



Coherent structure identification in turbulent channel flow using latent Dirichlet allocation

Mohamed Frihat, Bérengère Podvin, Lionel Mathelin, Yann Fraigneau, François Yvon

► To cite this version:

Mohamed Frihat, Bérengère Podvin, Lionel Mathelin, Yann Fraigneau, François Yvon. Coherent structure identification in turbulent channel flow using latent Dirichlet allocation. *Journal of Fluid Mechanics*, 2021, 920, pp.A27. <10.1017/jfm.2021.444>. <hal-03258870>

HAL Id: hal-03258870

<https://hal.science/hal-03258870v1>

Submitted on 5 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Coherent structure identification in turbulent channel flow using Latent Dirichlet Allocation

Mohamed Frihat, Bérengère Podvin, Lionel Mathelin,
Yann Fraigneau, François Yvon

(Received xx; revised xx; accepted xx)

Identification of coherent structures is an essential step to describe and model turbulence generation mechanisms in wall-bounded flows. To this end, we present a clustering method based on Latent Dirichlet Allocation (LDA), a generative probabilistic model for collections of discrete data. The method is applied to a set of snapshots featuring the Reynolds stress (Q_- events) for a turbulent channel flow at a moderate Reynolds number $R_\tau = 590$. Both 2D and 3D analysis show that LDA provides a robust and compact flow description in terms of a combination of *motifs*, which are latent variables inferred from the set of snapshots. We find that the characteristics of the motifs scale with the wall distance, in agreement with the wall-attached eddy hypothesis of [Townsend \(1961\)](#). LDA motifs can be used to reconstruct fields with an efficiency that can be compared with the POD. Moreover, the LDA model makes it possible to generate a collection of synthetic fields that is statistically closer to the original dataset than its POD-generated counterpart. These findings highlight the potential of LDA for turbulent flow analysis, compression and generation.

1. Introduction

The introduction of coherent structures ([Kline *et al.* 1967](#); [Townsend 1947](#)) has represented a major paradigm shift for turbulence theory and has had a significant impact in various related fields, ranging from geophysical flows to industrial applications. Coherent structure identification has become a key step towards modelling and controlling wall-bounded turbulent flows. However a recurrent stumbling block is the absence of a precise definition of structures, as is apparent from several comprehensive reviews ([Cantwell 1981](#); [Robinson 1991](#); [Jimenez 2013](#); [Dennis 2015](#)).

Studies originating in the 1960's ([Kline *et al.* 1967](#); [Kim *et al.* 1971](#)) have established that most of the turbulence in the near-wall region occurred in a highly intermittent manner in both space and time, during what was originally termed “bursting events”. Quadrant analysis of the Reynolds stress in the plane of streamwise and wall-normal fluctuation (u', v') was introduced by [Wallace *et al.* \(1972\)](#); [Willmarth & Lu \(1972\)](#) to characterize these events. Bursting events were found to be associated with low-speed streaks being lifted away from the wall, as well with sweeping motions of high-speed fluid towards the wall, which respectively correspond to Quadrant II ($u' < 0, v' > 0$) and Quadrant IV ($u' > 0, v' < 0$) events. The two quadrants corresponding to $-u'v' > 0$ can be termed Q_- events and represent the major contribution to the Reynolds stress ([Wallace 2016](#)). An interpretation of these bursts is that they are the signature of coherent structures or eddies advected by the mean field. Determining the characteristics of these structures has been the object of considerable effort, [Jimenez \(2018\)](#).

A central element of wall turbulence theory is the attached eddy model, reviewed in

detail by [Marusic & Monty \(2019\)](#). The model is based on the idea that turbulence arises as a field of randomly distributed eddies, identified as organized flow patterns which extend to the wall, in the sense that their characteristics are influenced by the wall. Further assumptions require that the entire geometry of the eddies scales with the wall distance, with a constant characteristic velocity scale. The model was extended by [Perry & Chong \(1982\)](#), who introduced the idea of a hierarchy of discrete scales, with an inverse-scale probability distribution. [Woodcock & Marusic \(2015\)](#) showed that this inverse probability distribution was in fact a direct consequence of the self-similarity of the eddies. Further extensions of the model for the logarithmic layer include a wider variety of structures, such as wall-detached ones ([Perry & Marusic 1995](#); [Hu *et al.* 2020](#)).

Detection of self-similarity in boundary layers has been the focus of several experimental studies, such as [Baars *et al.* \(2017\)](#)'s, who used spectral coherence analysis to provide evidence of self-similar structures in the streamwise velocity fluctuations of pipe flow. Numerical simulation has proved a powerful tool to explore three-dimensional flow fields using a clustering approach. Examples include the work of [Alamo *et al.* \(2006\)](#), who showed that the logarithmic region of turbulent channel was organized in self-similar vortex clusters, and [Lozano-Duran *et al.* \(2012\)](#) developed a three-dimensional extension of quadrant analysis to detect self-similarity in numerical data at various Reynolds numbers. More recently, wall-attached structures were identified in the streamwise fluctuations of a turbulent boundary layer ([Hwang & Sung 2018](#)) as well as in pipe flow ([Hwang & Sung 2019](#)). The structures were shown to scale with the wall distance while their population density scales inversely with the distance to the wall. [Cheng *et al.* \(2020\)](#) detected the signature of wall-attached eddies in the streamwise and spanwise velocity fluctuations in turbulent channel flow simulations at low Reynolds numbers. Evidence of self-similarity has been found as well in the context of resolvent analysis, [Sharma & McKeon \(2013\)](#). It has also emerged from Proper Orthogonal Decomposition (POD) results, such as channel flow simulations at low Reynolds numbers ([Podvin *et al.* 2010](#); [Podvin & Fraigneau 2017](#)), or pipe flow experiments ([Hellström *et al.* 2016](#)).

The increase of available data, whether through numerical simulation or experiment, has strengthened the need for new identification methods, such as those provided by machine learning (see [Brunton *et al.* \(2020\)](#) for a review). The challenge is to extract structural information about the data without pre-existing knowledge, which defines an *unsupervised learning* problem. Solutions to this problem should be robust, easy to implement and scalable. One example of unsupervised learning method that meets these criteria is Proper Orthogonal Decomposition ([Lumley 1967](#)), a now classical approach to decompose turbulent fields. POD is a statistical technique which provides an objective representation of the data as a linear combination of spatial eigenfunctions, which can be hierarchized with respect to a given norm. Although the reconstruction is optimal with respect to this norm ([Holmes *et al.* 1996](#)), a potential limitation of the decomposition is that the physical interpretation of the eigenfunctions is not clear. In particular, in the case of homogeneous statistics, the eigenfunctions are spatial Fourier modes over the full domain (see [Holmes *et al.* \(1996\)](#) for a proof), even though instantaneous patterns are strongly localized in space. The connection between POD spatial eigenfunctions with observed coherent structures is therefore not necessarily straightforward. Moreover, the amplitudes of the spatial eigenfunctions are generally strongly inter-dependent, even though they are by construction uncorrelated. This makes it difficult to give a physical meaning to individual amplitudes, especially in the absence of a probabilistic framework in which to interpret them.

In this paper we consider such a framework to explore an alternative unsupervised learning approach called Latent Dirichlet Allocation (LDA), which can be derived from

POD (Hofmann 1999). LDA is a generative probabilistic model, that is a probabilistic model that mimics the characteristics of a collection of data. It is based on a soft clustering approach, which was first developed for text mining applications (Blei *et al.* 2003), but has been extended to other fields in recent years (Aubert *et al.* 2013). The goal of LDA (Blei *et al.* 2003) is to find short descriptions of the members of a collection that enable efficient processing of large collections while preserving the essential statistical relationships that are useful for basic tasks such as classification, novelty detection, summarization, and similarity and relevance. LDA is a three-level hierarchical Bayesian model, in which each member of a collection is modeled as a finite mixture over an underlying set of topics or motifs.

In the field of natural language processing, the dataset to which LDA is applied consists of a set of documents, each of which is considered as a “bag-of-words”, that is an unordered set of words taken from a finite vocabulary. A particular word may appear several times in the document, or not appear at all. The number of occurrences of each vocabulary word in a document can be seen as an entry of a sparse matrix where the lines correspond to the vocabulary words and the columns to the documents. Based on this typically sparse word count matrix, the classification method returns a set of N_T topics, where the topics are latent variables inferred from the word counts in the documents and the number of topics N_T is a user-defined parameter.

Unlike “hard” clustering, such as the K-means approach (MacQueen 1967), where each document is assigned to a specific topic, LDA represents each document as a mixture of topics, where the coefficients of the mixture represent the probability of the topic in the document. An interesting application of the LDA method was carried out for a dataset containing images by Griffiths & Steyvers (2004). The dataset considered was a collection of gray-scale images where each image consists of an array of pixels, each of which is associated with a gray level. In this framework, each image is the equivalent of a document, each pixel represents an individual vocabulary word, and the gray-level intensity measured at each pixel is taken as the analog of the word count matrix entry (the lines of the matrix now represent the pixels, while the columns represent the snapshots). The sum of the intensities over the pixels, which will be called throughout the paper the *total intensity*, is the analog of the total number of words observed in the document. Given a set of original patterns constituting the topics or *motifs*, a collection of synthetic images was generated from random mixtures of the patterns. It was shown that LDA was able to recover the underlying patterns from the observations of the generated images.

Following Griffiths & Steyvers (2004), the idea of the paper is to look for evidence of coherent structure in turbulent flow snapshots by identifying LDA topics or *motifs*. The relevant gray-level intensity is based on the value of Q_- (unlike in Griffiths & Steyvers (2004)’s work, it corresponds to a physical field.) We thus propose the following analogy: each scalar field observed in a collection of snapshots results from a mixture of N_T spatial topics that will be referred to as *motifs* in the remainder of the paper. This can be compared with the standard view that each realization of a turbulent flow is constituted of a random superposition of discrete eddies, characterized by a hierarchy of scales.

The paper is organized as follows. We show in Section 2 how the POD method of snapshots, which is equivalent to Latent Semantic Allocation (LSA), can be generalized to a probabilistic framework (Probabilistic Latent Semantic Allocation or PLSA) which is then further extended into Latent Dirichlet Allocation (LDA) in Section 3. Application to the extraction of motifs for a turbulent channel flow is introduced in Section 4 and results are discussed in Section 5. The potential of the approach for flow reconstruction and flow generation is considered in Section 6 before Section 7 closes the paper.

2. A probabilistic extension of Proper Orthogonal Decomposition

To suitably introduce and contextualize the Latent Dirichlet Allocation, several established approaches to represent data are first briefly discussed.

2.1. Proper Orthogonal Decomposition

2.1.1. General formulation

The Proper Orthogonal Decomposition (POD) is arguably the most popular tool for representation and analysis of turbulent flow fields. It relies on a method rediscovered and revisited several times in different scientific domains and comes by several names (Principal Component Analysis, Empirical Mode Decomposition, Karhunen-Loève decomposition, Latent Semantic Allocation (LSA) ...) although they are not all strictly equivalent. It was introduced for turbulent flows and adapted by Lumley (1967).

The POD method allows to derive an orthogonal basis for the (sub)space of the fluctuations of a multi-dimensional quantity \mathbf{f} of finite variance. One can show that a basis for the space of fluctuations, defined as $\mathbf{f}'(t) := \mathbf{f}(t) - \langle \mathbf{f} \rangle$, with $\langle \cdot \rangle$ the statistical mean, is given by the set of elements $\{\phi_n\}_n$, eigenvectors of the following eigenvalue problem (Holmes *et al.* 1996):

$$C\phi_n = \lambda_n\phi_n, \quad (2.1)$$

with λ_n the eigenvalue and $C \in \mathbb{R}^{N_x \times N_x}$ the empirical 2-point covariance matrix:

$$C = \frac{1}{N_s} \sum_{i=1}^{N_s} \mathbf{f}'(t_i) \mathbf{f}'(t_i), \quad (2.2)$$

with $\{t_i\}_i$ the time instants for which the field \mathbf{f} is available. Some conditions on the temporal sampling scheme apply for the empirical covariance \hat{C} to be an accurate approximation of C (Holmes *et al.* 1996). POD modes are identified as the eigenvectors ϕ_n .

2.1.2. Method of snapshots

The above method is a quite natural implementation of the underlying Hilbert-Schmidt decomposition theory. However, the algorithmic complexity associated with the eigenvalue problem (2.1) scales as $\mathcal{O}(N_s N_x^2)$, where the number of field instances N_s was assumed to be lower than the size N_x of the discrete field, $N_s \leq N_x$. For large field vectors (large N_x), the computational and memory cost is hence high. For this widely encountered situation, a possible workaround was suggested in Sirovich (1987) and consists in solving the following eigenvalue problem:

$$\tilde{C} \mathbf{a}_n = \lambda_n \mathbf{a}_n, \quad \mathbf{a}_n \in \mathbb{R}^{N_s}, \quad (2.3)$$

with

$$\tilde{C}_{i,i'} \propto \langle \mathbf{f}'(t_i), \mathbf{f}'(t_{i'}) \rangle_{\Omega}, \quad \forall i, i' \in [1, N_s] \subset \mathbb{N}, \quad (2.4)$$

and $\langle \cdot, \cdot \rangle_{\Omega}$ the Euclidean inner product. Since the correlation matrix \tilde{C} is Hermitian, its eigenvalues are real and non-negative, $\lambda_n \geq 0$, $\forall n$, and its eigenvectors $\{\mathbf{a}_n\}_n$ are orthogonal and can be made orthonormal in an Euclidean sense, $\mathbf{a}_n^T \mathbf{a}_{n'} \propto \delta_{n,n'}$, with δ the Kronecker delta. The spatial POD modes are finally retrieved via projection as follows:

$$\phi_n = \lambda_n^{-1/2} F' \mathbf{a}_n, \quad \forall n. \quad (2.5)$$

where the i -th column of the matrix F' is the snapshot \mathbf{f}'_i .

The algorithmic complexity is now $\mathcal{O}(N_s^3)$ and scales much better than the standard

POD approach ($\mathcal{O}(N_s N_x^2)$) in the usual situation where $N_s \ll N_x$. In this work, we rely on this so-called method of snapshots to implement POD.

Formally the decomposition of the snapshot matrix F' is equivalent to a singular value decomposition SVD

$$F' = \Phi \Sigma A^\top, \quad (2.6)$$

where Φ is the matrix constituted by the n columns ϕ_n ,

A is the matrix containing the n columns \mathbf{a}_n and Σ is a diagonal matrix whose entries are $\lambda_n^{-1/2}$. The snapshot matrix can thus be decomposed into a snapshot-mode matrix A and into a cell-mode matrix Φ . The spatial modes or *structures* can be seen as latent variables allowing optimal reconstruction of the data in the L_2 norm or an equivalent. The decomposition can be truncated to retain only the N_T largest values corresponding to the N_T first columns of each matrix.

2.2. Probabilistic Latent Semantic Analysis

In all that follows we will consider a collection of N_s scalar fields $\{\mathbf{f}_i\}_{i=1,\dots,N_s}$. Each field is of dimension N_x and consists of either positive or zero integer values on each grid cell. For each snapshot i , the value of \mathbf{f}_i on grid cell l indicates that the grid cell i has been detected or activated $f_{l,i}$ times. Probabilistic Latent Semantic Analysis (PLSA) tackles the problem of finding latent variables using a probabilistic method instead of SVD. This representation assumes that each snapshot \mathbf{f}_i consists of a mixture of structures \mathbf{z}_n .

PLSA adds a probabilistic flavor as follows:

- given a snapshot \mathbf{f}_i , the structure \mathbf{z}_n is present in that snapshot with probability $p(\mathbf{z}_n|\mathbf{f}_i)$,
- given a structure \mathbf{z}_n , the grid cell \mathbf{x}_l is activated with probability $p(\mathbf{x}_l|\mathbf{z}_n)$.

Formally, the joint probability of seeing a given snapshot \mathbf{f}_i and activating a grid cell \mathbf{x}_l is:

$$p(\mathbf{f}_i, \mathbf{x}_l) = p(\mathbf{f}_i) \sum_n p(\mathbf{z}_n|\mathbf{f}_i) p(\mathbf{x}_l|\mathbf{z}_n). \quad (2.7)$$

$p(\mathbf{f}_i)$, $p(\mathbf{z}_n|\mathbf{f}_i)$, and $p(\mathbf{x}_l|\mathbf{f}_i)$ are the parameters of the model: $p(\mathbf{f}_i)$ is the probability to obtain such a snapshot \mathbf{f}_i and is constant in our case, $p(\mathbf{f}_i) = 1/N_s$. $p(\mathbf{z}_n|\mathbf{f}_i)$ and $p(\mathbf{x}_l|\mathbf{z}_n)$ can be inferred using the Expectation-Maximization (EM) algorithm of [Dempster et al. \(1977\)](#).

Using Bayes' rule, $p(\mathbf{f}_i, \mathbf{x}_l)$ can be equivalently written as:

$$p(\mathbf{f}_i, \mathbf{x}_l) = \sum_n p(\mathbf{z}_n) p(\mathbf{x}_l|\mathbf{z}_n) p(\mathbf{f}_i|\mathbf{z}_n). \quad (2.8)$$

This alternative formulation shows a direct link between PLSA model and POD model (as mentioned above, POD is called Latent Semantic Allocation or LSA in text mining). If we compare equations (2.6) and (2.8), we see that the structure probability $p(\mathbf{z}_n)$ corresponds to the diagonal matrix Λ_n , the probability of the snapshot \mathbf{f}_i given the structure \mathbf{z}_n corresponds to the snapshot-mode matrix entry $A_{i,n}$, and the probability to activate the cell \mathbf{x}_l given the structure \mathbf{z}_n corresponds to the matrix entry $\Phi_{l,n}$.

3. Latent Dirichlet Allocation

Latent Dirichlet Allocation (LDA) extends PLSA to address its limitations. Its specificity is:

- the introduction of a probabilistic model for the collection of snapshots: each snap-

shot is now characterized by a distribution over the structures which will be now called *motifs*.

- the use of Dirichlet distributions to model both motif-cell and snapshot-motif distributions.

The Dirichlet distribution is a multivariate probability distribution over the space of multinomial distributions. It is parametrized by a vector of positive-valued parameters $\alpha = (\alpha_1, \dots, \alpha_N)$:

$$p(x_1, \dots, x_N; \alpha_1, \dots, \alpha_N) = \frac{1}{B(\alpha)} \prod_{n=1}^N x_n^{\alpha_n - 1},$$

where B is a normalizing constant, which can be expressed in terms of the Gamma function Γ :

$$B(\alpha) = \frac{\prod_{n=1}^N \Gamma(\alpha_n)}{\Gamma(\sum_{n=1}^N \alpha_n)}.$$

The support of the Dirichlet distribution is the set of N -dimensional discrete distributions, which constitutes the $N - 1$ simplex. Introduction of the Dirichlet distribution allows us to specify the prior belief about the snapshots. The Bayesian learning problem is now to estimate $p(\mathbf{z}_n, \mathbf{f}_i)$ and $p(\mathbf{x}_l, \mathbf{z}_n)$ from F given our prior belief α , and it can be shown that Dirichlet distributions offer a tractable, well-posed solution to this problem (Blei *et al.* 2003).

LDA is therefore based on the following representation:

(i) Each motif \mathbf{z}_n is associated with a multinomial distribution φ_n over the grid cells ($p(\mathbf{x}_l | \mathbf{z}_n) = \varphi_{l,n}$). This distribution is modeled with a Dirichlet prior parametrized with a N_x -dimensional vector β . The components β_l of β control the sparsity of the distribution: values of β_l larger than 1 correspond to evenly dense distributions, while values lower than 1 correspond to sparse distributions. In all that follows, we will assume a non-informative prior, meaning that $\beta = \beta \mathbf{1}_{N_x}$.

(ii) Each snapshot \mathbf{f}_i , is associated with a distribution of motifs θ_i such that $\theta_{n,i} = p(\mathbf{z}_n | \mathbf{f}_i)$. The probabilities of each motif add up to 1 in each snapshot. This distribution is modelled with a N_T -dimensional Dirichlet distribution of parameter α . The magnitude of α characterizes the sparsity of the distribution (low values of α_n correspond to snapshots with relatively few motifs). The same assumption of a non-informative prior leads us to assume $\alpha = \alpha \mathbf{1}_{N_T}$.

The generative process performed by LDA with N_T motifs is the following:

- (i) For each motif \mathbf{z}_n , a cell-motif distribution φ_n is drawn from the Dirichlet distribution of parameter β .
- (ii) For each snapshot \mathbf{f}_i :
 - a snapshot-motif distribution θ_i is drawn.
 - each intensity unit $1 \leq j \leq N_i$ where N_i is the total intensity with $N_i = \sum_l f_{l,i}$ is then distributed among the different cells as follows:
 - a motif \mathbf{z}_n is first selected from θ_i (motif \mathbf{z}_n occurs with probability $\theta_{n,i}$ in the snapshot),
 - for this motif, a cell l is chosen among the cells using $\varphi_{l,n}$ and the intensity associated with cell l is incremented by 1.

The generative process can be summarized as follows:

Algorithm 1: LDA Generative Model.

```

for each of the  $N_T$  motifs  $n$  do
  | sample  $\varphi_n \sim \text{Dir}(\beta)$ 
end
for each of the  $N_s$  snapshots  $i$  do
  | sample  $\theta_i \sim \text{Dir}(\alpha)$ 
  | for each of the  $N_i$  intensity units do
  |   | 1. sample a motif  $z_n$  from  $\theta_{n,i}$ 
  |   | 2. for this motif sample a cell  $l$  from  $\varphi_{l,n}$ 
  | end
end

```

The snapshot-motif distribution θ_i and the cell-motif distribution φ_n are determined from the observed f_i . They are respectively N_T - and N_x -dimensional categorical distributions. Finding the distributions θ_i and φ_n that are most compatible with the observations is an inference problem that can be solved by either a variational formulation (Blei *et al.* 2003) or a Gibbs sampler (Griffiths & Steyvers 2002). In the variational approach, the objective function to minimize is the Kullback-Leibler divergence. The solution *a priori* depends on the number of motifs and on the values of the Dirichlet parameters α and β .

We conclude this section with two remarks.

(i) LDA can generalize to new snapshots more easily than PLSA, due to the snapshot-motif distribution formalism. In PLSA, the snapshot probability is a fixed point in the dataset, which cannot be estimated directly if it is missing. In LDA, the dataset serves as training data for the Dirichlet distribution of snapshot-motif distributions. If a snapshot is missing, it can easily be sampled from the Dirichlet distribution instead.

(ii) An alternative viewpoint can also be adopted in interpreting the LDA in the form of a regularized matrix factorization method. This is further discussed in Appendix A.

4. Application of LDA to turbulent flows

4.1. Numerical configuration

The idea of this paper is to apply this methodology to snapshots of turbulent flows in order to determine latent motifs from observations of Q_- events. We will consider the configuration of turbulent channel flow at a moderate Reynolds number of $R_\tau = u_\tau h / \nu = 590$ (Moser *et al.* 1999; Muralidhar *et al.* 2019), where R_τ is the Reynolds number based on the fluid viscosity ν , channel half-height h and friction velocity u_τ . Wall units based on the friction velocity and fluid viscosity will be denoted with a subscript $+$. The streamwise, wall-normal and spanwise directions will be referred to as x, y and z respectively. The horizontal dimensions of the numerical domain are $(\pi, \pi/2)h$. Periodic boundary conditions are used in the horizontal directions. The resolution of $(256)^3$ points is based on a regular spacing in the horizontal directions and a hyperbolic tangent stretching function for the vertical direction. The configuration is shown in Figure 1. More details about the numerical simulation can be found in Muralidhar *et al.* (2019).

4.2. LDA inputs

In this section, we introduce the different parameters of the study. The python library `scikit-learn` (Pedregosa *et al.* 2011) was used to implement LDA. The sensitivity of the results to these parameters will be examined in a subsequent section.

We first focus on 2-D vertical subsections of the domain, then present 3-D results. The vertical extent of the domain of investigation was the half-channel height. Since this is an exploration into a new technique, a limited range of scales was considered in the horizontal dimensions: the spanwise dimension of the domain was limited to 450 wall units. The streamwise extent of the domain was in the range of 450-900 wall units. The number of realizations considered for 2-D analysis was $N_s = 800$, with a time separation of 60 wall time units. The number of snapshots was increased to 2400 for 3-D analysis.

The scalar field \mathbf{f} of interest corresponds to Q_- events. It is defined as the positive part of the product $-u'v'$, where fluctuations are defined with respect to an average taken over all snapshots and horizontal planes. The LDA procedure requires that the input field consists of integer values: it was therefore rescaled and digitized and the scalar field f was defined as:

$$f = [A\tau_-],$$

where $\tau_- = \max(-u'v', 0)$ and $[\cdot]$ represents the integer part. The rescaling factor A was chosen in order to yield a sufficiently large, yet still tractable, total intensity. In practice we used $A = 40$, which led to a total intensity $\sum_i \sum_l f_{l,i}$ of about 10^8 for plane sections. The effect of the rescaling factor will be examined in a subsequent section.

LDA is characterized by a user-defined number of motifs N_T , a parameter α which characterizes the sparsity of prior Dirichlet snapshot-motif distribution, and a parameter β which characterizes the sparsity of the prior Dirichlet motif-cell distribution. Results were obtained assuming uniform priors for α and β with a default value of $1/N_T$. The sensitivity of the results to the priors will be evaluated in Section 5.2.

4.3. LDA outputs

For a collection of N_s snapshots and a user-defined number of motifs N_T , LDA returns N_T motif-cell distributions φ_n and N_s snapshot-motif distributions θ_i . Each *motif* is defined by a probability distribution φ_n associated with each grid cell. It is therefore analogous to a structure or a portion of structure since it contains spatial information - note however that its definition is different from standard approaches. The motif-snapshot distribution θ_i characterizes the prevalence of a given motif in the snapshot.

As will be made clear below, the motifs most often consist of single connected regions, although occasionally a couple of distinct regions were identified. In most cases, the motifs can thus be characterized by a characteristic location \mathbf{x}^c and a characteristic dimension in each direction L_j , $j \in \{x, y, z\}$.

To determine these characteristics, we first define for each motif a mask associated with a domain D . The origin of the domain was defined as the position \mathbf{x}_m corresponding to its maximum probability $p_m = \varphi_n(\mathbf{x}_m)$. The dimensions of the domain in each direction (for instance L_x) were defined as the segment extending from the domain origin over which the probability remained larger than 1% of its maximum value p_m . The position and characteristic dimension of a motif for instance in the x -direction are then defined as:

$$x^c = \frac{\int_D x \varphi_n dD}{\int_D \varphi_n dD}, \quad (4.1)$$

$$L_x^2 = 2 \frac{\int_D (x - x^c)^2 \varphi_n dD}{\int_D \varphi_n dD}. \quad (4.2)$$

Analogous definitions can be given for y^c and z^c .




Figure 1: Numerical domain D . The shaded surfaces correspond to the two types of planes used in the analysis. The volume considered for 3D analysis is indicated in bold lines.

5. Results

5.1. Vertical planes

In order to investigate in detail the vertical organization of the flow, LDA was first applied to vertical sections of the flow. Both cross-flow (y, z) and longitudinal (x, y) sections were considered. Due to the horizontal homogeneity of the flow, we do not expect significant changes in the cell-motif and the motif-document distributions when the sections are translated in the horizontal direction.

5.1.1. Cross-flow planes

The dimensions of the cross-sections were $d_{z+} = 450$ and $d_{y+} = 590$. Figure 2 shows selected motifs for a total number of motifs $N_T = 96$ on a vertical plane at $x = 0$. The motifs consist of isolated regions, the dimensions of which increase with the wall distance. This is confirmed by Figure 3, which represents characteristic sizes of LDA motifs of a succession of four vertical planes separated by a distance of 100 wall units (+). We point out that observing motifs which are detached from the wall does not inform the presence of wall-attached structures, as they would be consistent with a cross-section of a wall-attached structure elongated in the streamwise direction. Results for several motif numbers (three different motif numbers $N_T = 48, 96, 144$ are shown in Figure 3), it was found that both spanwise and vertical dimensions increase linearly with the wall distance in the region $y_+ > 100$. Again, this is in agreement with [Townsend \(1961\)](#)'s hypothesis of a hierarchy of structures of increasing dimensions, which was also confirmed numerically by [Flores & Jimenez \(2010\)](#).

The aspect ratio L_z/L_y is constant with the wall distance above $y_+ > 100$, with a typical value of about 1. We note that [Lozano-Duran et al. \(2012\)](#) found with a different definition that Q_- events were characterized by nearly equal spanwise and vertical sizes $\Delta z \sim \Delta y$, while [Alamo et al. \(2006\)](#) found a scaling of $\Delta z \sim 1.5\Delta y$ for vortex clusters.

Figure 4 (left) shows the distribution of the vertical location $p(y_m)$ of the motif maximum probability. Comparison of two different plane locations x confirms that results do not depend on the location of the plane, which reflects the statistical homogeneity of the flow in the horizontal direction. The probability decreases as the inverse of the wall distance on all planes. This is in agreement with Townsend's self-similarity hypothesis that the number of structures decreases with the wall distance in $1/y$ ([Townsend 1961](#);

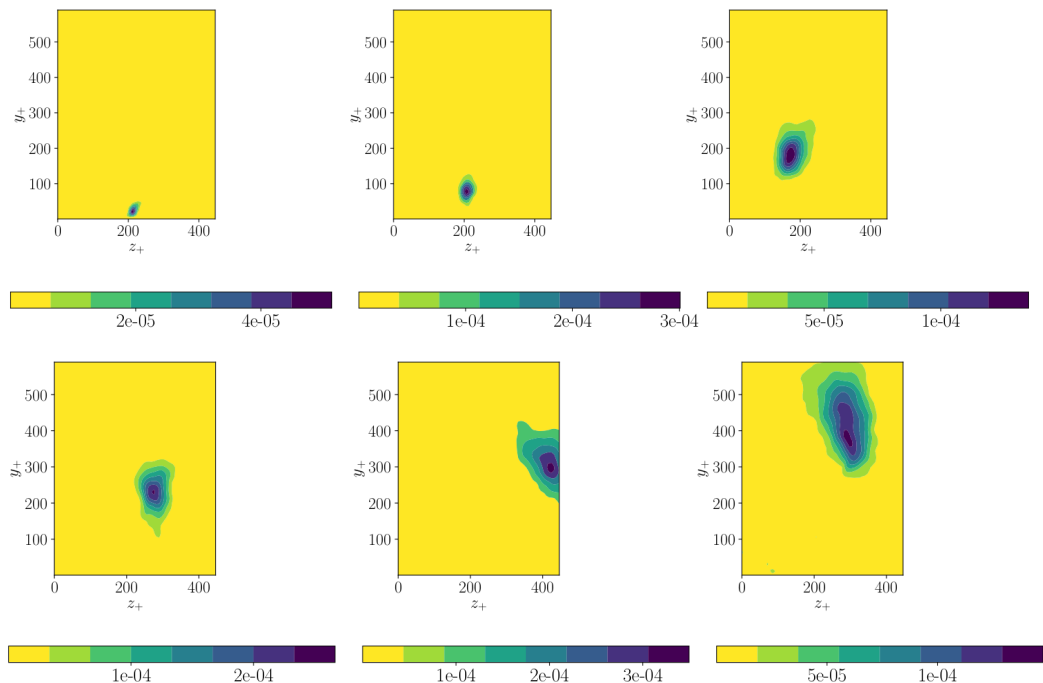


Figure 2: Selected motifs in a cross-flow plane for a number of motifs $N_T = 96$.

Woodcock & Marusic 2015). Figure 4 (right) shows that a good fit is $p(y) \simeq \frac{c}{y} - \gamma$, with $\gamma = 0.0006$ and $c = 0.4$.

5.1.2. Longitudinal planes

We now examine results for the longitudinal sections (x, y) . The streamwise and vertical dimensions of the sections are respectively $d_{x+} = 900$ and $d_{y+} = 590$ wall units, although some tests were also carried out for a streamwise extent of 450 units. Figure 5 presents selected motifs for the longitudinal planes for $N_T = 96$. As in the cross-flow plane, the dimensions of the motifs increase with the wall distance, which is confirmed by Figure 6. The characteristic dimensions seem essentially independent of the total number of motifs (see also next section). There is a wide disparity in streamwise characteristic dimensions near the wall. The motif aspect ratio is highest near the wall and decreases sharply in the region $0 < y_+ < 50$. The vertical dimension increases linearly with the wall distance in the region $y_+ > 100$, as well as the streamwise dimension, with an aspect ratio of L_x/L_y on the order of 2.

Figure 7 shows the distribution of the motif maximum probability location for two different sets $N_T = 48, 96$, and for two domain lengths. The shape of the distribution does not appear to change, and again fits well with the distribution $p \simeq \frac{c}{y} - \gamma$ with $c = 0.4$ and $\gamma = -0.0006$ (Figure 7 right).

5.2. Sensitivity of the results

In this section we examine if and how the characteristics of the motifs depend on the various parameters of LDA. We point out that the probabilistic framework of the model makes exact comparison difficult, since there is no convergence in the L_2 sense, and the

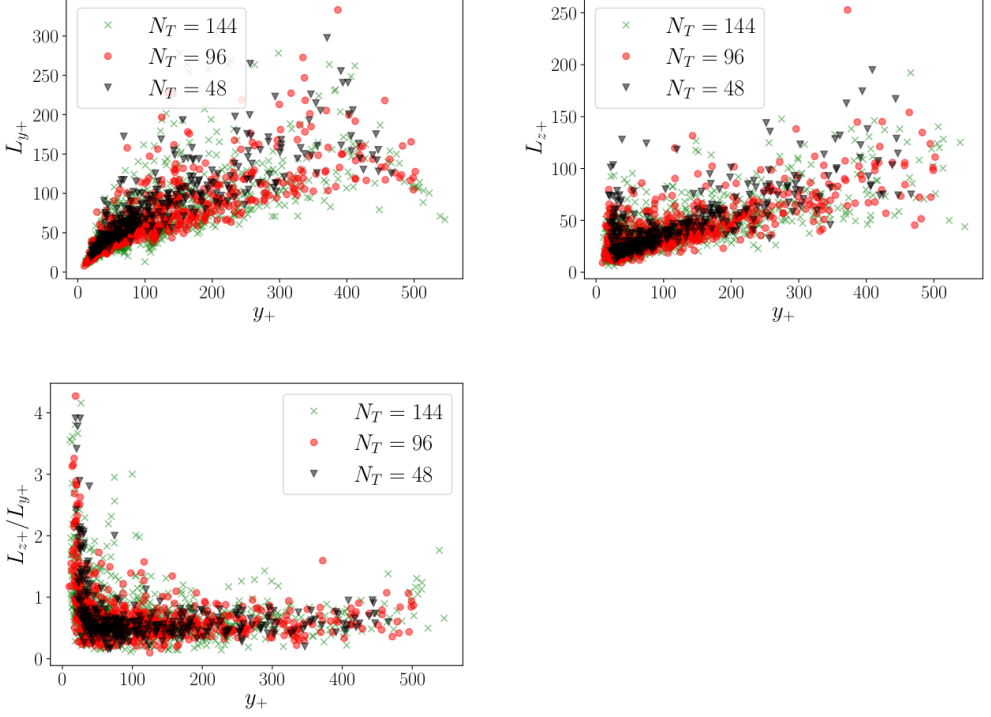


Figure 3: Cross-plane motif characteristic sizes; Left: Vertical dimension L_y ; Right : Spanwise dimension L_z ; Bottom: Aspect ratio L_y/L_z . Each dot corresponds to a motif.

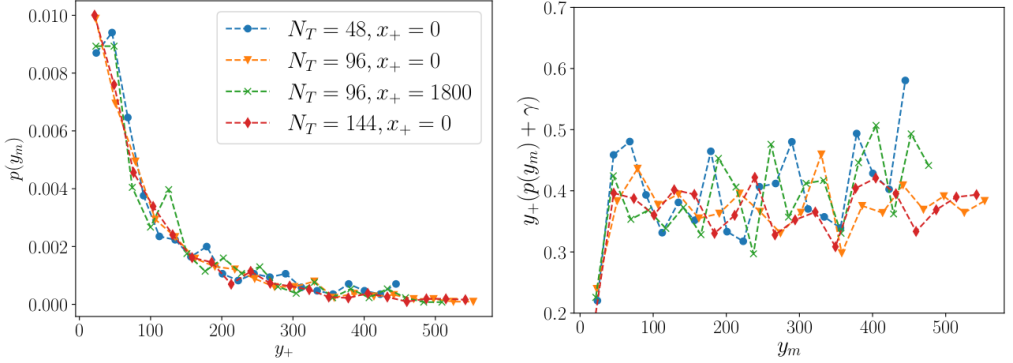


Figure 4: Left: Distribution of the motif maximum location y^c ; Right: Compensated plot of the distribution for different sets of motifs and different subdomains. The legend is the same for the two figures.

Kullback-Leibler divergence, which measures the difference between two distributions is not a true metric tensor (see Appendix).

The criteria we chose to assess the robustness of the results were the characteristic size of the topics and the distribution of their locations. We first examine the influence of various LDA parameters on the results obtained for cross-flow sections for a constant




Figure 5: Selected motifs for a longitudinal plane with $N_T = 96$ motifs.

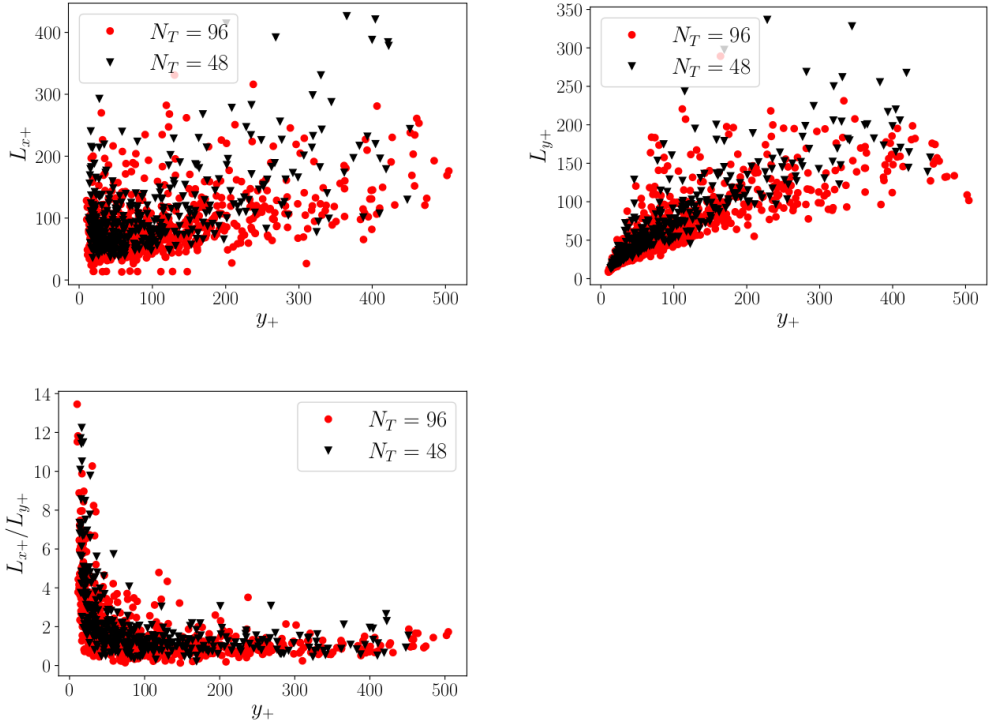


Figure 6: Longitudinal motif characteristic dimensions; Left: Streamwise dimension L_x ; Right: Vertical dimension L_y ; Bottom : Aspect Ratio L_x/L_y . Each dot corresponds to a motif.

number of topics $N_T = 48$. The reference case corresponded to an amplitude $A = 40$, prior values of $\alpha = \beta = 1/N_T$ and a total number of snapshots $N_s = 800$.

Figure 8 (top row) shows that the characteristic dimension is not modified when the number of snapshots was reduced by 50%, indicating that the procedure has converged. Figure 8 (bottom row) shows the characteristic vertical dimension L_y of the structures when the rescaling parameter A was varied. Similar results (not shown) were found for L_z .

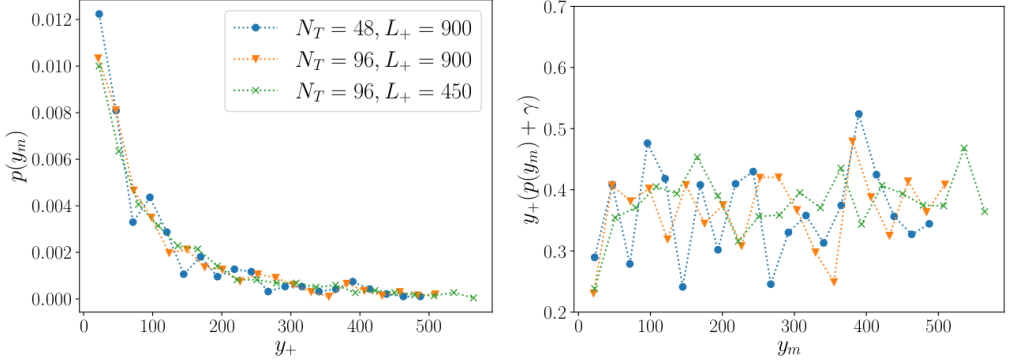


Figure 7: Left: Histogram of the motif location y_m^c ; Right: Compensated plot of the histogram for different sets of motifs and different subdomains. The legend is the same for the two figures.

Although some fluctuations were observed in the individual characteristic dimensions, no significant statistical change was observed. Figure 9 shows the characteristic dimensions of the structures for different prior choices for α and β , which govern the sparsity of the representation. No significant statistical trend was modified when α and β were made to vary within $1/10$ and up to 10 times their default values of $1/N_T$. Figure 10 shows that the distribution of the maximum location of the motifs follows the same inverse law and does not depend on the choice of parameters chosen for LDA.

We now study the sensitivity of the motifs to the choice of N_T for both types of vertical planes. We have seen in the previous sections that the motif dimensions appear essentially independent of the number of motifs considered. To quantify this more precisely, we first define a characteristic motif size L_T as $L_T = \sqrt{\langle A_T \rangle}$ where A_T is the area corresponding to the ellipsoid with the same characteristic dimensions as the motif and $\langle \cdot \rangle$ represents the average over the motifs. Figure 11 summarizes how the motif size evolves with the number of motifs for both vertical and longitudinal planes. In all cases, it was found that the characteristic size varies slowly around a minimal value (Figure 11, left), and that the characteristic area of the motif was minimum when the sum of the motif characteristic areas N_TA_T was comparable with the total domain area A_D (Figure 11, right).

5.3. 3-D Analysis

LDA was then applied to a volumic section of the flow of size $450 \times 590 \times 450$ wall units. Figure 12 shows the cross-sections views of three 3-D motifs. One can note the streamwise coherence of the topics over different heights. We note that the small dimensions of the volume may make it difficult to capture full-length structures, even at this comparatively low Reynolds number, and results should be confirmed by a more extensive investigation which is outside the scope of this paper.

The characteristic dimensions of the motifs are reported in Figure 13. Two different regions can be identified. For $y_+ < 100$ the region is characterized by a wide distribution of L_x , with large values that can extend over the whole domain. Some relatively large values of L_z can occasionally be observed. For $y_+ > 100$ values of L_x are lower and L_z grows linearly. L_y appears to grow linearly over both regions.

The ratio between the horizontal dimensions L_x and L_z is reported in Figure 13 (right). We can see that the streamwise to spanwise aspect ratio decreases over $0 < y_+ < 100$ from an average value of 5 at the wall, which corresponds to the typical aspect ratio of

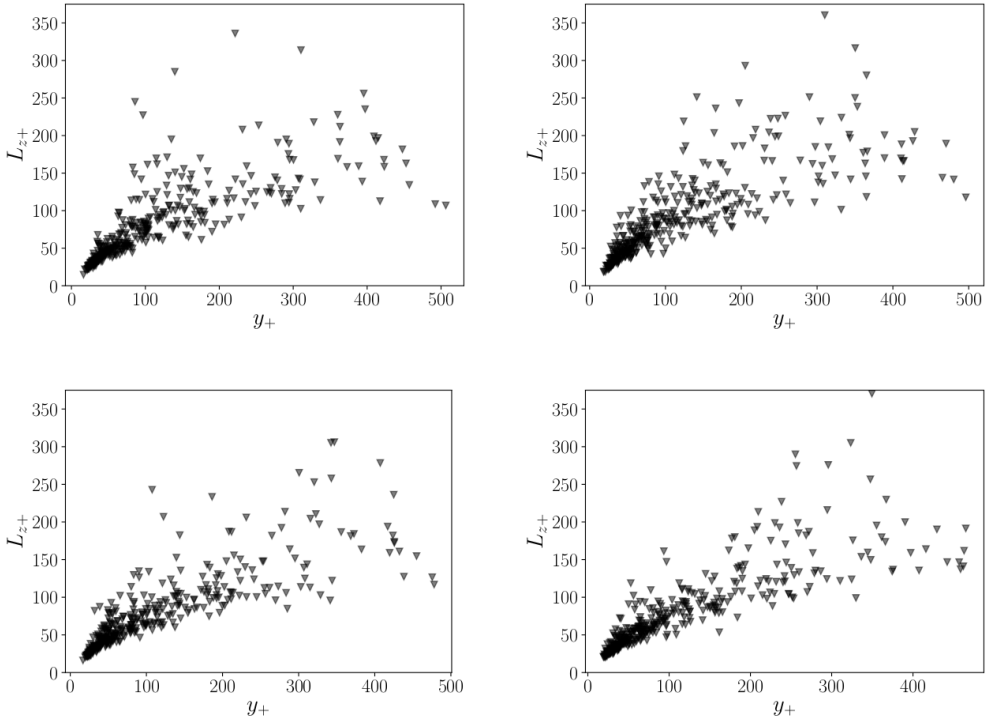


Figure 8: Motif characteristic vertical dimension for $N_T = 48$. Top row: Influence of dataset size; N_s : $N_s = 800$ (left), $N_s = 400$ (right); Bottom row: Effect of rescaling factor; $A = 60$ (left); $A = 20$ (right).

the streaks (Dennis 2015). It then decreases more slowly towards an aspect ratio of about 2 in the region $100 < y_+ < 500$. This ratio is consistent with results from analysis of POD eigenfunctions in Podvin *et al.* (2010), as well as from vortex cluster analysis from Alamo *et al.* (2006). 3-D motif characteristic sizes are consistent with those obtained for vertical planes, which shows that information about the 3-D organization of the flow can be obtained from analysis performed on 2-D sections. This is of particular interest as it suggests that the LDA method could be usefully applied to PIV experimental data.

6. Field reconstruction and generation

6.1. Reconstruction

We now examine how the flow can be reconstructed using LDA. In all that follows, without loss of generality, we will focus on one of the cross-flow planes examined in Section 5, specifically the cross-section at $x = 0$ of dimensions $d_{y+} = 590$ and $d_{z+} = 450$. As described in the algorithm presented in Section 3, both the motif-snapshot and the cell-motif distributions can be sampled for the total intensity $N_i = \sum_l f_{l,i}$ in the i -th snapshot. This total intensity is defined as the rescaled integral value of the Reynolds stress (digitized and restricted to Q_- events) over the plane. Since results were found to be essentially independent of the rescaling, we can make the simplifying assumption that N_i is large enough so that the distribution φ_n is well approximated by the samples. For

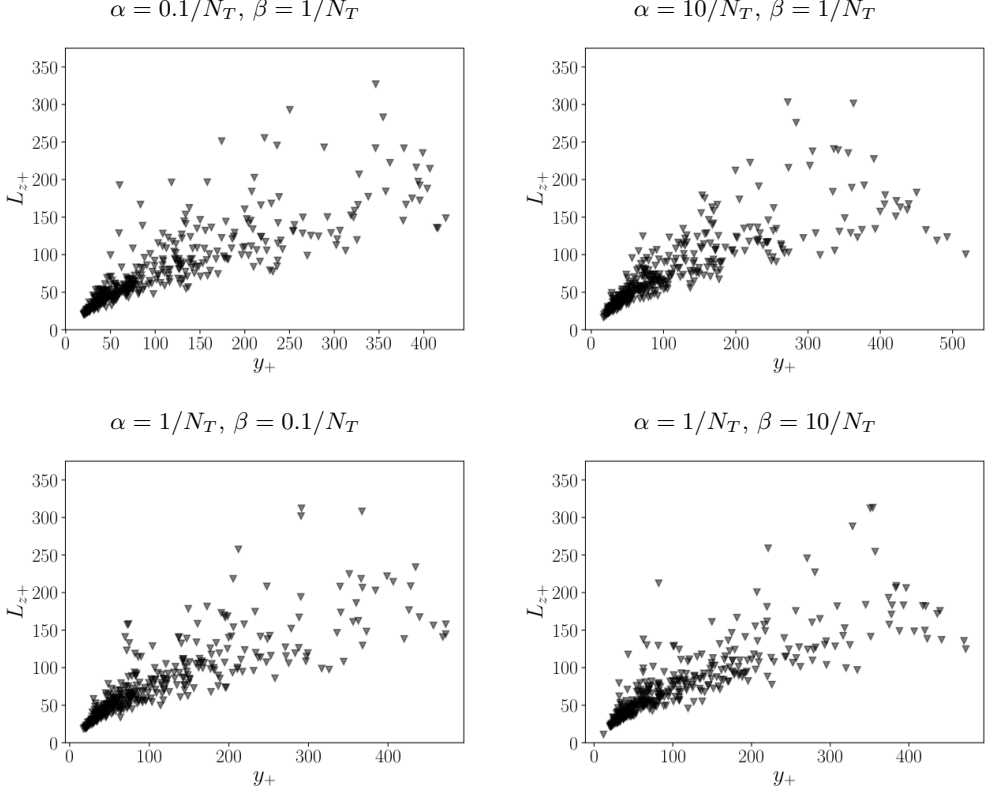


Figure 9: Characteristic vertical motif length for different LDA priors, $N_T = 48$.




Figure 10: Distribution p of motif/cell distribution maximum y_m for different parameters.

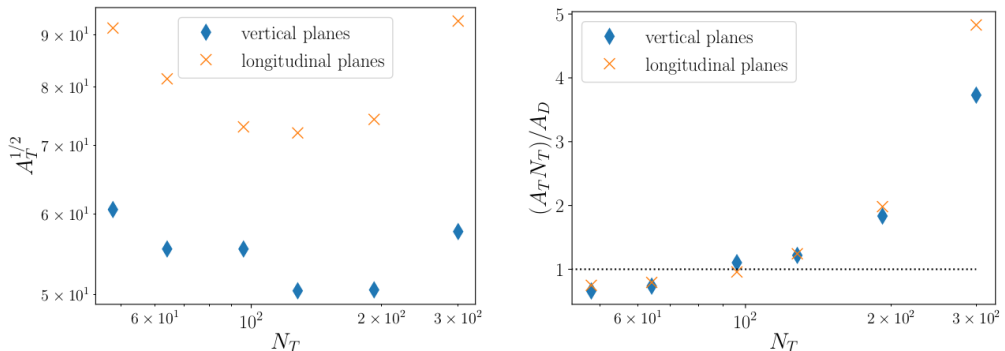


Figure 11: Left: motif characteristic dimension L_T for different datasets as a function of the number of motifs; Right: relative fraction of the area captured by the sum of the topics $N_T A_T / A_D$.

a given total intensity N_i , a reconstruction of the i -th snapshot can then be obtained at each grid cell \mathbf{x}_l from

$$\tau^{\text{R-LDA}}(\mathbf{x}, t_i) = \frac{1}{A} f_i(\mathbf{x}) = \frac{N_i}{A} \sum_{n=1}^{N_T} \theta_{n,i} \varphi_n(\mathbf{x}),$$

where

- $\varphi_n(\mathbf{x})$ is the motif-cell distribution,
- the snapshot-motif distribution $\theta_{n,i}$ represents the likelihood of motif \mathbf{z}_n in the i -th snapshot.

It seems natural to compare this reconstruction with the POD representation of the field which has a similar expression

$$\tau^{\text{R-POD}}(\mathbf{x}, t_i) = \sum_{n=0}^{N_{\text{POD}}-1} a_{n,i} \phi_n(\mathbf{x}),$$

where

- $\phi_n(\mathbf{x})$ are the POD eigenfunctions extracted from the autocorrelation tensor $C_{i,i'}$ obtained from the N_s snapshots,
- $a_{n,i}$ corresponds to the amplitude of the n -th POD mode in the i -th snapshot.

The first six fluctuating POD modes are represented in Figure 14. We note that the 0-th POD mode represents the temporal average of the field. As expected, the fluctuating POD modes consist of Fourier modes in that spanwise direction (due to homogeneity of the statistics), and their intensity reaches a maximum at around $y_+ \simeq 25$.

If the number of POD modes is equal to the number of motifs N_T , by construction, POD will provide a better representation of the statistics at least up to second-order (Holmes *et al.* 1996). We note that, in terms of computational requirement, POD may appear less expensive than LDA, as it requires solving an SVD problem versus implementing an iterative Expectation Maximization algorithm (Dempster *et al.* 1977). However the performance of the EM algorithm can be improved, in particular with online updates (Hofmann 1999).

In terms of storage, a reconstructed snapshot requires N_{POD} modes for POD and N_T topics for LDA. However, storage reduction could be obtained in the case of LDA by filtering out the motifs with a low probability $\theta_{n,i}$, *i.e.*, lower than a threshold κ . We

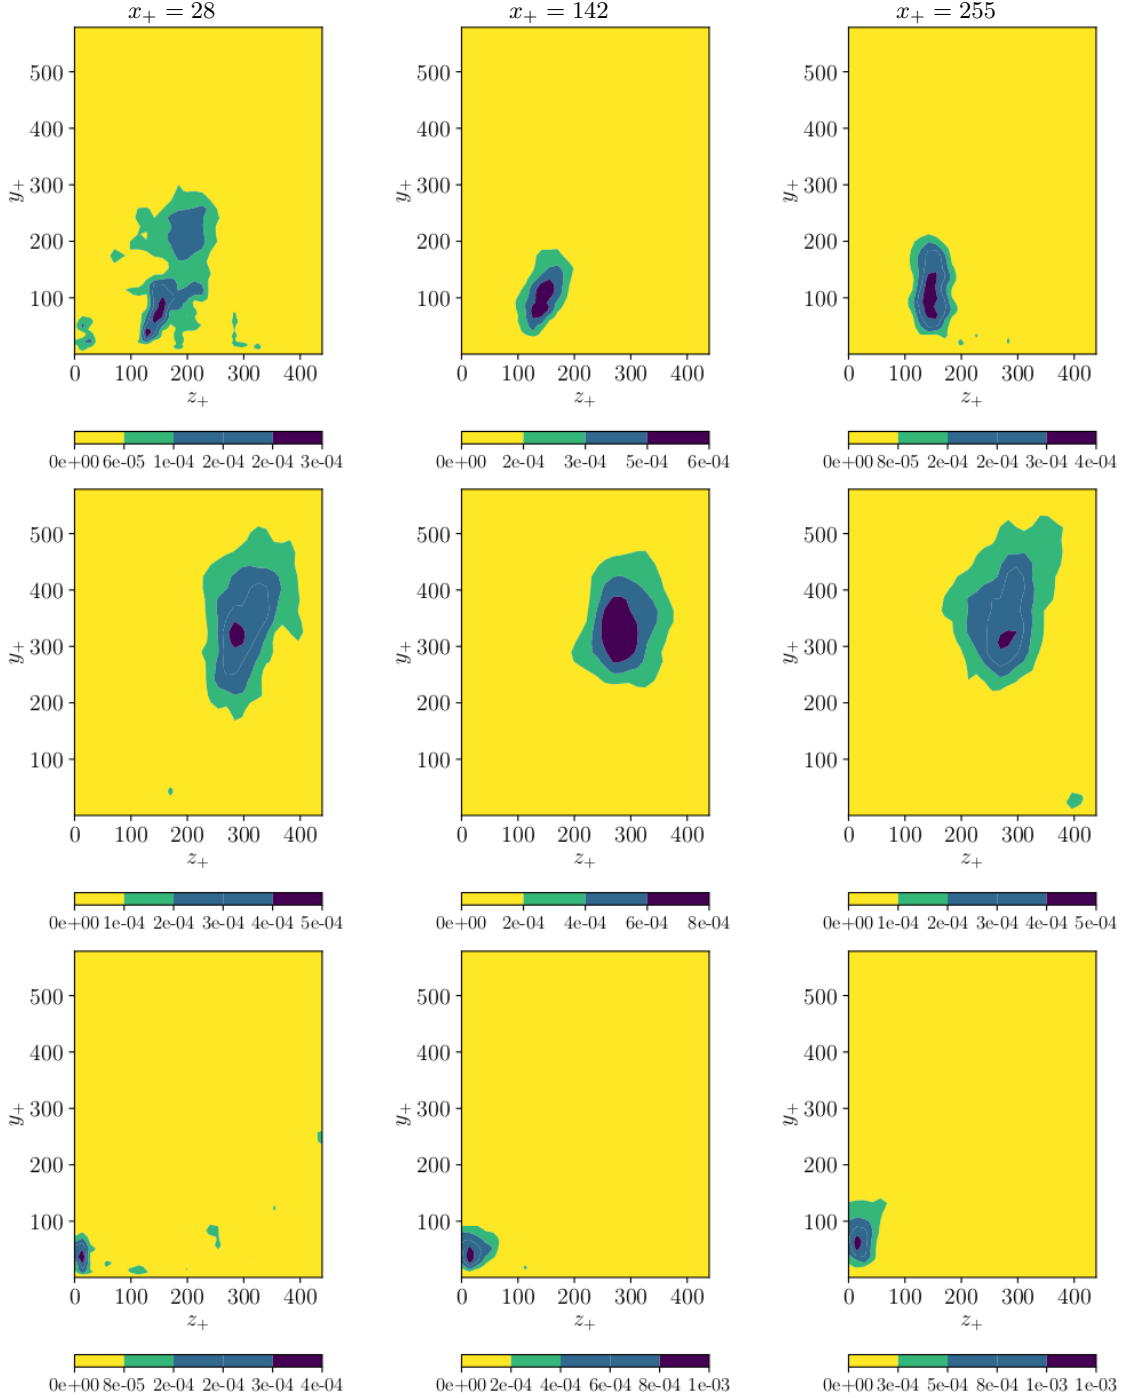


Figure 12: Cross-sections at different streamwise locations of three different 3D motifs obtained for $N_T = 144$; Top row: Motif index $n = 34$; Middle row: Motif index $n = 7$; Bottom row: Motif index $n = 24$.

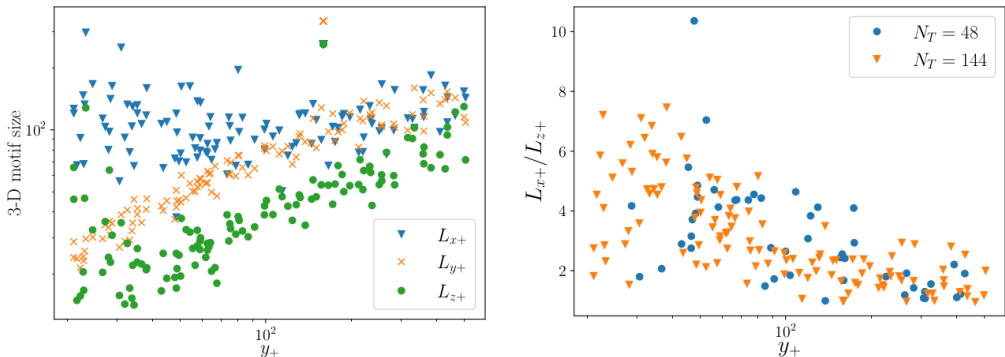


Figure 13: Left: Characteristic dimensions of the 3D motifs, $N_T = 144$; Right : Evolution of ratio L_x/L_z with height for $N_T = 144$ and $N_T = 48$.

note that, in this case, it is necessary to store the indices n of the motifs as well as the value of $\theta_{n,i}$, so that if n modes (resp. topics) are kept, storage will consist of $2n$ variables per snapshot. We see that storage reduction can be achieved if the fraction of retained modes $\eta = n/N_T$ is sufficiently small. The LDA storage data length per snapshot $2\eta N_T$ should then be compared with the POD data length N_{POD} .

For $N_T = 96$, choosing a threshold of $\kappa = 0.015$ resulted in less than 8% difference between the filtered and unfiltered LDA reconstructions (the L_2 norm was used). The average value for η was 0.2, which means that the number of POD modes that would represent a storage equivalent to that of LDA with $N_T = 96$ is $N_{\text{POD}} \simeq 2\eta N_T \simeq 40$. We note that the total storage cost should further take into account the size of the LDA basis $\{z_n\}_n$, which will be larger than the POD basis $\{\phi_n\}_n$ since they are respectively equivalent to N_T and N_{POD} fields. However efficient storage of the LDA basis can be achieved by making use of the limited spatial support of z_n , in particular for motifs located close to the wall.

In the remainder of this section we will compare a filtered LDA reconstruction of 96 motifs (where values of $\theta_{n,i}$ lower than $\kappa = 0.015$ are excluded from the reconstruction), with a POD representation of $N_{\text{POD}} = 48$ modes, which captures about 75% of the total energy. Figure 15 compares an instantaneous field with its LDA reconstruction and its POD reconstruction. A more general assessment is provided by Figure 16, which shows the correlation coefficient between each snapshot and its reconstruction based on POD as well as that based on LDA. Although POD appears to be slightly superior, the correlation coefficients are very close with respective average values of 0.75 for LDA and 0.77 for POD.

6.2. Generation

LDA is a generative model, so it is straightforward to generate synthetic snapshots by sampling from distributions θ and φ for a total intensity $N_i = \sum_l f_{l,i}$, which is modeled as a Poisson process with the same mean and standard deviation as the original database.

In contrast, POD is not a generative model *per se*. We will use a simplified version of the probabilistic extension of POD (PPCA) derived by [Tipping & Bishop \(1999\)](#), which is presented in Appendix B, where we will make the additional assumption that no noise is present in the model, POD-based synthetic fields will be reconstructed from deterministic spatial POD modes ϕ_n and random POD amplitudes \mathbf{a}_n which are assumed to be Gaussian variables. Examination of Figure 17, which represents the distribution of




Figure 14: Contour plot of the first six fluctuating normalized POD spatial modes; Contour values go from -0.03 to 0.03 . Negative values are indicated by dashed lines.




Figure 15: Instantaneous Reynolds stress field (limited to Q_- events) Left: True field; Middle: POD-reconstructed field using 48 POD modes; Right: LDA-reconstructed field using 96 modes.

the first fluctuating POD coefficients $n \geq 1$, suggests that it is quite acceptable as a first approximation to assume Gaussian distributions for the amplitudes \mathbf{a}_n — alternatively, the amplitudes could be sampled from the empirical distributions. The amplitude of the 0-th mode, which corresponds to the average of the field over the snapshots, will be assumed to be constant for all snapshots.

We can therefore compare the databases reconstructed from and generated with LDA with those obtained from POD. The generated databases consist of N_s snapshots corresponding to arbitrary instants \tilde{t}_i . Overall, the statistics of five different databases can be compared:

- the true database $\tau_-(y, z, t_i)$ corresponding to the actual values of the Q_- events




Figure 16: Distribution of the correlation coefficient between each original snapshot and its reconstruction based on LDA (top) or POD (bottom).

- the POD-reconstructed (R-POD) or POD-projected database

$$\tau_{-}^{\text{R-POD}}(y, z, t_i) = \sum_{n=0}^{N_{\text{POD}}-1} a_{i,n} \phi_n(y, z),$$

where ϕ_n are the POD eigenfunctions and $a_{i,n}$ are the amplitudes of the n -th POD mode in the i -th snapshot.

- the POD-generated (G-POD) database

$$\tau_{-}^{\text{G-POD}}(y, z, \tilde{t}_i) = \sum_{n=0}^{N_{\text{POD}}-1} \tilde{a}_{i,n} \phi_n(y, z),$$

where $\tilde{a}_{i,0} = \langle a_{i,0} \rangle$, with $\langle \cdot \rangle$ the average over all snapshots and $\tilde{a}_{i,n}$, $n \geq 1$, centered Gaussian random variables with variance $\langle \tilde{a}_{i,n}^2 \rangle$.

- the LDA-reconstructed database (R-LDA)

$$\tau_{-}^{\text{R-LDA}}(y, z, t_i) = \frac{N_i}{A} \sum_{n=1}^{N_T} \theta_{n,i} \varphi_n(y, z),$$

where N_i is the total intensity measured in the i -th snapshot, $\theta_{n,i}$ is the distribution of motif n on the i -th snapshot and $\varphi_n(y, z)$ is the identified distribution of the cell at (y, z) on motif n .

- the LDA-generated database (G-LDA)

$$\tau_{-}^{\text{G-LDA}}(y, z, \tilde{t}_i) = \frac{\tilde{N}_i}{A} \sum_{n=1}^{N_T} \tilde{\theta}_{n,i} \varphi_n(y, z),$$

where \tilde{N}_i is the total intensity, which is sampled from a Poisson process, $\varphi_n(y, z)$ is the identified distribution of the cell at (y, z) on motif n and $\tilde{\theta}_{n,i}$ is sampled for each n from the empirical distribution $\theta_{n,i}$ over the snapshots.

Figure 18 shows the statistics of the different databases as a function of the wall

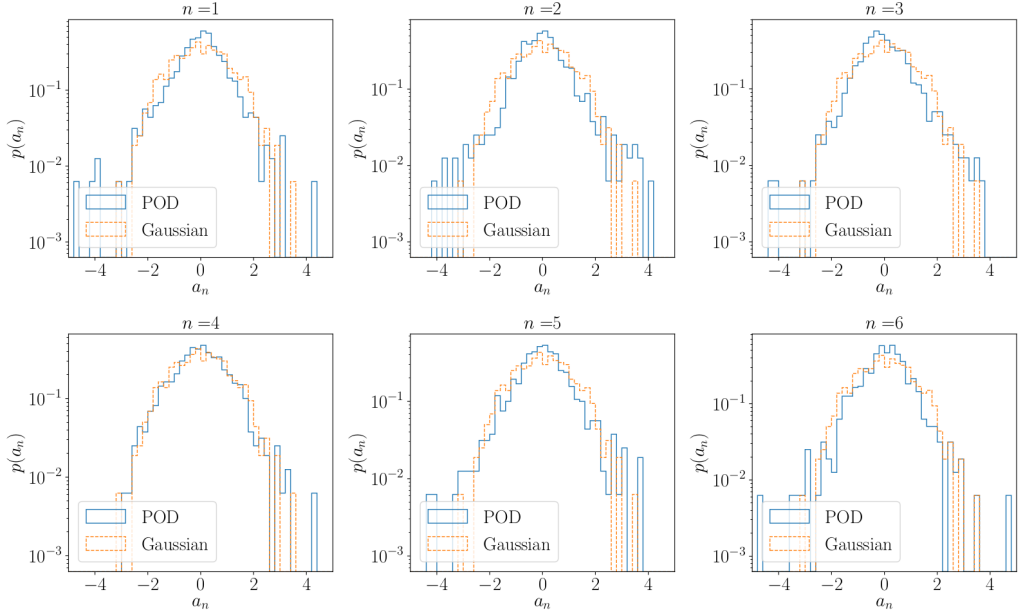


Figure 17: Histograms of the normalized amplitudes of the first six fluctuating POD modes and comparison with a sampled Gaussian distribution.

distance. Averages are taken over all snapshots and in the streamwise direction. The mean value of the Reynolds stresses is correctly recovered by all methods. The second-order statistics are slightly better recovered by the POD-reconstructed and POD-generated snapshot sets, but both LDA approaches also capture a significant portion of the variance. The POD databases capture 75% of the total variance, while the reconstructed and generated LDA databases respectively capture 68% and 60% of the variance. Figure 19 shows the vertical spatial autocorrelation of τ_- defined as $R(y, y') = \langle \tau_-(x, y, z, t) \tau_-(x, y', z, t) \rangle$ (where $\langle \cdot \rangle$ represents an average taken in time and in the spanwise position). We can see that the generated LDA autocorrelation is very similar to its reconstructed POD counterpart, which shows that the LDA synthetic fields capture as much as the spatial structure as the POD reconstructed ones. We note that the autocorrelation at large separations is well reproduced by all datasets.

Figure 20 shows histograms of the fields at different heights. We note that unlike the LDA approach, which is a non-negative decomposition (since it is based on probabilities), some negative values are observed for the POD approach, even though the original field values considered are always positive. We can see that at different wall distances the POD-reconstructed database reproduces well the distribution of the original database, but the POD-generated database does not. This failure is due to the fact that although POD amplitudes are uncorrelated by construction, they are not independent. We note that the same failure was observed when sampling the POD coefficients from their data-observed distributions $a_{i,n}$ instead of Gaussian processes. In contrast, both reconstructed and generated LDA methods yield very similar distributions, which reproduce the main features of the original Reynolds stress values, such as the intermittency (sharp peak at zero) and the asymptotic decay for positive values.

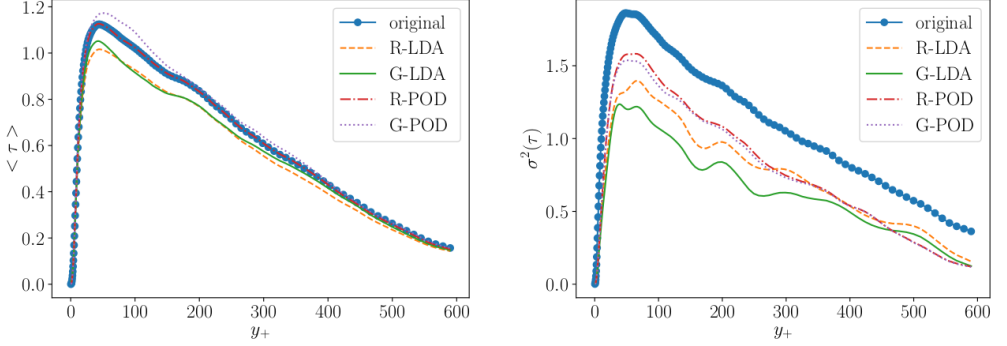


Figure 18: Statistics of the different databases averaged over the spanwise direction and the number of snapshots. Left: Mean value; Right: Standard deviation.




Figure 19: Spatial autocorrelation of the Reynolds stress (limited to Q_- events) in the vertical direction at different heights. The average is taken over snapshots and in the spanwise direction.

7. Conclusion

This paper presents exploratory work about the application of Latent Dirichlet Allocation (LDA) to the identification of coherent structures in turbulent flows. In the probabilistic framework of LDA, latent factors or motifs are inferred from a collection of snapshots. Each snapshot is characterized by a motif distribution, and each motif itself is distributed over space. Implementation was carried out for a scalar field representing Reynolds stress Q_- events. Evidence of self-similarity was found in the motifs: the spanwise and vertical dimensions of the motifs increase linearly with the wall distance in the logarithmic region, and the number of structures evolves inversely with the wall distance. This is in agreement with the eddy attached model hypotheses. The characteristics of the motifs were established to be robust with respect to the LDA parameters.

LDA yields a sparse, efficient reconstruction of the snapshots that compares reasonably well with POD representation. Adding in the fact that the motifs have a local spatial support, even when statistics are homogeneous, could make the LDA representation of interest for estimation and control purposes. Further, a strong benefit of LDA is

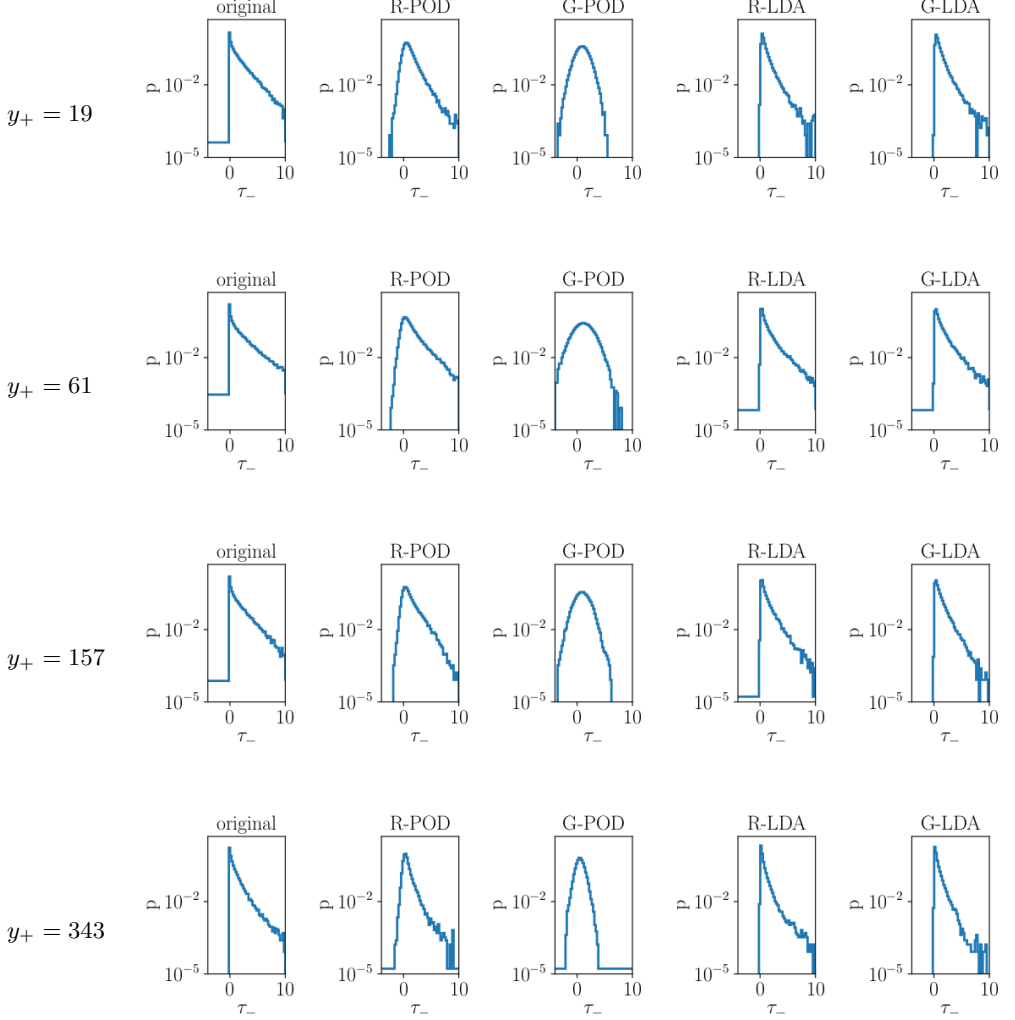


Figure 20: Histograms of the Reynolds stress (limited to Q_- events) corresponding to the different databases at different heights.

its inherent generative property, which makes it possible to generate a set of synthetic snapshots which is statistically similar to the original one.

The first results obtained with the LDA method open up exciting prospects for data analysis and modeling of turbulent flows. We plan to study larger domains at higher Reynolds numbers in future work. Moreover, while the investigation was limited to a positive scalar field in the present implementation, it would be useful to extend the capabilities of LDA to fully real, as well as multi-dimensional fields. Finally, since the technique appears well suited to describe intermittent phenomena, it would be interesting to apply it to strongly inhomogeneous flow regions such as the turbulent/non-turbulent interface (Philip *et al.* 2014) or other types of intermittency (Johnson & Meneveau 2017).

Acknowledgments

This work was supported by the Center of Data Science from the Paris-Saclay University. Computations were carried out at IDRIS-GENCI (project 02262). The authors are grateful to the anonymous Referees for their helpful comments on the first version of the manuscript.

Declaration of Interests

The authors report no conflict of interest.

Appendix A. LDA as a factorization method

To further shed light on the interpretation of LDA, we now adopt a different viewpoint and briefly explore the connections between the decomposition methods discussed above in the framework of Matrix Factorization (MF). Specifically, we now explain how model decomposition methods, such as POD, K-means and LDA, can be interpreted in terms of Matrix Factorization.

A.1. Matrix factorization

Letting $F \in \mathbb{R}^{N_x \times N_s}$ be a data matrix to be approximated, MF consists in the following decomposition:

$$F = XY, \quad (\text{A } 1)$$

with $X \in \mathbb{R}^{N_x \times N_T}$ and $Y \in \mathbb{R}^{N_T \times N_s}$ two real-valued matrices. Compression is achieved whenever $N_T < \min(N_x, N_s)$, which is considered hereafter. MF can be formulated as an optimization problem:

$$(X, Y) \in \arg \min_{\tilde{X} \in \mathcal{S}_X, \tilde{Y} \in \mathcal{S}_Y} \left\| F - \tilde{X}\tilde{Y} \right\|^2 + \mathcal{R}(\tilde{X}, \tilde{Y}), \quad (\text{A } 2)$$

with $\|\cdot\|$ a given norm, \mathcal{S}_X and \mathcal{S}_Y admissibility sets for X and Y respectively, and \mathcal{R} a regularization term.

A.2. POD-MF equivalence

Let the singular value decomposition (SVD) of the $N_x \times N_s$ real-valued data matrix F be

$$F = \Psi \Sigma B^\top, \quad (\text{A } 3)$$

with Ψ and B two orthonormal matrices and Σ being diagonal. The Eckart-Young theorem makes precise in which sense this decomposition is optimal, [Eckart & Young \(1936\)](#). In particular, it follows that

$$\Psi_{N_T}, (\Sigma B^\top)_{N_T} \in \arg \min_{\tilde{\Psi}^\top \tilde{\Psi} = I_{N_T}} \left\| F - \tilde{\Psi} \left(\widetilde{\Sigma B^\top} \right) \right\|_F, \quad \forall N_T \leq \min(N_x, N_s), \quad (\text{A } 4)$$

where $(\Sigma B^\top)_{N_T} = \Sigma_{N_T} B_{N_T}^\top$ and with Ψ_{N_T} and B_{N_T} the restriction of Ψ and B to their columns associated with the dominant N_T singular values $\text{diag}(\Sigma_{N_T})$.

From Eq. (A 3), it comes

$$F F^\top \Psi = \Psi_{N_T} \Sigma_{N_T} B_{N_T}^\top B_{N_T} \Sigma_{N_T}^\top \Psi_{N_T}^\top \Psi_{N_T} = \Psi_{N_T} \Sigma_{N_T}^2 = C_{N_T} \Psi_{N_T}. \quad (\text{A } 5)$$

Referring to Eqs. (2.1) and (2.2), the diagonal matrix $\Sigma_{N_T}^2$ and Ψ_{N_T} then directly

identify with the N_T dominant eigenvalues Λ and POD modes Φ , respectively. Denoting the Moore-Penrose pseudo-inverse with a $^+$ superscript, the POD projection coefficients are:

$$A = \Phi^+ F = \Phi^\top F = \Psi_{N_T}^\top F = \Sigma_{N_T} B_{N_T}^\top, \quad (\text{A } 6)$$

so that the POD decomposition is finally seen to satisfy the following matrix factorization problem:

$$\Phi, A \in \arg \min_{\Phi^\top \Phi = I_{N_T}} \|F - \Phi A\|_F, \quad (\text{A } 7)$$

of the form of Eq. (A 2) with $\mathcal{R} \equiv 0$ and \mathcal{S}_X such that $X^\top X = I_{N_T}$.

A.3. *K-means-MF equivalence*

Clustering is an unsupervised learning technique aiming at identifying groups (clusters) in the data such that data points in the same group have similar features, while data points in different groups have highly dissimilar features.

K-means is one of the simplest and popular clustering methods, [MacQueen \(1967\)](#); [Lloyd \(1982\)](#). The algorithm tries to iteratively partition the dataset into N_T predefined distinct non-overlapping clusters $\{\mathcal{C}_n\}_n$. In its standard deterministic version, each data point belongs to only one cluster. The key idea consists in assigning each data point to the closest centroid (arithmetic mean of all the data points that belong to that cluster). The distance is defined in terms of some chosen norm $\|\cdot\|$. Setting the number of clusters N_T , the algorithm starts with an initial guess for the N_T centroids $\{\mathbf{c}_n\}_n$, by randomly selecting N_T data points from the data set without replacement. It then iterates between the data assignment step, assigning each data point \mathbf{f}_i to the closest cluster $\mathcal{C}_{n_i^*}$ and the centroid update step, which computes the centroid of each cluster:

$$n_i^* \leftarrow \arg \max_{1 \leq \tilde{n} \leq N_T} \|\mathbf{c}_{\tilde{n}} - \mathbf{f}_i\|^2, \quad \forall 1 \leq i \leq N_s, \quad (\text{A } 8)$$

$$\mathbf{c}_n \leftarrow \frac{1}{\text{card}[\mathcal{C}_n]} \sum_{\mathbf{f}_i \in \mathcal{C}_n} \mathbf{f}_i, \quad \forall 1 \leq n \leq N_T. \quad (\text{A } 9)$$

K-means is guaranteed to converge to a local optimum but not necessarily to a global optimum. Therefore, we choose to run the algorithm with different initializations of centroids and retain the solution that yielded the lowest loss \mathcal{L} :

$$\mathcal{L} = \sum_{n=1}^{N_T} \sum_{\mathbf{f}_i \in \mathcal{C}_n} \|\mathbf{f}_i - \mathbf{c}_n\|^2. \quad (\text{A } 10)$$

Solving a clustering problem in the L^2 -sense means finding a set of $\{\mathcal{C}_n\}_{n=1}^{N_T}$ disjoint clusters ($\mathcal{C}_n \cap \mathcal{C}_{n'} = \{\emptyset\}$, $n \neq n'$), that minimizes the following cost function:

$$\mathcal{L} = \sum_{n=1}^{N_T} \sum_{\mathbf{f}_i \in \mathcal{C}_n} \|\mathbf{f}_i - \mathbf{c}_n\|_2^2 = \sum_{i=1}^{N_s} \|\mathbf{f}_i\|_2^2 - \sum_{n=1}^{N_T} \sum_{\mathbf{f}_i, \mathbf{f}_{i'} \in \mathcal{C}_n} n_n^{-1} \mathbf{f}_i^\top \mathbf{f}_{i'}, \quad (\text{A } 11)$$

where $\{\mathbf{c}_n\}_{n=1}^{N_T}$ are the cluster centroids, $\mathbf{c}_n := \sum_{\mathbf{f}_i \in \mathcal{C}_n} \mathbf{f}_i / n_n$, $n_n := \text{card}[\mathcal{C}_n]$.

Let $Y \in [0, 1]^{N_s \times N_T}$ be the normalized cluster indicator matrix, $\mathbf{y}_n = n_n^{-1/2} \mathbb{1}_{\{\mathbf{f}_i \in \mathcal{C}_n\}}$. Disjointness of clusters implies that columns of Y are orthonormal, $Y^\top Y = I_{N_T}$. The

clustering problem (A 11) may now reformulate in terms of $Y \geq 0$ as, Ding *et al.* (2005):

$$\begin{aligned}
Y &\in \arg \min_{\tilde{Y} \geq 0, \tilde{Y}^\top \tilde{Y} = I_{N_T}} \text{Tr} [F^\top F] - \text{Tr} [\tilde{Y}^\top F^\top F \tilde{Y}], \\
&\in \arg \min_{\tilde{Y} \geq 0, \tilde{Y}^\top \tilde{Y} = I_{N_T}} \|F^\top F\|_F^2 - 2\text{Tr} [\tilde{Y}^\top F^\top F \tilde{Y}] + \|\tilde{Y}^\top \tilde{Y}\|_F^2, \\
&\in \arg \min_{\tilde{Y} \geq 0, \tilde{Y}^\top \tilde{Y} = I_{N_T}} \|F^\top F - \tilde{Y} \tilde{Y}^\top\|_F^2.
\end{aligned} \tag{A 12}$$

The Euclidean hard-clustering K-means problem hence stems from an orthogonal non-negative matrix factorization form and the clusters are given by $\mathbf{c}_n = n_n^{-1/2} F \mathbf{y}_n$, $\forall n$.

A.4. LDA-MF equivalence

We now focus on LDA and discuss the fact that, similarly to POD and K-means, it can also be interpreted as a matrix factorization technique, under certain conditions.

Let us consider the variational LDA flavor, where inferring the LDA parameters from maximizing the posterior distribution p is substituted with an approximated posterior q , easier to sample from. The inference problem then consists in minimizing the approximation error, which is equivalent to maximizing the Evidence Lower Bound (ELBO) \mathcal{L} :

$$\mathcal{L} = \mathbb{E}_{\mu_q} [p] - \mathbb{E}_{\mu_q} [q]. \tag{A 13}$$

Provided suitable approximations in the inference problem are made, and under a symmetric Dirichlet priors hypothesis ($\alpha = \alpha \mathbf{1}$), Faleiros & Lopes (2016) have derived an upper bound for the ELBO associated with variational LDA:

$$\max \mathcal{L} \lesssim \min \sum_l^{N_x} \sum_i^{N_s} \left(F_{l,i} \log \frac{F_{l,i}}{(XY)_{l,i}} + \sum_n^{N_T} \mathcal{R}(Y_{n,i}, \alpha_n) \right), \tag{A 14}$$

where $X \geq 0$ and $Y \geq 0$ are variational parameters to infer, normalized as $\sum_l X_{l,n} = \sum_n Y_{n,i} = 1$, and regarded as normalized probability distributions. \mathbf{x}_n is related to β while \mathbf{y}_i is related to the distribution θ_i of a document \mathbf{f}_i . The term $\mathcal{R}(Y_{n,i}, \alpha_n) := (Y_{n,i} - \alpha_n)(\log Y_{n,i} - Y_{n,i}(\log Y_{n,i} - 1))$ corresponds to the prior influence and induces sparsity over the document-topic distribution.

From Eq. (A 14), it follows that maximizing the ELBO \mathcal{L} under certain approximations takes the form of a non-negative matrix factorization problem (NMF) of $F \approx XY$ expressed in terms of the Kullback-Leibler divergence $D_I(F \| XY) := \sum_{l,i} \left(F_{l,i} \log \frac{F_{l,i}}{(XY)_{l,i}} - F_{l,i} + (XY)_{l,i} \right)$, supplemented with a regularization term.

Details of the derivation are beyond the scope of this paper and one should refer to Faleiros & Lopes (2016) for a more complete discussion.

Appendix B. Probabilistic PCA/ POD

In this section, we give a brief review of Probabilistic PCA (PPCA) (Tipping & Bishop 1999) which provides a density estimation framework for POD (or PCA/LSA), under hypotheses that are different from those given in section 2.2 for PLSA.

We will assume that the data is zero-centered without loss of generality. The basic idea of PPCA is to assume a Gaussian probability model for the observed data f' . In that formulation (see section 2.2), the motif-cell matrix $\tilde{\Phi}$ of dimension $N_x \times N_T$ does

not have a probabilistic interpretation, but relates each noisy observation to a set of N_T independent normalized Gaussian variables following

$$f' = \tilde{\Phi} \tilde{a} + \epsilon \quad (\text{B } 1)$$

where the variables \tilde{a} are defined to be independent and Gaussian with unit variance and ϵ represents noise.

An important assumption to proceed is that the model for the noise should be isotropic

$$\langle \epsilon \epsilon \rangle = \sigma^2 I,$$

so that all the dependences between the observations are going to be contained in $\tilde{\Phi}$. On can then show using equation (B 1) that

$$p(f') = \mathcal{N}(0, C)$$

where $C = \tilde{\Phi}^T \tilde{\Phi} + \sigma^2 I$ is the observation covariance matrix of dimension N_x^2 .

The issue is to determine $\tilde{\Phi}$ and $\tilde{\sigma}$, given the observations of f' . Under the assumption of isotropic Gaussian noise, [Tipping & Bishop \(1999\)](#) showed that the maximum likelihood estimators $\hat{\Phi}$ and $\hat{\sigma}^2$ can be obtained from standard POD analysis on the N_s snapshots. They showed that

$$\hat{\Phi} = \Phi (\Lambda_{nmode} - \sigma^2 I_{N_T})^{1/2} R \quad (\text{B } 2)$$

where Φ contains the first N_T eigenvectors of the sampled covariance matrix \tilde{C} where \tilde{C} was defined in equation 2.4 (note that the dimension of \tilde{C} is N_s^2), Λ_{N_T} is a diagonal matrix containing the N_T first eigenvalues of \tilde{C} and R is an arbitrary rotation matrix.

An estimate for the error variance can then be given by

$$\hat{\sigma}^2 = \frac{1}{N_s - N_T} \sum_{j=N_T+1}^{N_s} \lambda_j, \quad (\text{B } 3)$$

which represents the variance lost in the project and averaged over the lost dimensions.

REFERENCES

- ALAMO, J.C. DEL, JIMENEZ, J., ZANDONADE, P. & MOSER, R.D. 2006 Self-similar vortex clusters in the turbulent logarithmic region. *J. Fluid Mech.* **561**, 329–358.
- AUBERT, A. H., TAVENARD, R., EMONET, R., DE LAVENNE, A., MALINOWSKI, S., GUYET, T., QUINIOU, R., ODOBEZ, J.-M., MEROT, P. & GASCUEL-ODOUX, C. 2013 Clustering flood events from water quality time series using latent Dirichlet allocation model. *Water Resources Research* **49** (12), 8187–8199.
- BAARS, W. J., HUTCHINS, N. & MARUSIC, I. 2017 Self-similarity of wall-attached turbulence in boundary layers. *Journal of Fluid Mechanics* **823**, R2.
- BLEI, D., NG, A. & JORDAN, M.I. 2003 Latent Dirichlet allocation. *Journal of Machine Learning Research* **3**, 993–1022.
- BRUNTON, S. L., NOACK, B. R. & KOUMOUTSAKOS, P. 2020 Machine learning for fluid mechanics. *Annual Review of Fluid Mechanics* **52** (1), 477–508.
- CANTWELL, B.J. 1981 Organized motion in turbulent flow. *Ann. Rev. Fluid Mech.* **13**, 457–515.
- CHENG, CHENG, LI, WEIPENG, LOZANO-DURÁN, ADRIÁN & LIU, HONG 2020 Uncovering townsend’s wall-attached eddies in low-reynolds-number wall turbulence. *Journal of Fluid Mechanics* **889**, A29.
- DEMPSTER, A. P., LAIRD, N. M. & RUBIN, D. B. 1977 Maximum likelihood from incomplete data via the em algorithm. *JOURNAL OF THE ROYAL STATISTICAL SOCIETY, SERIES B* **39** (1), 1–38.
- DENNIS, J.C.D. 2015 Coherent structures in wall-bounded turbulence. *An. Acad. Bras. Cienc.* **87** (2), 161–193.

- DING, C., HE, X. & SIMON, H.D. 2005 On the equivalence of nonnegative matrix factorization and spectral clustering. In *in SIAM International Conference on Data Mining*.
- ECKART, C. & YOUNG, G. 1936 The approximation of one matrix by another of lower rank. *Psychometrika* **1** (3), 211–218.
- FALEIROS, T. & LOPES, A. 2016 On the equivalence between algorithms for non-negative matrix factorization and latent dirichlet allocation. In *ESANN European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*. Bruges (Belgium).
- FLORES, O. & JIMENEZ, J. 2010 Hierarchy of minimal flow units in the logarithmic layer. *Physics of Fluids* **22**, 071704.
- GRIFFITHS, T. L. & STEYVERS, M. 2002 A probabilistic approach to semantic representation. In *Proceedings of the 24th Annual Conference of the Cognitive Science Society*.
- GRIFFITHS, T. L. & STEYVERS, M. 2004 Finding scientific topics. *Proceedings of the National Academy of Sciences* **101** (suppl 1), 5228–5235.
- HELLSTRÖM, LEO H. O., MARUSIC, IVAN & SMITS, ALEXANDER J. 2016 Self-similarity of the large-scale motions in turbulent pipe flow. *Journal of Fluid Mechanics* **792**, R1.
- HOFMANN, T. 1999 Probabilistic latent semantic analysis. In *Proceedings of Uncertainty in Artificial Intelligence*. UAI99.
- HOLMES, P., LUMLEY, J.L. & BERKOOZ, GAL 1996 *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge University Press.
- HU, R., YANG, X. I. A. & ZHENG, X. 2020 Wall-attached and wall-detached eddies in wall-bounded turbulent flows. *Journal of Fluid Mechanics* **885**, A30.
- HWANG, J. & SUNG, H. J. 2018 Wall-attached structures of velocity fluctuations in a turbulent boundary layer. *Journal of Fluid Mechanics* **856**, 958–983.
- HWANG, J. & SUNG, H. J. 2019 Wall-attached clusters for the logarithmic velocity law in turbulent pipe flow. *Physics of Fluids* **31** (5), 055109.
- JIMENEZ, J. 2013 Near-wall turbulence. *Physics of Fluids* **25** (1), 101302.
- JIMENEZ, J. 2018 Coherent structures in wall-bounded turbulence. *Journal of Fluid Mechanics* **842**, P1.
- JOHNSON, P. L. & MENEVEAU, C. 2017 Turbulence intermittency in a multiple-time-scale navier-stokes-based reduced model. *Phys. Rev. Fluids* **2**, 072601.
- KIM, H.T., KLINE, S.J. & REYNOLDS, W.C. 1971 The production of turbulence near a smooth wall in a turbulent boundary layer. *J. Fluid Mech.* **50** (1), 133–160.
- KLINE, S.J., REYNOLDS, W.C., SCHRAUB, F.A. & RUNSTADLER, P.W. 1967 The structure of turbulent boundary layers. *J. Fluid Mech.* **30** (4), 741–773.
- LLOYD, S. P. 1982 Least squares quantization in PCM. *IEEE Transactions on Information Theory* **28** (2), 129–137.
- LOZANO-DURAN, A., FLORES, O. & JIMENEZ, J. 2012 The three-dimensional structure of momentum transfer in turbulent channels. *J. Fluid Mech.* **524**, 1–31.
- LUMLEY, J.L. 1967 The structure of inhomogeneous turbulent flows. In *Atmospheric Turbulence and Radio Wave Propagation* (ed. A.M Iaglom & V.I Tatarski), pp. 221–227. Nauka, Moscow.
- MACQUEEN, J. 1967 Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, pp. 281–297. Berkeley, Calif.: University of California Press.
- MARUSIC, I. & MONTY, J. P. 2019 Attached eddy model of wall turbulence. *Annual Review of Fluid Mechanics* **51** (1), 49–74.
- MOSER, R., KIM, J. & MANSOUR, N.N. 1999 Direct numerical simulation of turbulent channel flow up to $re_\tau = 590$. *Phys. Fluids* **11** (4), 943.
- MURALIDHAR, S., PODVIN, B., MATHÉLIN, L. & FRAIGNEAU, Y. 2019 Spatio-temporal proper orthogonal decomposition of turbulent channel flow. *J. Fluid Mech.* **864**, 614–639.
- PEDREGOSA, F., VAROQUAUX, G., GRAMFORT, A., MICHEL, V., THIRION, B., GRISEL, O., BLONDEL, M., PRETTENHOFER, P., WEISS, R., DUBOURG, V., VANDERPLAS, J., PASSOS, A., COURNAPEAU, D., BRUCHER, M., PERROT, M. & DUCHESNAY, E. 2011 Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830.
- PERRY, A.E. & CHONG, M.S. 1982 On the mechanism of wall turbulence. *J. Fluid Mech.* **119**, 173–217.

- PERRY, A.E. & MARUSIC, I. 1995 A wall-wake model for the turbulent structure of boundary layers part 1. extension of the attached eddy hypothesis. *J. Fluid Mech.* **298**, 361.
- PHILIP, J., MENEVEAU, C., DE SILVA, C. M. & MARUSIC, I. 2014 Multiscale analysis of fluxes at the turbulent/non-turbulent interface in high reynolds number boundary layers. *Physics of Fluids* **26** (1), 015105.
- PODVIN, B. & FRAIGNEAU, Y. 2017 A few thoughts on proper orthogonal decomposition in turbulence. *Physics of Fluids* **29**, 531.
- PODVIN, B., FRAIGNEAU, Y., JOUANGUY, J. & LAVAL, J.P. 2010 On self-similarity in the inner wall layer of a turbulent channel flow. *Journal Fluids Engineering* **132** (4), 41202.
- ROBINSON, S.K. 1991 Coherent motions in the turbulent boundary layer. *Ann. Rev. Fluid Mech.* **23**, 601–639.
- SHARMA, A. S. & MCKEON, B. J. 2013 On coherent structure in wall turbulence. *Journal of Fluid Mechanics* **728**, 196–238.
- SIROVICH, L. 1987 Turbulence and the dynamics of coherent structures part i: Coherent structures. *Quart. Appl. Math.* **45** (3), 561–571.
- TIPPING, M. E. & BISHOP, C. M. 1999 Probabilistic principal component analysis. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **61** (3), 611–622.
- TOWNSEND, A.A. 1947 Measurements in the turbulent wake of a cylinder. *Proc. R. Soc. London Ser. A* **190**, 551–561.
- TOWNSEND, A. A. 1961 Equilibrium layers and wall turbulence. *Physics of Fluids* pp. 97–120.
- WALLACE, J.M., ECKELMANN, H. & BRODKEY, R.S. 1972 The wall region in turbulent shear flow. *J. Fluid Mech.* **54** (1), 39–48.
- WALLACE, J. M. 2016 Quadrant analysis in turbulence research: History and evolution. *Annual Review of Fluid Mechanics* **48** (1), 131–158.
- WILLMARTH, W. W. & LU, S. S. 1972 Structure of the reynolds stress near the wall. *J. Fluid Mech.* **55** (1), 65–92.
- WOODCOCK, J.D. & MARUSIC, I. 2015 The statistical behavior of attached eddies. *Physics of Fluids* **27**, 015104.