $Re_\tau = 2003$ DNS results shown in the third column) positioned at three different streamwise locations across the velocity field. Figure 3(*b*) shows the fixed window situated at the left edge of the velocity field, in which it overlaps only the leftmost part of the UMZ edge. The velocity p.d.f. does not yield any

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

anti-mode because the UMZ contribution is negligible, and thus no iso-velocity contours were detected and no trace of the UMZ was found.

Figure 3(*c*) shows the fixed window positioned farther downstream, in which it overlaps approximately half of the UMZs. The resulting velocity p.d.f. now shows a weak anti-mode, because a sufficient fraction of the velocities in the FOV are associated with the UMZs. The anti-mode was enough to identify an iso-velocity contour across the full FOV, shown in black. It is clear that most of this iso-contour is associated with the random noise fluctuations that were previously ignored in panel (*b*), but now are included because they share the same nominal value of velocity as the anti-mode. Most of the UMZ edge in panel (*c*) therefore exhibits a high fractal dimension (resulting from the random fluctuations) that does not actually correspond to the physically smooth interface between the synthetic UMZs.

Figure 3(*d*) shows the fixed window situated even farther downstream, where it overlaps the entire UMZ edge. In this case, we see that the anti-mode (marked by red circle) in the velocity p.d.f. is stronger and the resulting iso-velocity contour mostly overlaps with the true UMZ edge, with only small amounts of high fractal dimension contamination from the noise on either side.

To summarise the results from this synthetic field: the traditional procedure for identifying UMZs via the velocity p.d.f. identifies UMZ edges that can be substantially contaminated by random velocity fluctuations, and thus have a higher fractal dimension than that associated with the true UMZ edge. This contamination effect gets worse as the size of the FOV becomes larger relative to the size of the true UMZ. So not only can the fixed FOV not fully resolve UMZ edges that are smaller or larger than the FOV, but edges that are smaller will necessarily become contaminated because iso-velocity contours that span the entire FOV are selected and these naturally include random contamination.

Figure 3(*e*) shows results from the moving window approach applied to the synthetic field in panel (*a*). The moving window succeeded in identifying precisely the physical UMZ edge with only a few, very short segments of contamination. The moving window excluded most of the noise by two mechanisms. First, when the moving window does not substantially overlap a UMZ, no anti-modes are detected and thus no local anti-modal velocity points are included, just like in the fixed window in panel (*b*). Second, unlike the fixed window approach, when the moving window overlaps a UMZ enough to produce a local anti-mode in the p.d.f., like in panels (*c*) and (*d*), the coherency criteria eliminate any anti-modal velocity points associated with noise. Because these noise points are not actually part of a coherent structure and are just the result of randomly matching the nominal anti-modal velocity, they are often separated by a distance greater than the coherency criterion, $d_y = 0.05h$ (based on the Taylor microscale that demarcates the physical size of UMZ edges) and thus they will not be connected to neighbouring anti-modal points. The d_y scale is marked by the blue scale bar in panels (*e*) and (*j*).

However, this raises a question: if the coherency criteria exclude noise points near a smooth UMZ edge, could they incorrectly exclude points associated with a rough UMZ edge? To answer this, in the second column starting with figure 3(*f*), we constructed the same base flow as in panel (*a*), but this time, we used a random UMZ edge between the two UMZs, where the edge amplitude still falls within a Taylor microscale (since the microscale was found to characterise the thickness of the edge based on velocity and independent of edge geometry considerations (de Silva *et al.* 2017; Heisel *et al.* 2018)). Figure 3(*g*,*h*,*i*) shows the UMZ edges extracted from the three fixed windows. As before, the detected UMZ edges included significant contamination regions beyond the location of the true edge. However, in figure 3(*j*), we see that the moving window succeeded in identifying the random but real UMZ edge with minimal noise contamination.

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

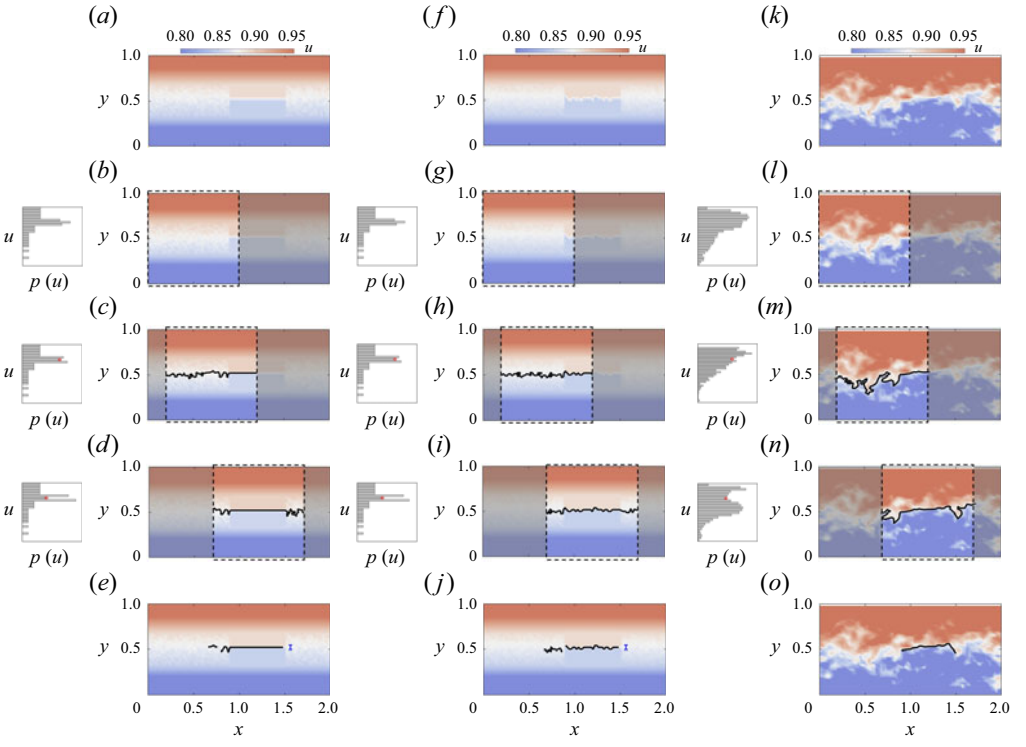


Figure 3. Synthetic UMZ fields analysed by fixed window and moving window approaches to illustrate the vulnerability of the fixed window iso-velocity contour to noise contamination. The first column analyses a velocity field with two synthetic UMZs (0.18×0.6 in outer units, fixed above and below $y = 0.5$) separated by a flat interface; the second column has the same UMZs with a rough interface; the third column is an instantaneous field from the channel flow DNS. The first row is the velocity field. The second row is the fixed window analysis for an FOV upstream of the UMZs, with the velocity p.d.f.s to the left. The third row is the fixed window for an FOV overlapping half of the UMZ. The fourth row is the fixed window for an FOV overlapping all of the UMZs. The fifth row is the moving window analysis. The same coherence criteria are used in all three moving window cases and d_y is marked by the blue scale bars in panels (e,j).

Therefore, the coherency requirements do not bias the detection against edges with high fractal dimension. The coherency requirements eliminate random noise that is not confined to the Taylor microscale, but preserve edges of any shape – smooth or rough – that fall within this scale.

In the third column of figure 3, we applied the same procedure to a representative snapshot from the channel flow data. As in the synthetic cases, we see that the fixed window in panel (m) identified a UMZ edge that is mostly composed of random points that were not detected as anti-modal in panel (l). The moving window method in panel (o) eliminated most of these spurious points and revealed a UMZ edge that appears relatively smooth. However, unlike in the synthetic cases shown in the first two columns, for the real data, there is no way to decide which of these methods better captures an underlying objective interface because no such interface is known *a priori*. Thus, we argue that both algorithms appear to perform consistently across synthetic and real data, and because we validated the moving window approach for the synthetic cases, we can have reasonable confidence in its identification of interfaces for the real channel flow data.

Because the moving window method can equally detect smooth or rough edges, the fact that most detected edges from the channel flow are smooth supports the hypothesis that

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

the high fractal dimension previously associated with UMZ edges is mostly a product of the iso-contour algorithm, and does not necessarily reflect physical edges. UMZ interfaces actually tend to have low fractal dimension when they are detected using local p.d.f.s, because the additional local information and coherence criteria avoid the spurious contours.

Therefore, the moving FOV approach results in the appearance of smoother UMZ interfaces – not through a loss of information by smoothing data, but rather due to an increase in local information and the reduction of spurious, non-local contour values. From this perspective, the absence of physical UMZs can generate spurious interfaces with high fractal dimension when calculated using a fixed window approach and, therefore, segments of UMZ interfaces with high fractal dimension should be interpreted cautiously. We calculated the fractal dimension, D, for the synthetic interfaces shown in figure 3 using the box-counting method following de Silva *et al.* (2017). For the smooth UMZ in the first column, the fixed window approach yielded an inflated dimension, $D \approx 1.18 \pm 0.08$ (similar in magnitude to the dimension $D \approx 1.2$ obtained via the fixed window approach of de Silva *et al.* (2017) and Heisel *et al.* (2022)), whereas the moving window yielded a more accurate dimension close to unity, $D \approx 1.01 \pm 0.03$. For the rough UMZ interface in the second column, both techniques detected the higher fractal dimension: the fixed window approach yielded 1.18 ± 0.05 and the moving window yielded 1.16 ± 0.06. When the true interface has a high fractal dimension, both techniques yield accurate estimates, but if the true interface is smooth, the fixed window approach inflates the dimension as if it were rough, whereas the moving window approach reports the dimension accurately. These fractal dimension calculations were also performed over an ensemble of 1000 randomly generated synthetic frames for both smooth ($D \approx 1.15 \pm 0.06$ for fixed versus 0.98 ± 0.03 for moving window) and rough (1.16 ± 0.06 for fixed versus 1.13 ± 0.08 for moving window) interfaces, and the results were nearly the same as those reported for the single frames in figure 3. Furthermore, the fractal dimension of the DNS data for the fixed window method (with window size $1h$) was found to be $D \approx 1.19 \pm 0.05$ (again, similar to the value $D \approx 1.2$ reported by de Silva *et al.* (2017) and Heisel *et al.* (2022)), compared with the lower value, $D \approx 1.04 \pm 0.04$, for the subset of edges with roughly the same extent as the fixed window ($1 \pm 0.1h$) detected by the moving window approach.

2.5. *Streamwise extent*

The streamwise extent, Δx, of the UMZ interfaces detected by the moving window approach ranges far beyond the fixed field of view, $\mathcal{L}^+ \approx 2000$, used in most previous studies, including many interfaces at the scale of LSMs. In fact, even compared with the two previous moving-window studies, the present study finds a more significant population of large-scale UMZ interfaces. Figure 4 shows the p.d.f. of the streamwise extents for all detected UMZ edges (roughly 1.9×10^6 in total) for the current analysis based on anti-modal velocities (black), in comparison with previous moving-window studies by Heisel *et al.* (2018) (blue dashed) and Laskari *et al.* (2018) (red) based on modal velocities. Both of these earlier distributions terminate for $\Delta x/h < 1$ and they appear to follow a power-law distribution, whereas the present results appear to follow an exponential distribution, at least for $1 < \Delta x/h < 3$, after which the tail is not well converged. For smaller scales, the present distribution shows a very slight upward curve, although much less prominent than the two other studies.

An exponential size distribution was also reported in an unrelated study of LSMs in channel flow at $Re_\tau = 930$ by Lee *et al.* (2014). Even if attached eddies are traditionally

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

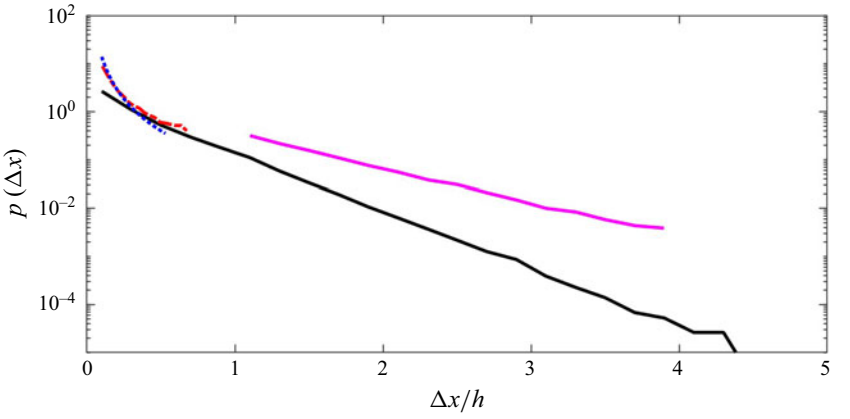


Figure 4. Probability densities p of the streamwise extent Δx for UMZ interface detected by the present approach (black solid), compared with Laskari *et al.* (2018) (red, dashed) and Heisel *et al.* (2018) (blue dashed). Laskari *et al.* (2018) excluded the very small scales and reported their streamwise size distribution starting from $0.1h$, and we applied this cutoff to the present results and those of Heisel *et al.* (2018), re-normalising the p.d.f. appropriately. In the present work, 370 edges are larger than $3h$, which accounts for 0.02 % of the total edges. The exponential distribution of LSMs from Lee *et al.* (2014) (magenta, solid) is for comparison.

assumed to follow a power-law distribution (Hu, Dong & Vinuesa 2023), it seems likely that LSMs should be represented by an exponential distribution as a consequence of superposition. When small-scale attached eddies merge together to form large-scale motions, the number of unique small scales is reduced relative to the number of larger scales (excluding double-counting) and this deforms the power law into an exponential distribution. This sort of superposition process is commonly observed in natural systems including bubble coalescence (Lovejoy, Gaonac'h & Schertzer 2004) and is usually described by a truncated power law (or a power law with an exponential cutoff, see Clauset, Shalizi & Newman (2009)), in which the probability distribution is given by $p \sim (\Delta x)^{-\alpha} e^{-\Delta x/\Delta x^*}$, and Δx^* represents some characteristic scale at which the small structures merge to form LSMs or at which the outer geometry limits the growth of the large-scales. In the semi-log plot in figure 4, the characteristic merging scale is given by $\Delta x^* \approx 0.36h$. In other words, past this size, the attached eddies tend to join together into super-structures, thereby reducing the relative count of small-scale, attached eddies.

The fact that Heisel *et al.* (2018) and Laskari *et al.* (2018) detected smaller, individual structures with $\Delta x/h < 1$ and not superposed LSMs therefore makes the power-law trend they identified quite natural, since their structures did not approach the size needed to exhibit the exponential cutoff. In fact, it seems likely that they report power-law p.d.f.s due to the rarity of LSM events in their limited data records and the inability to converge the p.d.f. in an exponential regime. Laskari *et al.* (2018) recorded 4125 turnover times (δ_{99}/U_∞) worth of data, while Heisel *et al.* (2018) had only about 40 turnover times (assuming the atmospheric boundary layer was 100 m), whereas the current study includes approximately $5.8 \times 10^5 h$ in total streamwise extent (divided in segments $8\pi h$ in length), and thus was able to converge the p.d.f. farther along the tail describing LSMs to reveal the exponential behaviour. Heisel *et al.* (2018) reported that the longest feature they detected was approximately 1.6δ, but this does not appear in their distribution for convergence reasons. Arguably, the slope of the two power-law trends approaches the same exponential regime as the scales become larger.

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

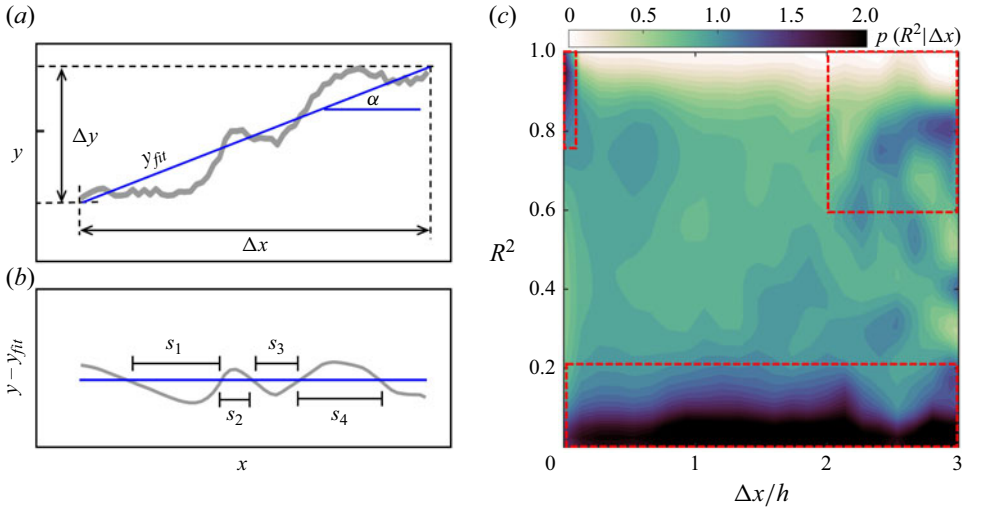


Figure 5. (*a*) Geometric properties of a typical UMZ interface (grey curve), including: the streamwise extent, Δx; the wall-normal extent, Δy; the least-squares best-fit line, y_{fit} (blue), its inclination angle α, measured with respect to the wall, and its corresponding R^2. (*b*) Residual of the smoothed interface from the linear fit, along with the streamwise distances between adjacent zero-crossings, s_i. (*c*) Joint p.d.f. of R^2, conditioned on streamwise edge extents, Δx.

3. UMZ edge properties

In this section, we characterise the UMZ interfaces detected by the moving window approach in terms of their geometric properties. In contrast to the fixed window approach, the moving window method identifies significantly longer and more linear UMZ edges, from which we can measure well-defined inclination angles. Thus, for the first time, we can meaningfully ask how UMZ edge inclinations are related to the streamwise extent of the edges and we offer a simple model to explain this relationship. We also explore the waviness/segmentation of the edges and their wall-normal distribution.

The geometry of the detected UMZ interfaces was characterised in terms of spatial extent, linearity and inclination as illustrated in figure 5(*a*). The spatial extents Δx and Δy were defined by the bounding boxes of the edges in the streamwise and wall-normal directions, respectively.

3.1. *Linearity*

To assign a meaningful inclination angle to the UMZ interfaces, we need to identify interfaces that are sufficiently linear so as to be characterisable by a single angle. A highly curved UMZ interface would exhibit a varying inclination angle along its extent, without any single meaningful value. Therefore, we first examine the linearity of the UMZ interfaces, quantified by the coefficient of determination, R^2, which represents the fraction of the variation in the interface that is explained by a linear model, y_{fit}. In addition to R^2, we examined the residuals of the linear fit to the interface, shown in figure 5(*b*), and we observed that they tended to form a non-periodic, wavy pattern. Therefore, R^2 represents the relative balance between the linear and wavy contributions to the interface shape. The inclination angle, α, for each interface was calculated based on the linear model best fit. Because the residuals tended to oscillate symmetrically about the best-fit line, this angle appears to be a meaningful measure of edge orientation even for edges with low R^2.

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

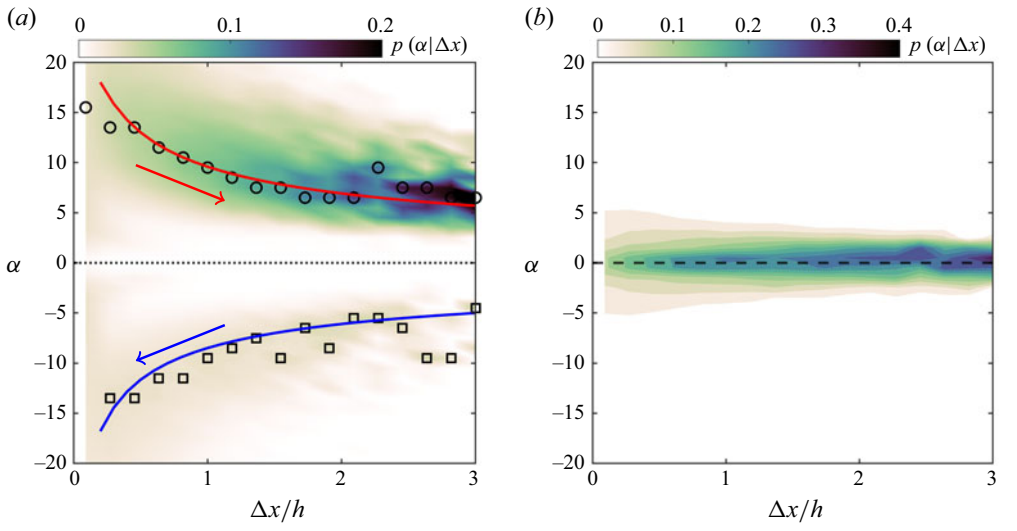


Figure 6. UMZ interface inclination angle p.d.f. conditioned on streamwise extent for (*a*) strongly linear $R^2 > 0.6$ and (*b*) weakly linear $R^2 < 0.2$ interfaces, with modal points marked as circles. The red curve indicates the lift-up/shear model from (3.10) with fitted parameters $\Delta x_0 \approx 0.2h$ and $\gamma \approx 0.5$. The blue line represents the fitting of the same model from (3.10) to the upstream (negative) inclined interfaces, with $\Delta x_0 \approx -3h$ and $\gamma = -0.09$. The arrow represents the direction of evolution under shear in each case.

Figure 5(*c*) is a p.d.f. of R^2 conditioned on the interface extent, Δx (to avoid bias due to the rarity of large-scale interfaces) and indicates that wavy interfaces (small R^2) exist across all scales, but strongly linear interfaces (large R^2) are mostly confined to the smallest and largest scales. The high linearity of the smallest interfaces is just a trivial artefact of line-fitting over very short extents, but for the largest scales, with sizes $2 < \Delta x/h < 3$, the linear structure is not trivial and indicates a high degree of organisation in the largest scale UMZs. (The figure is truncated at $\Delta x/h = 3$ due to convergence limitations.) These strongly linear, large-scale interfaces appear most prominent for $R^2 \gtrsim 0.6$, whereas the weakly linear interfaces at all scales are most prominent for $R^2 \lesssim 0.2$. Therefore, we examine the inclination and structure of these two regions of interfaces separately. The subsequent analysis is relatively insensitive to the precise choice of these linearity thresholds in R^2.

3.2. *Inclination*

The distribution of inclination angles for strongly linear UMZ interfaces is shown in figure 6(*a*), conditioned on the interface extent. The interfaces tended to be inclined downstream with angles that start around 10°–15° for the shortest extents, and then decrease with increasing edge extent until the angles saturate at around 5°–10° for the longest interfaces. There are a negligible number of upstream inclined interfaces that follow a mirrored distribution.

The 10°–15° inclinations are typical of the ramp-like structures that were first reported by Head & Bandyopadhyay (1981) and subsequently confirmed over a wide range of Reynolds numbers by Marusic & Heuer (2007). Christensen & Adrian (2001) associated these angles with the aligned heads of hairpin vortex packets, and according to Bandyopadhyay (1980), the shallow angle of the packet is a consequence of the stretching and growth of individual hairpins that are each inclined with steeper angles, closer to 45°, but collectively stretch at a shallower angle. The 5°–10° inclinations are consistent

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

with competing reports about the ramp-like structures of Liu, Adrian & Hanratty (2001) and Vesely *et al.* (2009), who both reported 6°–8° inclinations from 2-point correlation maps. The discrepancy between these reports has not, to the best of our knowledge, been explained, although these earlier studies did not examine the variation of inclination angle with structure size.

Subsequent studies sought to isolate attached eddies to examine their scale-dependent inclination angles without any contamination from the surrounding flow. Deshpande, Monty & Marusic (2019) performed this discrimination using a spanwise size filter to detect only large-scale attached features, whereas Cheng, Shyy & Fu (2022) used linear stochastic estimation to filter for wall-attached structures. Cheng *et al.* (2022) found that the largest scale structures have an inclination of 14°, consistent with previous reports by Baars, Hutchins & Marusic (2016). However, as the scales become smaller, the inclination angle was found to decrease and then reverse sign, and the authors dismissed these results as the result of contamination by detached motions with random orientations. Deshpande *et al.* (2019) also found the same 14° inclination for structures without any spanwise size filter, but as they filtered out structures with small spanwise extents, they found that the inclination angle increased towards 45°. They also attributed the shallow inclination angle to contamination from smaller, detached scales and argued that the true inclination angle of individual attached eddies was closer to 45°.

The conclusions of Deshpande *et al.* (2019) and Cheng *et al.* (2022) that the shallow inclination angles are the result of bias from detached, small-scale features differs from the physical picture of Head & Bandyopadhyay (1981), in which the shallow inclination angles are the natural consequence of attached hairpins growing and aligning into packets. The current results seem consistent with the picture of Head & Bandyopadhyay (1981), where the larger the superstructures become, the shallower their inclination.

3.3. *Modelling the size-dependent inclination*

To explain why the inclination angle appears to decrease past the consensus value of 14° for ramp-like structures as the streamwise extent increases, we first note that this trend is only visible due to the inclusion of relatively rare LSM-sized edges, which were not studied in sufficient quantity in most previous studies. We hypothesised that the increased tilting of these structures with increasing size is due to the effect of the mean shear on tilting these structures closer to the wall, impeded by a lift-up effect due to the induced velocity of the hairpin vortex packets.

We developed a simple quasi-2-D model to describe the expected effects of vortex lift-up and shear-driven stretching on the inclination angle for the back-bone of a hairpin packet. The model involves analysing two planar projections of a classical hairpin vortex, illustrated in figure 7(*a*). The hairpin is oriented in the streamwise (x) direction, with its feet taken to be at the origin, $x = 0$, from which eventually we will calculate the inclination angle based on its bounding box. The initial streamwise extent of the vortex is Δx_0, which is positive for downstream-inclined hairpins (with their heads downstream of their feet) and negative for upstream-inclined hairpins (with their feet downstream of their heads). The initial wall-normal extent of the vortex is Δy_0.

We considered the flow induced by the vortex legs on the head by projecting the entire structure onto the spanwise/wall-normal (z–y) plane in figure 7(*b*) and then treating the problem in terms of planar, point-vortex dynamics. Then, we considered the stretching of the vortex by the mean shear by projecting the entire structure onto the streamwise/wall-normal (x–y) plane in figure 7(*c*). Finally, we combined these two effects to obtain their influence on the inclination angle of the hairpin.

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

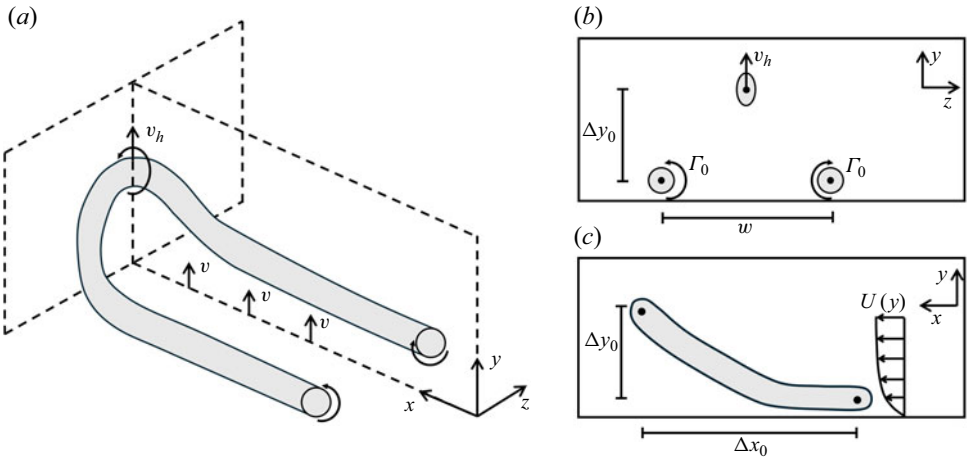


Figure 7. Illustration of the quasi-2-D model for lift-up and shear effects. (*a*) 3-D hairpin with two legs parallel to the wall, joined at a prograde vortex head. (*b*) Projection of the hairpin in the spanwise/wall-normal plane, showing the circulations of the two legs, $\pm\Gamma_0$ separated by distance w, and the hairpin head, initially separated from the legs by height Δy_0. The legs induce an upward velocity at the head, v_h for downstream-inclined hairpins. (*c*) Projection of the hairpin in the streamwise/wall-normal plane, showing the initial streamwise extent between the head and feet, Δx_0, and the mean velocity profile, $U(y)$.

For the induced velocity analysis in the spanwise/wall-normal plane, we represented the two legs as line vortices, parallel to the wall, intersecting the plane in figure 7(*b*) with fixed, equal and opposite circulations, Γ_0. One is located at $(-w/2,\, y_0)$. The tangential velocity induced by this vortex at any location $(z,\, y)$ is then given by

$$v(z,\, y) = \frac{\Gamma_0}{2\pi r}, \qquad r = \sqrt{(z + w/2)^2 + (y - y_0)^2}. \tag{3.1}$$

The parallel vortex is located at $(+w/2,\, y_0)$. The vortex head is at $(0,\, y)$ in this plane. By superposition, the total induced velocity in the wall-normal direction (ignoring the curvature of the realistic 3-D hairpin) is just the sum of the velocities induced by the two legs:

$$v_h = v(0,\, y) = \frac{\Gamma_0}{\pi r}, \qquad r = \sqrt{(w/2)^2 + (y - y_0)^2}. \tag{3.2}$$

This planar configuration is illustrated in figure 7(*b*). We assumed that the vortex head is prograde, which results in a positive lift-up velocity, $v_h > 0$, induced between legs upstream of the head (or a negative velocity, $v_h < 0$, induced between legs downstream of the head). This sign is accounted for in the sign of the circulation, Γ_0, where $\Gamma_0 > 0$ for downstream inclined vortices and $\Gamma_0 < 0$ for upstream inclined vortices.

We also assumed that the vortex is self-similar such that the distance between the legs w is proportional to the distance from the legs to the head, $(y - y_0)$ and remains so despite the lift-up and stretching, implying that $(w/2)^2 \approx (\beta^2 - 1)(y - y_0)^2$. Simplifying (3.2), we obtained

$$v_h = \frac{\Gamma_0}{\pi r}, \qquad r = \beta(y - y_0), \tag{3.3}$$

where $\beta > 1$ is just a constant representing the aspect ratio of the vortex packet.

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

The induced velocity is responsible for the lift-up of the vortex head located at position y, i.e. $v_h = \mathrm{d}y/\mathrm{d}t$, thus,

$$\frac{\mathrm{d}y}{\mathrm{d}t} = \frac{\Gamma_0}{\pi\beta(y - y_0)}, \tag{3.4}$$

and then solving the differential equation as the vortex head location shifts from y_i to y yields

$$\Delta y^2 = \Delta y_0^2 + \frac{2\Gamma_0}{\pi\beta}\Delta t, \tag{3.5}$$

where $\Delta y_0 = (y_i - y_0)$ and $\Delta y = (y - y_0)$. We note the nonlinear effect of the lift-up is because its magnitude depends on the distance over which it occurs.

We defined the time scale over which the lift-up occurs to be the same time scale over which the vortex gets stretched and tilted due to the mean velocity gradient, $\mathrm{d}U/\mathrm{d}y$, to compare the relative changes in the two coordinate directions over the same time period. This implicitly assumes that the vortex remains coherent over this time scale; if the vortex packet decoheres during this time, then we would expect different segments of the packet to lift-up out of sync with the stretching, resulting in a lower degree of linearity as measured by R^2.

For the shear-induced stretching analysis in the streamwise/wall-normal plane, we assumed that the stretching results in a change of the streamwise extent of the packet from an initial length of Δx_0 to Δx due to the velocity gradient (again ignoring the curvature of the hairpin, for simplicity), as shown in the planar slice in figure 7(c). Therefore,

$$\Delta x = \Delta x_0 + \Delta y_0 \frac{\mathrm{d}U}{\mathrm{d}y}\Delta t, \tag{3.6}$$

where the sign of Δx, Δx_0 corresponds to whether the vortex legs are upstream (+) or downstream (−) of the vortex head.

To combine the two planar effects, we solved for Δt in the shear stretching and substituted for the time scale of the vortex lift-up to obtain

$$\Delta y^2 = \Delta y_0^2 + \frac{2\Gamma_0}{\pi\beta}\left(\frac{\Delta x - \Delta x_0}{\Delta y_0 \frac{\mathrm{d}U}{\mathrm{d}y}}\right), \tag{3.7}$$

and after some re-arranging,

$$\frac{\Delta y}{\Delta x} = \sqrt{\left(\frac{\Delta y_0}{\Delta x_0}\right)^2 \frac{\Delta x_0^2}{\Delta x^2} + \gamma\left(\frac{\Delta y_0}{\Delta x_0}\right)\frac{\Delta x_0}{\Delta x}\left(1 - \frac{\Delta x_0}{\Delta x}\right)}, \tag{3.8}$$

where we defined a parameter γ to represent the relative balance between vortex lift-up and stretching, given by

$$\gamma \equiv \frac{2}{\pi\beta}\frac{\Gamma_0/\Delta y_0^2}{\frac{\mathrm{d}U}{\mathrm{d}y}} = \frac{\text{vortex lift-up}}{\text{shear stretching}}. \tag{3.9}$$

We can also write this expression in terms of: (i) the tangent of an initial inclination angle, $\tan(\alpha_0) = \Delta y_0/\Delta x_0$, where $\alpha_0 \approx 18°$ for typical packets, based on the theory of Bandyopadhyay (1980); and (ii) the tangent of the final inclination angle, $\tan(\alpha) = \Delta y/\Delta x$, after the lift-up and stretching have occurred,

$$\tan\alpha = \pm\sqrt{\left(\tan\alpha_0\frac{\Delta x_0}{\Delta x}\right)^2 + \gamma\left(\tan\alpha_0\frac{\Delta x_0}{\Delta x}\right)\left(1 - \frac{\Delta x_0}{\Delta x}\right)}, \tag{3.10}$$

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

where the sign depends on whether the vortex is downstream inclined (+) or upstream included (−), from the definition of Δx.

According to this model, the final inclination angle, α, after the effects of shear and lift-up, depends on the streamwise extent of the edge, Δx, and two fitting parameters: Δx_0 that represents some typical, initial length of a vortex packet, prior to lift-up and stretching; and γ, which represents the balance between the mean shear and the vorticity driving the lift-up. For a downstream-inclined structure, as the UMZ interface gets larger, the inclination angle should decrease due to shear, but that decrease will be mitigated by the opposed lift-up due to the vortex self-induction. In the limit of dominant shear, $\gamma \to 0$ and $\Delta x \gg \Delta x_0$, resulting in an inclination angle $\alpha \to 0°$, as expected.

The fit for this model is shown by the red line overlaying the upper part of figure 6(*a*), where $\Delta x_0 \approx 0.2h$ and $\gamma \approx 0.5$, indicating that the characteristic packet length is approximately $0.2h$ or 400 viscous units, and the mean shear effect is roughly twice as strong as the vorticity lift-up associated with the vortex packets.

For the rare, upstream inclined vortices, the same effect occurs, but the signs are reversed. The shear force will tend to increase the inclination angle, lifting the vortex away from the wall, whereas the vortex induction will operate in the opposite direction to bring the prograde vortex head closer to the wall, because the vortex legs are now downstream of the head. This orientation means that $\Delta x < 0$ and the resulting inclination angle $\alpha < 0$. Because in this scenario the shear tends to shorten the streamwise extent of the vortices, we assume (for fitting purposes) that the initial structures correspond to the largest scales, $\Delta x_0 \approx -3h$, with the smallest inclination angle, $\alpha_0 \approx -5°$. The resulting model fit is shown by the blue line overlaying the bottom half of figure 6(*a*), where $\gamma \approx -0.09$, indicating that the shear is even stronger relative to the induced vorticity for upstream-inclined vortices.

Salesky & Anderson (2020) constructed a similar model for the inclination angles of structures in a non-neutrally stable boundary layer, in which the angle resulted from a balance between characteristic velocities for the packet and the buoyancy effects, but they did not decompose the packet velocity scale in terms of a separate balance between linear shear and nonlinear lift-up effects, which is crucial to explain the behaviour observed here in the absence of buoyancy.

The current model for the inclination angle variation applies only if the UMZ interface remains coherent throughout the stretching and lift-up process, or, in other words, as long as the two effects are balanced uniformly across the interface extent. As soon as the interface decoheres, the two competing effects will result in a weakly linear interface as quantified by R^2. This perhaps explains the relative absence of upstream-oriented interfaces, which our model suggests are exposed to a very strong shear ($\gamma \ll 1$) that tends to tear them apart. For weakly linear interfaces, the distribution of inclination angles is narrowly distributed about a horizontal orientation, $\alpha = 0$, as shown in figure 6(*b*).

3.4. *Edge segmentation*

For both the strongly and weakly linear UMZ interfaces considered previously, the residuals from the linear fit always appeared wavy, as illustrated in figure 5(*b*). Stronger linearity corresponded to lower amplitudes of waviness and weaker linearity corresponded to higher waviness, but the waviness itself was a robust feature of all interfaces. This suggests that all of the interfaces can be thought of as superpositions of attached eddies (or superpositions of packets of eddies), where these coherent building blocks are arranged hierarchically for strongly linear interfaces, as sketched in figure 8(*a*). For the weakly linear interfaces, we suspect that the hierarchies of attached eddies have lost coherence due to

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

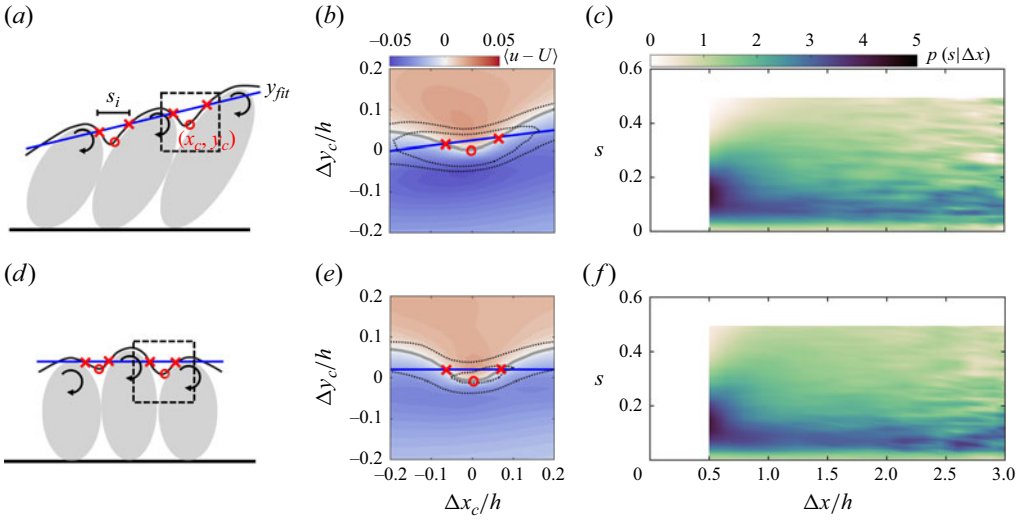


Figure 8. (*a*) Sketch of the hypothesised attached eddies (grey) responsible for generating the observed waviness in the UMZ edge, for strongly linear edges. The y_{fit} line corresponds to the average inclination, and the zero-crossings of that line, marked as ×, are used to identify the upstream- and downstream-inclined edges of the prograde vortex heads. The midpoint between the zero-crossings, (x_c, y_c), is marked as a ∘. (*b*) Conditionally averaged fluctuating velocity field about the midpoints, (x_c, y_c), corresponding to the dashed window sketched in panel (*a*). The dashed lines are iso-contours of signed swirling strength, $\lambda_{ci} = -0.6$(inner), -0.4(outer). (*c*) Segment length joint p.d.f. conditioned on interface extent. Panels (*d*), (*e*) and (*f*) correspond to the weakly linear edges.

imbalances between stretching and lift-up, and thus appear more randomly oriented, as sketched in figure 8(*d*).

To characterise the flow structures associated with this interface waviness, we identified the zero-crossing points of the smoothed residuals and recorded the segment lengths, s_i, between these zero-crossings. Zero-crossings have been widely used for the detection of boundaries between complicated coherent motions, including in the context of UMZs (Tang *et al.* 2021). The sketch in figure 8(*a*) marks these zero-crossings with × symbols and the midpoint between two zero-crossings, (x_c, y_c), is marked with a ∘. We then examined the fluctuating velocity field conditionally averaged on the midpoint between two zero-crossings – which corresponds to the region marked by the dashed black box – to observe what the zero-crossings correspond to, physically, in this scenario.

Figure 8(*b*) shows contours of the conditionally averaged flow field, where it is apparent that the zero-crossing points lie along upstream- and downstream-inclined segments of a region of coherent prograde swirl, marked in the overlaid, dashed contours. The waviness of the UMZ edges therefore captures the physical oscillations of prograde vortex structures, and each pair of zero-crossing points represents a repeating unit of this oscillation. Figure 8(*e*) shows the same conditional averaging for the weakly linear interfaces, where the zero-crossings still capture oscillatory prograde vorticity, but without any coherent streamwise inclination.

Finally, we measured the typical segment length, s, to characterise the length scale associated with the waviness. Figure 8(*c*) shows the distribution of segment lengths as a function of interface extent for the strongly linear interfaces and figure 8(*f*) shows the same distribution for the weakly linear interfaces. Despite the hierarchy breakdown for the weakly linear interfaces, the segment size distributions appear almost identical in both cases, supporting the idea that all of the large-scale edges are the result of superpositions

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

of similar, small-scale features. Additionally, in both distributions, the characteristic size of the individual wavy features constituting the overall hierarchy appears to be around $0.1h$ across most of the streamwise extents. This scale does not correspond to any of the classical length scales of the flow, but it is the same order of magnitude as the packet length scale parameter, Δx_0, fitted in the inclination angle modelling in the previous section which characterised individual structures along the UMZ interface.

3.5. *Wall-normal spacing*

In addition to the new investigations of the size, orientation and composition of UMZ interfaces, we also re-examined a previously studied aspect of UMZs – their wall-normal spacing (Heisel *et al.* 2020) – which is crucial to developing stochastic models of the instantaneous velocity field. Using UMZs identified by the traditional fixed window approach, Heisel *et al.* (2020) collected the distances between all vertically adjacent neighbour points on UMZ edges, i.e. the UMZ thicknesses, denoted here as H. They found that the modal value, H^*, of the distribution of H varied with height according to $H^*(y) \approx y$ up to $y/h \approx 0.5$, which includes all of the logarithmic layer. Based on this linear distribution, they argued that the layer thickness was consistent with Prandtl's mixing length theory, although the classic theory dictates a slope fixed by the von Karman constant, $\kappa \approx 0.4$, instead of unity.

We reconsidered the analysis of Heisel *et al.* (2020) using UMZ interfaces detected with the current moving window scheme. Unlike the iso-contours of the fixed window, in which UMZ edges extend across the entire FOV and thus the number of UMZ layers is constant in the streamwise direction within a single FOV, the moving window detection schemes mean that the number of UMZ layers varies in the streamwise direction. Therefore, any given point along a UMZ interface could have zero, one or two neighbour points on vertically adjacent interfaces and the number of neighbour points can vary along the streamwise extent of the edge. Because of the streamwise variability of the layers, we decided to count the UMZ thicknesses differently depending on whether the thickness was measured below the edge of interest, above the edge of interest or in both directions. In a case without streamwise variation, these distinctions are less important because every streamwise point has the same configuration, but with streamwise variation, we can account for the fact that some configurations are more common than others and thus some might be obscured by grouping them all together.

Figure 9(*a–c*) illustrates the differences between the thicknesses identified by these three interface configurations. Figure 9(*a*) represents cases where the point of interest at location y^{top} is strictly above a neighbouring interface, but not below any other interfaces (counting only regions between interfaces and not regions between an interface and the wall itself). The thicknesses, H, are shown in the purple filled colour, which represents a 'top-down' perspective, emphasising the thicker layers farther from the wall and under-counting the thin layers near the wall. So, for instance, the left side of the middle interface (II) has no edge above it and the thicknesses are measured only (strictly) below; whereas downstream, there is another edge (III) above (II), in which case no thicknesses below (II) are counted. In this downstream region, only the upper edge (III) is strictly above another interface, in which case we count the thickness directly below it.

Figure 9(*b*) illustrates the case where the point of interest, y^{bot}, is strictly below a neighbouring interface, but not above any other interfaces. The corresponding thicknesses are shown in the blue filled colour, which represents a 'bottom-up' perspective on the different UMZ layers that captures the individual, thinner layers closer to the wall. In this case, the thickness between (I) and (II) is captured along its full extent, whereas the

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

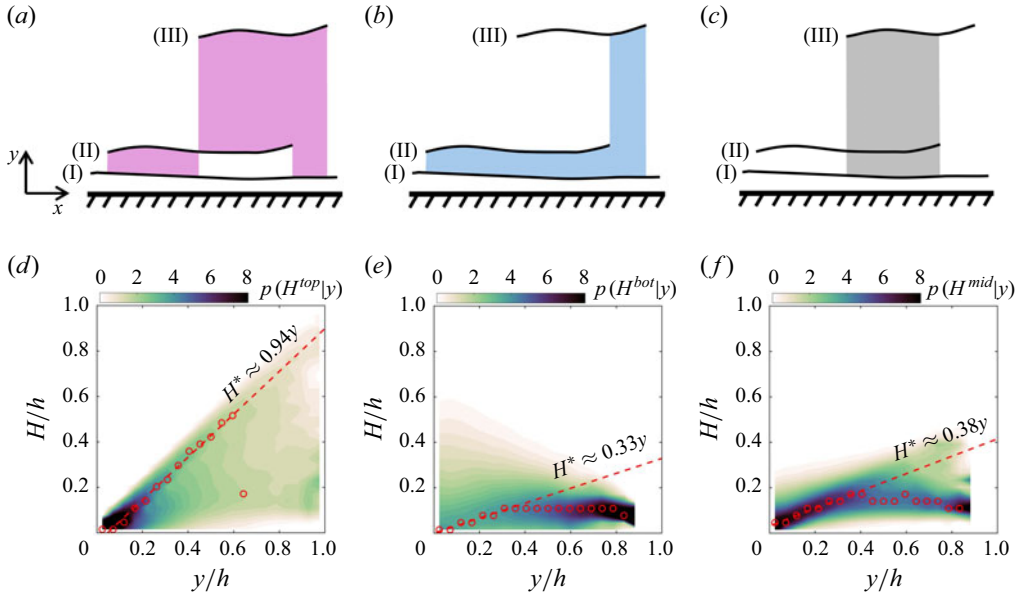


Figure 9. (*a–c*) Cartoons illustrating how different interface classification schemes affect the layer heights obtained: (*a*) y^{top} with thicknesses, H^{top}, in purple; (*b*) y^{bot} with thicknesses, H^{bot}, in blue; and (*c*) y^{mid} with thicknesses, H^{mid}, in grey. Only panel (*b*) is biased towards UMZ layers associate with near-wall structures. (*d–f*) Joint p.d.f.s of UMZ thicknesses, conditioned on y, for each of the three classifications, labelled with the slopes for the modal layer thicknesses with respect to wall-normal location.

thickness farther away between (II) and (III) is truncated. These configurations are thus biased towards the more numerous, near-wall layers that likely correspond to the more numerous, small-scale structures predicted by the attached eddy hypothesis. The first scheme in panel (*a*) undercounts these near-wall structures and allows for counting structures that have no footprint near the wall.

Finally, figure 9(*c*) illustrates the case where the point of interest, y^{mid}, is sandwiched between two neighbouring interfaces. The corresponding thicknesses are shown in grey, which represent regions of multiple, stacked layers, which are anticipated to be much less common in the moving window approach than in the fixed window approach, since the moving window removes any stacked layers that lack sufficient coherency and are likely the result of noise contamination.

These cartoons illustrate how the classification of points can drastically alter the resulting thickness distribution, compared with simply counting all thicknesses together. Out of all of the interface locations: 44 % are y^{top}, 47 % are y^{bot} and the remaining 9 % are y^{mid}. Therefore, if we are interested in analysing UMZ thicknesses associated with attached eddies or a log-layer, we should focus on those biased towards the wall, in panel (*b*), which constitute less than half of the total number of edge points. By including all of the configurations together, we would dilute any trends associated specifically with these attached structures.

The joint thickness distributions (conditioned on wall-normal location) for each of the three interface classifications are shown below the cartoons in figure 9(*d–f*), where the modal thicknesses are marked with red circles. The different interface classifications exhibit very different modal trends in layer thickness. The top-down perspective in panel (*d*) yields a roughly linear increase in thickness across all y, with slope, 0.94, close to that reported by Heisel *et al.* (2020). By contrast, the bottom-up perspective in panel (*e*) shows

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

a linear increase in thickness that is limited to the log-layer ($y \lesssim 0.2h$) and has a lower slope, 0.33, a bit closer to the typical values for the von Karman constant, and thus highly consistent with the mixing length proposal by Heisel *et al.* (2020). The middle perspective in panel (f) also shows a linear trend closer to the log-layer with a slope close to the von Karman constant.

The reason why the 'top-down' perspective yields a near unity slope, instead of a value closer to κ, is likely due to the fact that most of the UMZ interfaces in the ensemble are very close to the wall, so that the thickness recorded for any high (y^{top}) interface is roughly its height to the wall itself, y. Additionally, because y^{top} constitutes nearly half (44 %) of the total interfacial points in the ensemble, they need to be excluded to avoid a bias towards tall layers and a relative under-counting of the UMZ thicknesses in the log-layer itself. Since Heisel *et al.* (2020) did not exclude these top-down layers, their slope was likely inflated, resulting in the disagreement with the von Karman constant. By focusing on just the bottom-up perspective of layer measurements, we avoid this bias from the high interfaces and therefore provide stronger confirmation for the mixing length explanation proposed by Heisel *et al.* (2020).

3.6. *Reynolds number effects on edge properties*

The trends identified previously regarding the distribution of UMZ edge sizes, their linearity, inclination angles, segmentation and wall-normal spacing, were all validated using a second DNS channel flow at higher Reynolds number, $Re_\tau = 5200$, to confirm that these were all robust features of wall-bounded flows, and not merely an artefact of the $Re_\tau = 2003$ flow used previously. The higher Reynolds number DNS by Lee & Moser (2015) has the same domain size as the original case of Hoyas & Jiménez (2006): $L_x \times L_y \times L_z = 8\pi h \times 2h \times 3\pi h$, discretised with $N_x \times N_y \times N_z = 10\,240 \times 1536 \times 7680$ grid points. We used 15 360 streamwise/wall-normal planes from ten different time snapshots and performed the same linear interpolation to 50 wall unit resolution, consistent with our earlier analysis and works by de Silva *et al.* (2016, 2017). For the moving window UMZ algorithm, the identical coherence criteria listed in table 1 were employed, suggesting that these choices are robust to Reynolds number variation and potentially universal.

Figure 10(a) shows the distribution of streamwise extent of UMZ edges. The higher Reynolds number flow exhibits a similar exponential distribution as observed at the lower Reynolds number in figure 4, with a slightly smaller decay constant, i.e. characteristic scale, $\Delta x^* \approx 0.30h$, compared with $\Delta x^* \approx 0.36h$ for the lower Reynolds number case. The smaller scale potentially indicates that eddies begin to merge at smaller scales for high Reynolds numbers due to stronger shear. For the small scale power-law behaviour, the higher Reynolds number results align even better with those from Laskari *et al.* (2018) than did the lower Reynolds number case due to the nearly matched Reynolds numbers.

Figure 10(b) shows the linearity of the detected edges via the joint p.d.f. of R^2 conditioned on edge extent. The overall distribution is quite similar to that found in figure 5(c), but the highly linear region is more prominent in the higher Reynolds number flow, extending over a wider range of streamwise extents down to $\Delta x \approx 1h$. Based on our modelling analysis, this might be the result of increased coherence through stronger vortex stretching at higher Reynolds numbers, where the local velocity gradients are steeper.

In figure 11, we examine the relation between edge extent and inclination angle for the strongly linear edges in panel (a) and the weakly linear edges in panel (b), like we did previously at lower Reynolds number in figure 6(a,b). The trends at the two Reynolds numbers are almost identical, and thus the same modelling approach described above

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

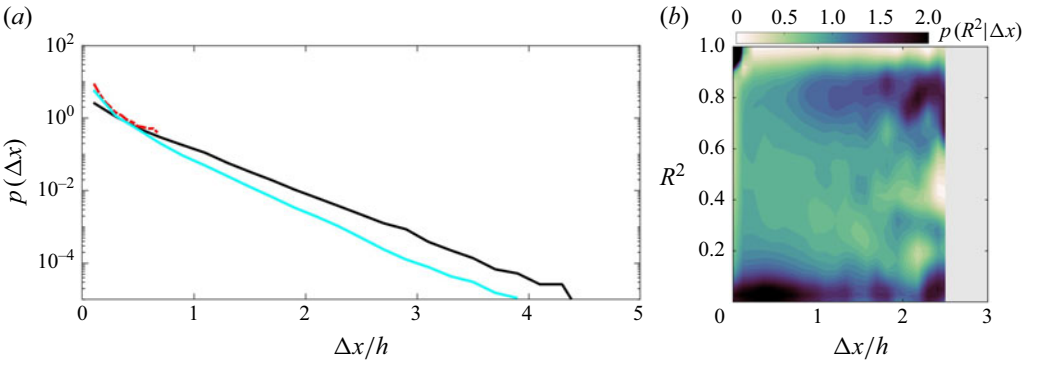


Figure 10. (*a*) Distribution of streamwise extents for UMZ interfaces: black is $Re_\tau = 2003$ shown in figure 4; cyan is $Re_\tau = 5200$. The higher Reynolds p.d.f. is very close to the results from Laskari *et al.* (2018) (red, dashed) that were measured at nearly equivalent Reynolds number, $Re_\tau = 5300$. (*b*) Joint p.d.f. of R^2 conditioned on edge extents, previously shown at lower Reynolds number in figure 5(*c*). The grey region excludes insufficiently converged p.d.f. tails: only 354 structures larger than $2.5h$ appear in this region.

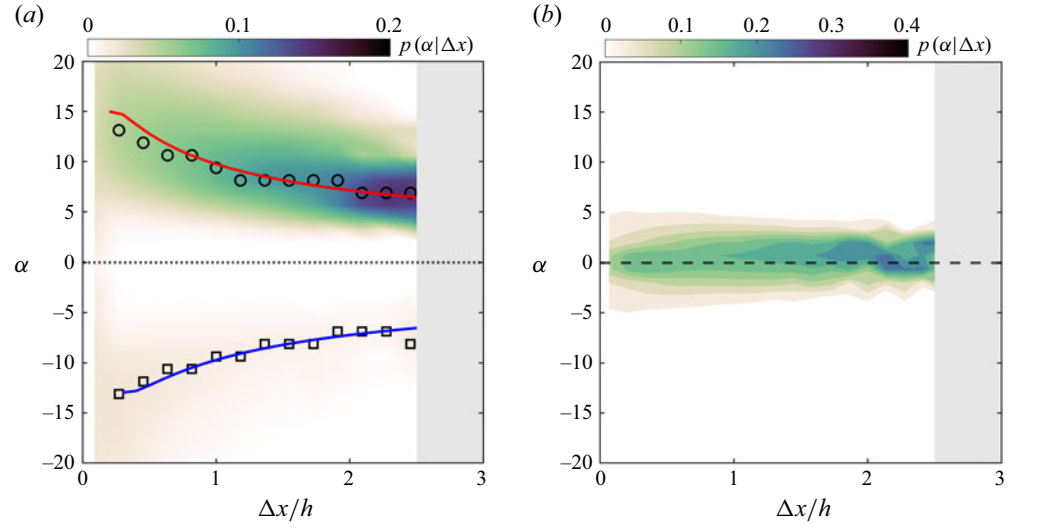


Figure 11. UMZ interface inclination angle p.d.f. conditioned on streamwise extent for (*a*) strongly linear $R^2 > 0.6$, and (*b*) weakly linear $R^2 < 0.2$ interfaces, with modal points marked as circles and squares. The red curve indicates the lift-up/shear model from (3.10) with fitted parameters $\Delta x_0 \approx 0.2h$ and $\gamma \approx 0.57$. The blue line represents the fitting of the same model to the upstream (negative) inclined interfaces, with $\Delta x_0 \approx -2.5h$ and $\gamma = -0.11$.

can capture both. For the higher Reynolds number, the fit parameters for the model in (3.10) are $\Delta x_0 \approx 0.2h$ and $\gamma \approx 0.57$, which differ just slightly from the lower Reynolds number fits. Based on the earlier interpretation of these parameters, this indicates that typical packet size and the relative strength of the shear versus vorticity lift-up effects are both largely independent of Reynolds number, although it is difficult to draw definitive general conclusions from just two cases. The segment length analysis and vertical spacing of edges are also nearly identical between the two Reynolds numbers, but are not shown here for brevity.

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

4. Conclusions

The uniform momentum zone (UMZ) has long been recognised as an important category of coherent structure in wall-bounded turbulent flows, and more recently has been employed as a building block for stochastic models of instantaneous and mean velocity fields. However, the traditional approach for identifying UMZs involved a fixed, limited field of view and thus precluded the identification of the full extent of large-scale structures. In this study, we refined and extended a moving-window approach to UMZ identification first introduced by Laskari *et al.* (2018) and Heisel *et al.* (2018), by carefully defining physically motivated coherency criteria and including far larger ensembles for better convergence of large-scale UMZ statistics. In the process, we used synthetic data to show how the moving window approach actually eliminates non-physical artefacts of the iso-velocity contour methodology, thereby reducing the fractal dimension of the resulting UMZ interfaces. These artefacts are eliminated by including local velocity information in the sliding p.d.f. and enforcing coherency criteria between the resulting locally defined anti-modal points.

The new, smoother UMZ interfaces were found to include a significant number of large-scale features, on the scale of LSMs, that were only rarely detected by previous investigators and not in sufficient quantity to perform statistical analysis of their geometric characteristics. We found that the streamwise extent of the UMZ interfaces showed an exponential distribution, consistent with the distribution of LSMs by Lee *et al.* (2014) and different from the power-law distribution found in previous UMZ studies. We argued that the exponential behaviour is the result of capturing more LSMs, which can be thought of as superpositions of smaller features, and thus are best described by a truncated power-law with an exponential tail.

We then analysed the geometry of the UMZ interfaces, with particular emphasis on their linearity, streamwise inclination and wall-normal spacing. The UMZ edges were divided into two groups based on their R^2 for a linear fit. The strongly linear interfaces exhibited a clear trend of decreasing downstream inclination angle with increasing size, while the weakly linear interfaces showed no streamwise inclination. A quasi-2-D model was proposed to explain the scale-dependent inclination of the highly linear interfaces, in which we assumed that the UMZ interfaces represented the backbones of hairpin packets that were experiencing simultaneous shear and vortex lift-up effects. The shear reduces the inclination angle with respect to the wall, while the vortex lift-up tends to increase it and the relative balance between the two effects was found to be size-dependent for self-similar hairpins. The resulting model predictions were shown to fit the trends for both the common, downstream-inclined and rare, upstream-inclined interfaces. In addition to the inclination, we also found the UMZ interfaces exhibited a waviness that indicated a secondary length scale associated with individual segments constituting each interface.

The wall-normal spacing of the UMZ interfaces was quantified by examining the wall-normal distances between points on neighbouring interfaces. Unlike previous studies, we classified the UMZ interfaces into those with neighbours above, below or on both sides, and we showed that the bottom-up perspective of sampling the UMZ edge population was able to reveal the strongest evidence in support of the mixing length theory of UMZ spacing proposed by Heisel *et al.* (2020).

Employing a moving-window approach to identifying UMZ interfaces allows the detection of physically realistic structures that correspond to classical intuitions about shear layers (without significant contamination from non-physical, high fractal dimension interfaces). The moving window approach also allows for the detection of large UMZ interfaces at the scale of LSMs. The longer UMZ interfaces also exhibit wall-inclination

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

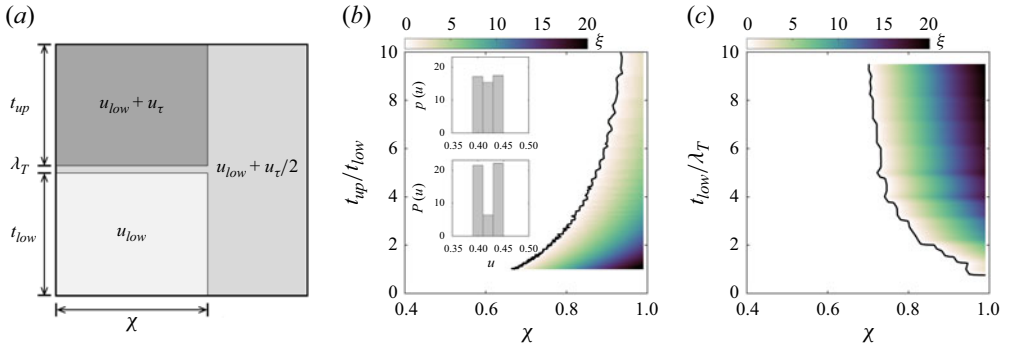


Figure 12. (*a*) Illustration of the two different UMZs on top of a background average velocity, with velocity difference u_τ, edge thickness λ_T and length χ relative to the window size. (*b*) Prominence, ξ, of the histogram anti-mode as a function of the length and asymmetry of the UMZ layer thickness. (*c*) Prominence, ξ, of the histogram anti-mode as a function of the length and ratio of UMZ to interface thickness. The black lines denote a threshold prominence cutoff of 0.1.

and wall-normal spacings that can be predicted from simple physical models associated with vortex lift-up, shear distortion and mixing length theories. The identification and modelling of these features of large-scale UMZ edges should make it possible to develop more accurate stochastic representations of instantaneous turbulent fields based on UMZs.

Acknowledgements. The authors thank the anonymous referees for very useful criticism and recommendations.

Funding. Both authors gratefully acknowledge the support of Israel Science Foundation grant 219/21 and I.J. acknowledges the Wenner-Gren Foundations grant GFOh2023-0014.

Appendix A. Kernel versus histogram

Kernel estimation has previously been criticised for under-detecting UMZ interfaces (Heisel *et al.* 2022). We can explain why this happens using the geometric argument reminiscent of the one employed in § 2.4 regarding the fractal dimension.

Kernel estimation smooths over low-prominence anti-modes in the p.d.f., and the prominence of an anti-mode is directly related to the geometry of the UMZ. Therefore, kernel estimation will necessarily under-detect UMZs with certain types of geometries. We illustrate this effect in figure 12. Figure 12(*a*) provides a sketch of two UMZs: a high-speed UMZ above a low-speed, separated by a velocity difference of u_τ, which varies linearly across an edge of thickness λ_T, and some background average velocity to their right. The length of the UMZ relative to the size of the interrogation window is χ, and the thickness of the two UMZ layers are t_{up} and t_{low}. If we employ the standard histogram approach with bin width $1/2u_\tau$, the histogram will contain only three distinct bins, the centre of which can indicate the presence of the UMZ edge via an anti-mode, if the anti-mode is sufficiently prominent. The prominence value, ξ, is defined as the depth of the histogram bin value relative to the lowest neighbouring peak bin value.

Figure 12(*b*) shows how the relative length of the UMZ interface and the asymmetry between the two UMZ thicknesses affect the resulting prominence. We see that longer interfaces and more symmetric layers result in higher prominences. Figure 12(*c*) shows the effect of the relative thickness of the UMZ layer compared with the edge thickness. Again, longer edges and thicker layers result in more prominent anti-modes.

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

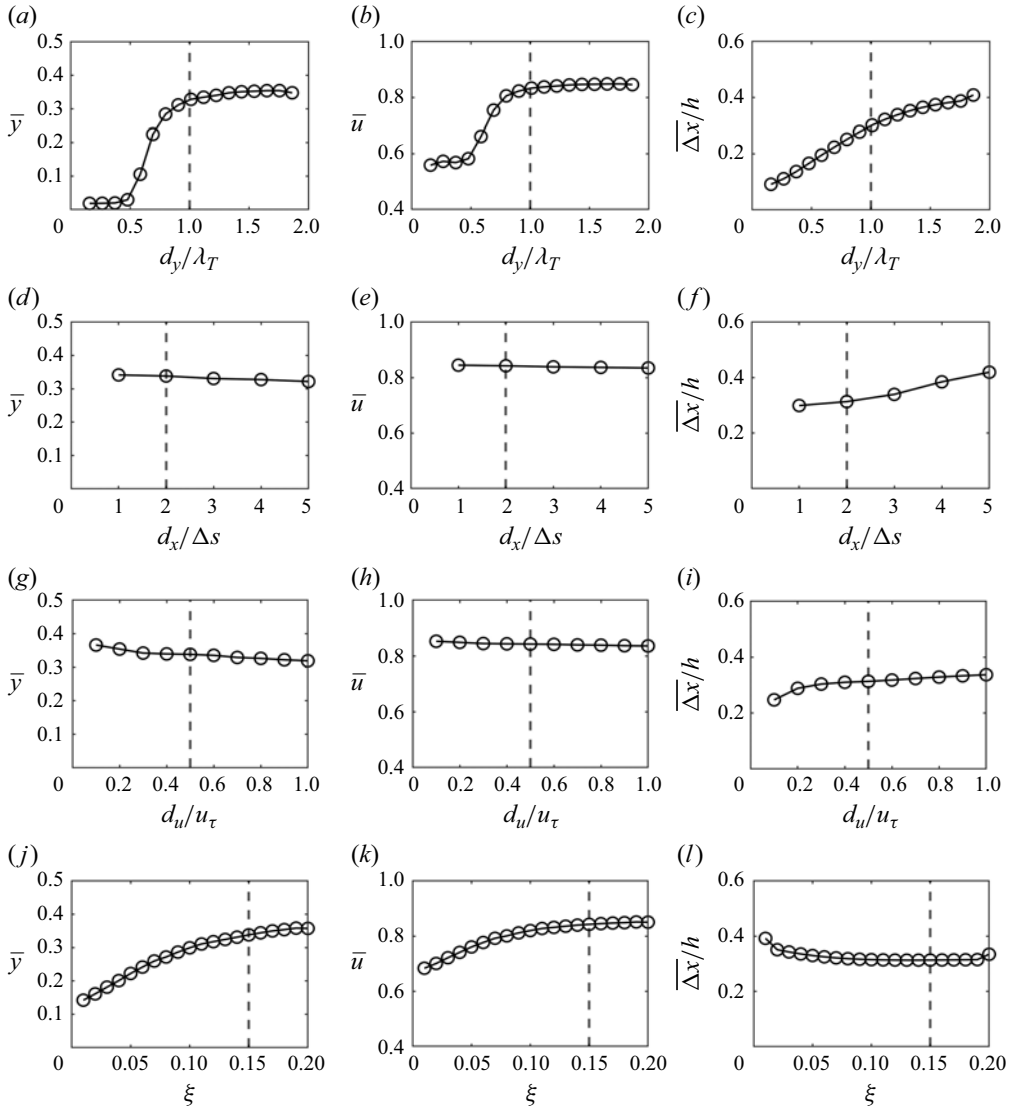


Figure 13. Percolation analysis with respect to the average wall-normal location, $\overline{y}$ in the first column, the average anti-modal velocity, $\overline{u}$, in the middle column and average streamwise extent, $\overline{\Delta x}$, in the third column, for each of the four coherency and detection parameters – d_y, d_x, d_u, ξ – in subsequent rows. The UMZ interfaces included in the plots for the first two columns are large scales with size $\Delta x/h = 2$. Dashed lines represent the value applied in the present study.

Physically, this means that the short, thin and highly anti-symmetric UMZs very near the wall (where the velocity gradient is highest) will tend to have lower prominences in a histogram approach and thus will be excluded from the kernel estimation approach. However, since our emphasis here is on the detection of large-scale UMZ edges that were previously excluded from fixed interrogation window approaches, we tolerate the loss of these small-scale edges in exchange for the added robustness of the kernel estimation. Of course, it is always possible to take the additional step to filter the histogram and remove these low prominence peaks, but that requires an additional arbitrary choice of filtering parameter, which is already accounted for in the kernel estimation bandwidth choice.

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

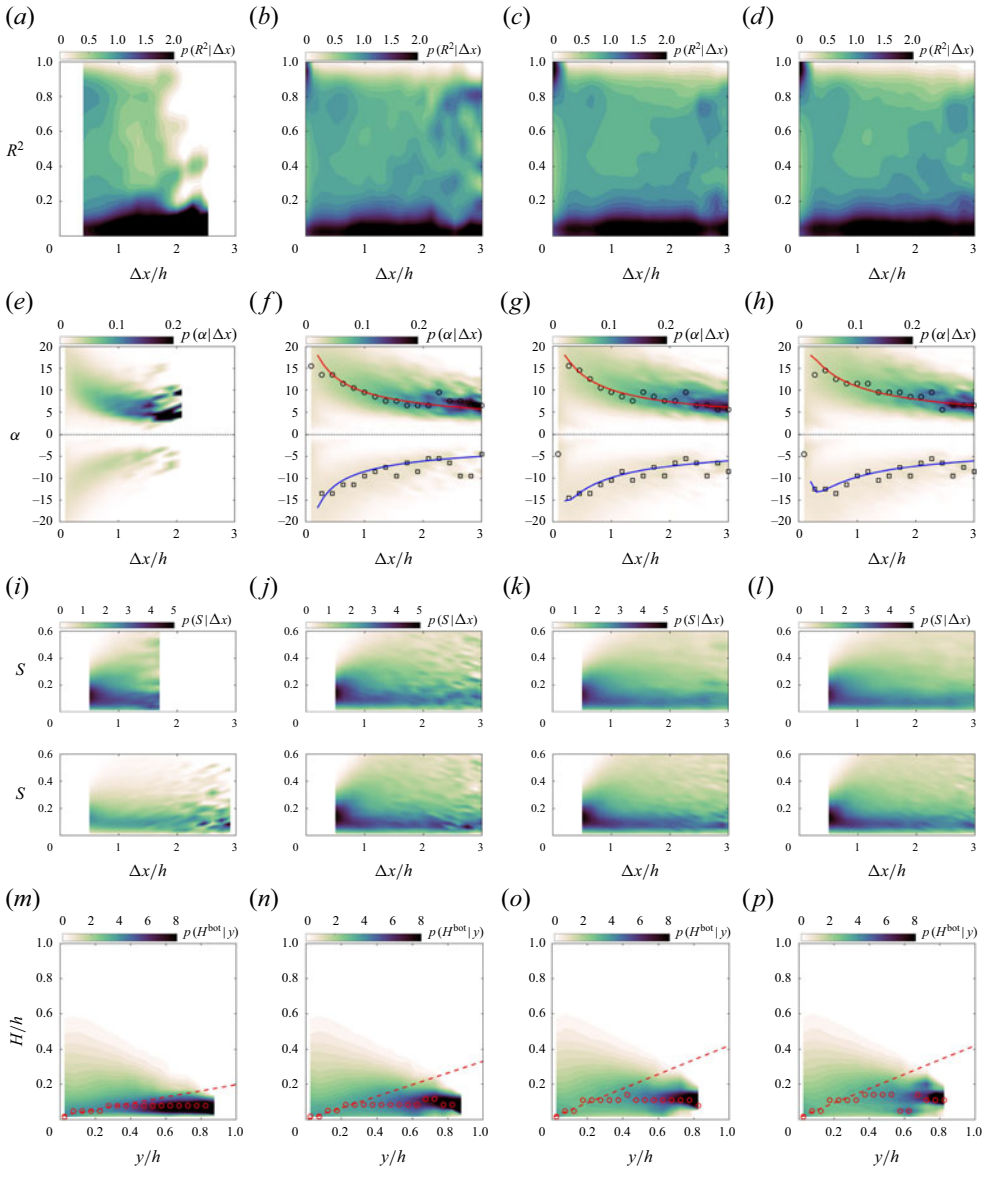


Figure 14. Percolation study of d_y: from left to right column, $d_y/\lambda_T = 0.5$, 1, 1.25, 1.5. (*a–d*) Conditional joint p.d.f. of streamwise size and R^2 (corresponding to figure 5*c*). (*e–h*) Conditional joint p.d.f. of streamwise size and inclination angles (figure 6, with the fitted parameters $\Delta x_0 = 0.2h$ and $\gamma \approx 0.50$ for all downstream inclined structures, and $\Delta x_0 = -3h$ and $\gamma \approx -0.10$ for all upstream inclined structures). (*i–l*) Conditional joint p.d.f. of streamwise size and segment lengths (figure 8*c*,*f*) for strongly linear (top) and weakly linear (bottom) edges. (*m–p*) Conditional joint p.d.f. of UMZ thickness conditioned on y (figure 9*e*).

Appendix B. Percolation analysis of the coherence criteria

Figure 13 shows the percolation analysis of the three parameters in the coherence criteria and the prominence parameter for the p.d.f. analysis. We examined the effect of these parameter choices on the large-scale UMZ interfaces with $\Delta x/h \approx 2$ with respect to three quantities: in the first column, the average y-location of the interfaces, $\overline{y}$; in the second

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

column, the average anti-modal velocity, $\overline{u}$; in the third column, the average streamwise extent of all detected edges, $\overline{\Delta x}$.

Figure 13(*a*,*b*,*c*) shows the effect of the wall-normal distance threshold d_y, normalised by the Taylor microscale, and we observe that for very small d_y, we identify only the UMZ interfaces very close to the wall, and as we increase d_y, the larger interfaces start to be detected in the log layer and above, and the statistics in the first two columns become robust with further increased d_y, with the dashed line representing the current choice, and only the average streamwise extent in panel (*c*) shows monotonical increasing trend with d_y. Panels (*d*,*e*,*f*) show the effect of streamwise distance threshold, normalised by the spatial resolution, where there is relatively little sensitivity to the parameter choice over a wide range. Panels (*g*,*h*,*i*) show the effect of choice of d_u, where again the sensitivity is low. Panels (*j*,*k*,*l*) show the effect of the choice of ξ, where the value is selected close to where the sensitivity levels off.

Figure 14 shows the effect of four different values of the wall-normal distance threshold d_y on maps of linearity (panels *a–d*), inclination angle (panels *e–h*), segment length (panels *i–l*) and UMZ thickness (panels *m–p*). The sensitivity of these maps to the other coherence criteria was also checked, but is not displayed here for brevity. The inclination angle trends and model fit parameters remain the same despite the variation in d_y, thus indicating a reasonable robustness to the moving window algorithm. Only the smallest value of d_y exhibits a noticeable impact on the detection of large-scale edges.

References

Adrian, R.J., Meinhart, C. & Tomkins, C.D. 2000 Vortex organization in the outer region of the turbulent boundary layer. *J. Fluid Mech.* **422**, 1–54.

Baars, W.J., Hutchins, N. & Marusic, I. 2016 Spectral stochastic estimation of high-Reynolds-number wall-bounded turbulence for a refined inner-outer interaction model. *Phys. Rev. Fluids* **1** (5), 1–23.

Bandyopadhyay, P. 1980 Large structure with a characteristic upstream interface in turbulent boundary layers. *Phys. Fluids* **23** (11), 2326–2327.

Cheng, C., Shyy, W. & Fu, L. 2022 Streamwise inclination angle of wall-attached eddies in turbulent channel flows. *J. Fluid Mech.* **946**, 1–18.

Christensen, K.T. & Adrian, R.J. 2001 Statistical evidence of hairpin vortex packets in wall turbulence. *J. Fluid Mech.* **431**, 433–443.

Clauset, A., Shalizi, C.R. & Newman, M.E.J. 2009 Power-law distributions in empirical data. *SIAM Rev.* **51** (4), 661–703.

Cremades, A. *et al.* 2024 Identifying regions of importance in wall-bounded turbulence through explainable deep learning. *Nat. Commun.* **15** (1), 1–12.

Deshpande, R., Monty, J.P. & Marusic, I. 2019 Streamwise inclination angle of large wall-attached structures in turbulent boundary layers. *J. Fluid Mech.* **877**, 1–13.

Ehsani, R., Heisel, M., Li, J., Voller, V., Hong, J. & Guala, M. 2024*a* Stochastic modelling of the instantaneous velocity profile in rough-wall turbulent boundary layers. *J. Fluid Mech.* **979**, A12.

Ehsani, R., Heisel, M., Puccioni, M., Hong, J., Iungo, V., Voller, V. & Guala, M. 2024*b* Stochastic modal velocity field in rough-wall turbulence. *J. Fluid Mech.* **999**, 1–31.

Encinar, M.P. & Jiménez, J. 2023 Identifying causally significant features in three-dimensional isotropic turbulence. *J. Fluid Mech.* **965**, 1–40.

Ester, M., Kriegel, H., Sander, J. & Xu, X. 1996 A density-based algorithm for discovering clusters in large spatial databases with noise. In *KDD'96: Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, pp. 226–231.

Fan, D., Xu, J., Yao, M.X. & Hickey, J. 2019 On the detection of internal interfacial layers in turbulent flows. *J. Fluid Mech.* **872**, 198–217.

Guala, M., Hommema, S.E. & Adrian, R.J. 2006 Large-scale and very-large-scale motions in turbulent pipe flow. *J. Fluid Mech.* **554**, 521–542.

Head, M.R. & Bandyopadhyay, P. 1981 New aspects of turbulent boundary-layer structure. *J. Fluid Mech.* **107**, 297–338.

Hearst, R.J., De Silva, C.M., Dogan, E. & Ganapathisubramani, B. 2021 Uniform-momentum zones in a turbulent boundary layer subjected to freestream turbulence. *J. Fluid Mech.* **915**, 1–28.

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press

HEISEL, M., DASARI, T., LIU, Y., HONG, J., COLETTI, F. & GUALA, M. 2018 The spatial structure of the logarithmic region in very-high-reynolds-number rough wall turbulent boundary layers. *J. Fluid Mech.* **857**, 704–747.

HEISEL, M., DE SILVA, C.M., HUTCHINS, N., MARUSIC, I. & GUALA, M. 2020 On the mixing length eddies and logarithmic mean velocity profile in wall turbulence. *J. Fluid Mech.* **887**, R1.

HEISEL, M., DE SILVA, C.M., HUTCHINS, N., MARUSIC, I. & GUALA, M. 2021 Prograde vortices, internal shear layers and the taylor microscale in high-Reynolds-number turbulent boundary layers. *J. Fluid Mech.* **920**, A52.

HEISEL, M., DE SILVA, C.M., KATUL, G.G. & CHAMECKI, M. 2022 Self-similar geometries within the inertial subrange of scales in boundary layer turbulence. *J. Fluid Mech.* **942**, A33.

HOYAS, S. & JIMÉNEZ, J. 2006 Scaling of the velocity fluctuations in turbulent channels up to $\mathrm{Re}_\tau = 2003$. *Phys. Fluids* **18** (1), 011702.

HU, R., DONG, S. & VINUESA, R. 2023 General attached eddies: scaling laws and cascade self-similarity. *Phys. Rev. Fluids* **8**, 044603.

JIMÉNEZ, J. 2018 Coherent structures in wall-bounded turbulence. *J. Fluid Mech.* **842**, P1.

LASKARI, A., DE SILVA, C.M., HUTCHINS, N. & MCKEON, B.J. 2022 Spatiotemporal characteristics of uniform momentum zones: experiments and modeling. *Phys. Rev. Fluids* **7** (10), 104603.

LASKARI, A., DE KAT, R., HEARST, R.J. & GANAPATHISUBRAMANI, B. 2018 Time evolution of uniform momentum zones in a turbulent boundary layer. *J. Fluid Mech.* **842**, 554–590.

LEE, J., LEE, J.H., CHOI, J. & SUNG, H.J. 2014 Spatial organization of large- and very-large-scale motions in a turbulent channel flow. *J. Fluid Mech.* **749** (3), 818–840.

LEE, M. & MOSER, R.D. 2015 Direct numerical simulation of turbulent channel flow up to. *J. Fluid Mech.* **774**, 395–415.

LIU, Z., ADRIAN, R.J. & HANRATTY, T.J. 2001 Large-scale modes of turbulent channel flow: transport and structure. *J. Fluid Mech.* **448**, 53–80.

LOVEJOY, S., GAONAC'H, H. & SCHERTZER, D. 2004 Bubble distributions and dynamics: the expansion-coalescence equation. *J. Geophys. Res. Solid Earth* **109** (11), 1–16.

MARUSIC, I. & HEUER, W.D.C. 2007 Reynolds number invariance of the structure inclination angle in wall turbulence. *Phys. Rev. Lett.* **99** (11), 3–6.

MEINHART, C.D. & ADRIAN, R.J. 1995 On the existence of uniform momentum zones in a turbulent boundary layer. *Phys. Fluids* **7** (4), 694–696.

ROBINSON, S.K. 1991 Coherent motions in the turbulent boundary layer. *Annu. Rev. Fluid Mech.* **23**, 601.

SALESKY, S.T. & ANDERSON, W. 2020 Revisiting inclination of large-scale motions in unstably stratified channel flow. *J. Fluid Mech.* **884**, R5.

SCOTT, D.W. 2015 *Multivariate Density Estimation: Theory, Practice, and Visualization*. John Wiley & Sons.

DE SILVA, C.M., GNANAMANICKAM, E.P., ATKINSON, C., BUCHMANN, N.A., HUTCHINS, N., SORIA, J. & MARUSIC, I. 2014 High spatial range velocity measurements in a high Reynolds number turbulent boundary layer. *Phys. Fluids* **26** (2), 025117.

DE SILVA, C.M., HUTCHINS, N. & MARUSIC, I. 2016 Uniform momentum zones in turbulent boundary layers. *J. Fluid Mech.* **786**, 309–331.

DE SILVA, C.M., PHILIP, J., HUTCHINS, N. & MARUSIC, I. 2017 Interfaces of uniform momentum zones in turbulent boundary layers. *J. Fluid Mech.* **820**, 451–478.

SILVERMAN, B.W. 2018 *Density Estimation for Statistics and Data Analysis*. Routledge (Taylor & Francis Group).

TANG, Z., FAN, Z., CHEN, L. & JIANG, N. 2021 Outer-layer structure arrangements based on the large-scale zero-crossings at moderate Reynolds number. *Phys. Fluids* **33** (8), 085121.

VESELY, L., HAIGERMOSER, C., GRECO, D. & ONORATO, M. 2009 Turbulent flow and organized motions over a two-dimensional rough wall. *Phys. Fluids* **21** (12), 1–19.

https://doi.org/10.1017/jfm.2025.10929 Published online by Cambridge University Press