1

Introduction

Figure 1.1 Oil distribution in the 2010 Deepwater Horizon oil spill in the Gulf of Mexico. Apparent barriers to the transport of oil are readily observable. Image: Daniel Beltrá (used with permission).

Flows in nature tend to generate striking patterns in the tracers they carry. An example is the oil spill distribution in Fig. 1.1, which is framed by apparent barriers to the spread of the oil in certain directions. More often than not, one's primary interest is to find such barriers and hence understand the overall direction and rate of transport without necessarily identifying pointwise oil concentration values with high accuracy. Figure 1.1 also conveys the strong technological and societal needs for uncovering, forecasting and shaping such barriers.

We think of *transport barriers* as observed inhibitors of the spread of substances in flows. They offer a simplified global template for the redistribution of those substances without the need to simulate or observe numerous different initial distributions in detail. Because of their simplifying role, transport barriers are broadly invoked as explanations for observations in several physical disciplines, including geophysical flows (Weiss and Provenzale, 2008), fluid dynamics (Ottino, 1989), plasma fusion (Dinklage et al., 2005), reactive flows (Rosner, 2000) and molecular dynamics (Toda, 2005).

Despite their frequent conceptual use, however, transport barriers are rarely defined precisely or extracted systematically from data. The purpose of this book is to survey effective and mathematically grounded methods for defining, locating and leveraging transport barriers

https://doi.org/10.1017/9781009225199.002 Published online by Cambridge University Press

in numerical simulations, laboratory experiments, technological processes and nature. In the rest of this Introduction, we briefly survey the main topics that we will be covering in later chapters.

1.1 The Mathematics of Transport Barriers

Throughout this book, we will adopt the geometric view of nonlinear dynamical systems theory on transport. That is, rather than focusing on individual fluid particle positions or pointwise concentration values, we seek to identify key invariant surfaces with a major impact on shaping transport patterns.

To illuminate the significance of such invariant surfaces, we note that models of transport phenomena are often tested based on their ability to reproduce the evolution of an initial condition or initial distribution. Due to their inherent sensitivity on initial conditions, parameters and uncertainties, however, even predictions from highly accurate models can ultimately display vast discrepancies with individual observations. A more meaningful way to test the validity of models is to assess their ability to reproduce transport barriers of the physical process accurately.

Figure 1.2 shows a conceptual example in which a highly accurate dynamical model (black dots) for the true dynamics (solid red) makes a vastly inaccurate prediction for the evolution of a single initial condition (blue), yet reproduces a transport barrier (the unstable manifold of a saddle point) of the true dynamics with high accuracy. Clearly, it is the latter metric in which the model should be assessed and found very accurate.

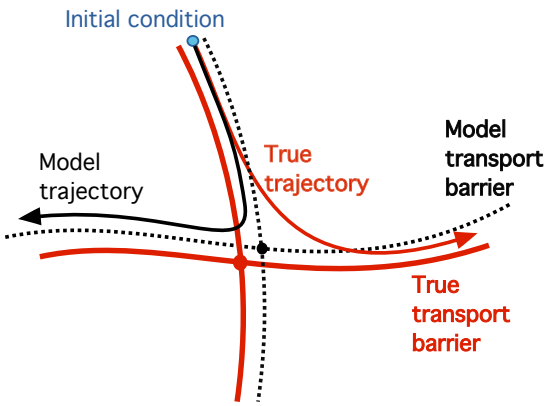

Figure 1.2 A good model for a transport process should accurately predict transport barriers but not necessarily individual trajectories.

A transport barrier can only block transport locally if its dimension is *precisely* one less than the dimension of the flow domain, as in Fig. 1.2. Such smooth surfaces are called codimension-one invariant manifolds in dynamical systems theory. They are curves in two-dimensional flows and two-dimensional surfaces in three-dimensional flows. While the impact of transport barriers is apparent in Figs. 1.1–1.9, their ever-shifting location, life span and intrinsic properties are not readily recoverable from instantaneous snapshots of an unsteady flow.

https://doi.org/10.1017/9781009225199.002 Published online by Cambridge University Press

Our underlying assumption will be that either the velocity field or a set of its trajectories is available either from observations or numerical simulations. With this assumption, the two fundamental questions we seek to answer are:

(1) *What defines the fundamental barriers to the transport of active and passive tracers in the flow?*
(2) *How can we locate the most influential transport barriers without solving their underlying transport equation?*

To review the technical material needed to answer these questions, Chapter 2 recalls relevant concepts from the geometric theory of dynamical systems and continuum mechanics in the context and language of fluid mechanics. In later chapters, we will build on these concepts to answer the two main questions posed above under increasingly fewer assumptions. The order of chapters will roughly mirror actual chronological developments in fluid mechanics and applied mathematics, each describing ideas and results of interest in their own right.

1.2 The Physics of Transport Barriers

While the techniques for transport-barrier detection we put forward in this book are quite diverse, we will set one common physical requirement for all of them: their predictions must be experimentally verifiable. Experimental verifiability for predicted transport barriers means that they should be detectable via material tracers such as dye, weakly diffusive smoke or small particles.

The requirement of unambiguous experimental reproducibility has an important implication: any self-consistent description of transport barriers must be independent of the frame of reference of the observer. For instance, barriers framing the evolution of an oil spill, such as the one shown in Fig. 1.1, are clearly identified as the same set of material points by an observer on the beach, another one on a ship and a third one in a circling helicopter, as illustrated in Fig. 1.3.

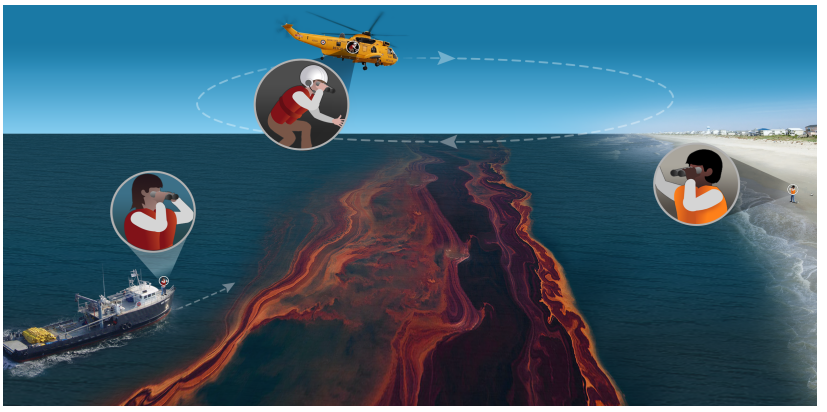

Figure 1.3 Observed transport barriers highlighted by tracers must be objectively defined, i.e., independently of the observer's frame of reference.

https://doi.org/10.1017/9781009225199.002 Published online by Cambridge University Press

These material barriers move on different orbits in each observer's frame of reference, yet all three observers agree exactly on where the edges of the oil spill are at any time instance. Accordingly, any self-consistent definition or detection method for transport barriers should be *objective*, i.e., indifferent to the observer and return the same set of material points forming the barrier in any observer frame. Beyond its theoretical significance, this litmus test for the self-consistency of barrier theories has a very practical motivation. If resource allocation or countermeasures are planned based on the location of transport barriers in a flow, there is no room for disagreement among observers.

In Chapter 3, we will discuss the principle of material frame indifference in detail and explore its use as a litmus test for the self-consistency of various views on transport barriers.

1.3 Idealized Transport Barriers vs. Finite-Time Coherent Structures

As we will see in Chapter 4, purely advective transport barriers can be defined unambiguously in flows that are known for all times. Such definitions, e.g., those requiring saddle-type behavior for a material surface for all times, exploit properties of the barriers that make them locally unique in the flow. The area of transport studies concerned with such idealized systems is generally referred to as chaotic advection. In its idealized setting, chaotic advection provides very helpful motivation for the main types of material barriers one encounters in nature. It also turns out that some approaches to diffusive and dynamically active transport barriers also yield steady velocity fields (barrier vector fields) that can be analyzed by the methods we review for steady flows in Chapter 4.

In practical problems, however, finite-time transport is of interest. Flow behavior over a finite time interval is continuous: all close enough trajectories remain close to each other. Consequently, infinitely many neighboring trajectories will satisfy any barrier definition phrased via inequalities that involve continuous quantities. As a result, the defining features of barriers exploited in chaotic advection no longer yield unique material surfaces.

One way to address this nonuniqueness problem is to identify barriers that act as observed centerpieces of material deformation over a finite time interval of interest. We perceive a surface to assume this role if it remains coherent, i.e., keeps its spatial integrity without developing smaller scales (filaments) during its temporal evolution. We will call such uniquely resilient material sets Lagrangian coherent structures (LCSs), as first coined by Haller and Yuan (2000). These LCSs act, for instance, as surfaces bounding the globally complex but temporally coherent cloud patterns in the atmosphere of Jupiter, as shown in Fig. 1.4.

Instantaneous limits of these material coherent structures are called objective Eulerian coherent structures (OECSs), which represent short-term, approximately material barriers. Unlike LCSs, OECSs are not material and hence can frame the creation, collision, break-up and disappearance of transport barriers in the flow. We will discuss LCSs and OECSs in detail in Chapter 5.

1.4 Transport Barriers in Flow Separation and Attachment

Flow separation involves the detachment of fluid from a rigid boundary, resulting in either the creation of a local recirculation zone (separation bubble) or the global breakaway of fluid particles from the boundary. Highly unsteady flows often display both types of separation, as illustrated by the separation patterns behind the airfoil shown in Fig. 1.5.

https://doi.org/10.1017/9781009225199.002 Published online by Cambridge University Press

Figure 1.4 Enhanced image of cloud patterns in the atmosphere of Jupiter, as seen by the Juno mission in 2020. As we show in later chapters, such patterns are created and shaped by hidden Lagrangian coherent structures (LCSs), such as material jets, eddies and fronts. Image: NASA/JPL-Caltech/SwRI/MSSS/Gerald Eichstädt /Seán Doran.

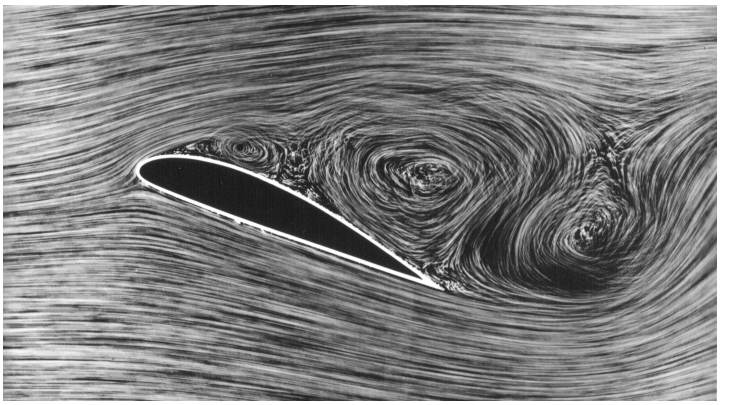

Figure 1.5 Visualization of aerodynamic separation and reattachment along transport barriers attached to an airfoil in a wind tunnel. Image: German Aerospace Center (DLR), CC-BY 3.0.

Separation depletes the kinetic energy content of the flow near the wall, leading to a degradation in the operational performance of engineering devices. Specifically, separation on a bluff body can increase the pressure drag significantly, whereas separation in a diffuser decreases the pressure recovery. Flow attachment is the opposite phenomenon, involving the sustained convergence of fluid elements toward a well-defined boundary location. Local separation necessarily involves a reattachment of the flow that forms the downstream boundary of the separation zone, as seen in Fig. 1.5.

While separation has a clear impact on aerodynamic forces acting on the flow boundary, its experimental detection exploits its impact of fluid particle transport. Indeed, flow visualizations of the type shown in Fig. 1.5 are only possible because particles in the separated and unseparated flow regions are kept apart from each other by material transport barriers. These separation surfaces collect nearby particles along the boundary and eject them into

https://doi.org/10.1017/9781009225199.002 Published online by Cambridge University Press

the mean flow. In contrast, reattachment surfaces guide particles from the mean flow to the wall and then repel them along the boundary.

While typical wall-anchored material surfaces will stretch and fold, separation and reattachment surfaces distinguish themselves by remaining coherent. Therefore, using the terminology introduced in §1.3, we can view separation surfaces as attracting LCSs and attachment surfaces as repelling LCSs.

Unlike general LCSs, however, separation and reattachment surfaces exert weak attraction and repulsion near flow boundaries characterized by no-slip boundary conditions. In dynamical systems terms, these surfaces are nonhyperbolic invariant manifolds, unless the boundary can be characterized with free-slip conditions. This makes most LCS diagnostic tools inefficient in a small neighborhood of a no-slip boundary. At the same time, the direct connection of these LCSs with the boundary facilitates their more detailed local analysis. This in turn enables the identification of separation and reattachment surfaces near the wall solely based on the derivatives of the velocity field along the boundary, without material advection. We will discuss these results for LCSs and OECSs forming separation surfaces and separation spikes in Chapter 6.

1.5 Transport Barriers in Inertial Particle Motion

Patterns formed by finite-size (or inertial) particles carried by fluid flows are often notably different from those seen for fluid particles in dye visualizations. Specifically, while fluid elements in an incompressible fluid cannot exhibit clustering or scattering, both phenomena are well documented in inertial particle motion in incompressible fluids. A spectacular albeit alarming example of inertial clustering is the Great Pacific Garbage Patch, which covers an area of about 1.6 million square kilometers (three times the size of France) (see Fig. 1.6).

Figure 1.6 A transport barrier formed by inertial particles: the Great Pacific Garbage Patch, a floating island of accumulating garbage in the Pacific Ocean. Image: Ocean Cleanup.

https://doi.org/10.1017/9781009225199.002 Published online by Cambridge University Press

Microplastics constitute about 94 % of the roughly 1.8 trillion pieces of plastic in this patch, even though they only make up 8% of the total mass of 80,000 tons.[1] Given the characteristic length scales of the ocean, all objects forming the garbage patch can be considered inertial particles.

In Chapter 7, we will discuss the fundamentals of transport barriers for inertial particle motion. It turns out that these barriers can be uncovered by applying the LCS methods of Chapter 5 to a modification of the carrier fluid velocity field that accounts for inertial effects.

1.6 Barriers to Diffusive and Stochastic Transport

In Chapter 8, we will discuss barriers to the transport of passive tracers whose diffusion cannot be ignored relative to their advective redistribution over the time scales of interest. An example is the long-term mixing of dye in a gently stirred glass of water, as shown in Fig. 1.7. To model such a mixing process accurately, we can no longer set the diffusivity to zero in the advection–diffusion equation governing the spread of the dye.

Figure 1.7 Diffusive mixing of dye framed by transport barriers in a glass of water. Image: Nathan Dumlao at unsplash.com.

One might wonder why we bother defining and identifying transport barriers in diffusive tracer fields, given that they are readily seen in numerical simulations and experiments such as the one shown in Fig. 1.7. To understand our motivation, one must remember that observed surfaces with large concentration drops across them are not necessarily intrinsic barriers to tracer transport. Rather, many of them are remnants of high gradients in the initial tracer distribution. Indeed, the large concentration drop across the boundary of the initial dye drop in the experiment shown in Fig. 1.7 will persist as a slowly diffusing dye-fluid interface for very long times.

Similarly, the lack of large concentration gradients in a given flow domain may simply be due to the concentration being fully mixed or fully absent in that domain. Indeed, there are also transport barriers in the region yet unpenetrated by the dye in the experiment shown in Fig. 1.7. Those barriers will impact other concentration fields (say, the water temperature) with a different initial distribution carried by the same flow.

[1] Most of that mass has been found to be made up of abandoned fishing gear (Lebreton et al., 2018).

https://doi.org/10.1017/9781009225199.002 Published online by Cambridge University Press

Finally, one should consider that while large concentration drops seen in finite-time tracer evolution will often reflect the impact of transport barriers, these drops will not coincide with the barriers themselves. The amount of discrepancy between concentration drop locations and actual transport barriers is generally unclear: convergence of the latter to the former can only be expected in flows with well-defined asymptotic behavior. Identifying intrinsic transport barriers that are free from these idiosyncrasies of a specific observed tracer distribution is essential for universal conclusions valid for the transport of all tracer fields.

A closely related passive transport problem, stochastic transport, is concerned with the transfer of material fluid elements across spatial domains by a velocity field subject to uncertainties. These uncertainties are most often modeled using Brownian motion.

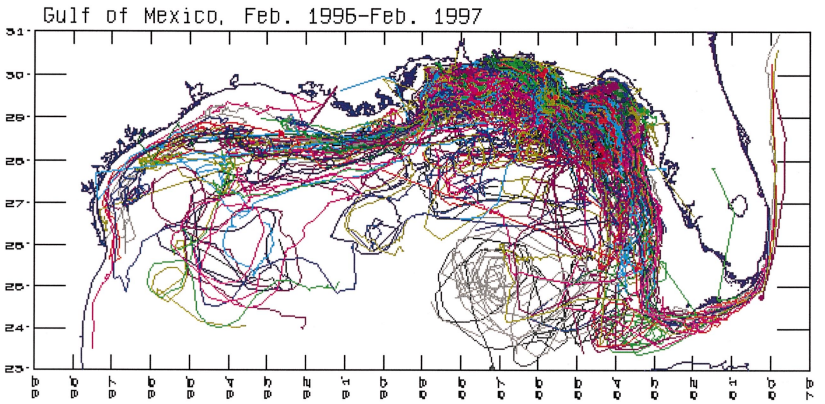

Figure 1.8 Drifter trajectories released and tracked in the Gulf of Mexico over a period of one year. The drifters do not cross an apparent transport barrier acting as the boundary of a *forbidden zone* west off the Florida coastline. (Occasional straight lines indicate captured and transported drifters.) Image: Yang, H. et al. (1999).

For example, drifter trajectories in the Gulf of Mexico, along with an apparent barrier they never cross (see Fig. 1.8), are complex enough to view as trajectories of a random flow.

The connection between trajectory distribution in such a random flow and the evolution of a diffusive scalar field under the deterministic component of the velocity field is given by a classic result: the probability density function of the random flow satisfies the Fokker–Planck equation, which can be recast as an advection–diffusion equation driven by the drift component of the velocity field. Using this connection, we will discuss the definition and properties of barriers to stochastic transport in Chapter 8.

1.7 Barriers to Dynamically Active Transport

In contrast to the passive transport concepts we have discussed so far, active transport involves the transfer of dynamically active tracers. Such active tracers are scalar or vector fields impacting the flow velocity directly rather than simply being carried by it. Active tracers of practical interest include the kinetic energy, vorticity and linear momentum.

https://doi.org/10.1017/9781009225199.002 Published online by Cambridge University Press

Framing the spatiotemporal evolution of these dynamically active fields informally using the concept of transport barriers is common practice in the literature. As a relevant example, Fig. 1.9 shows apparent barriers to vorticity transport in the turbulent flow near a spinning rotor.

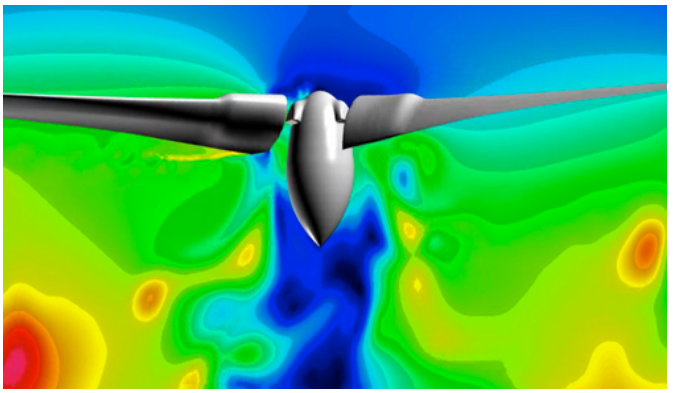

Figure 1.9 Apparent transport barriers in the instantaneous vorticity distribution on a cutting plane for a spinning rotor. Image: Neal M. Chaderjian, NASA Ames Research Center.

Studying dynamically active fields, however, is even more challenging than the transport problems we have already discussed. First, the transport equation for these quantities is a nonlinear partial differential equation (PDE) for the velocity field. For fluid flows, this PDE is the Navier–Stokes equation, whose solution structure in three dimensions is still not fully understood. Second, unlike passive scalar concentration fields, all physically relevant dynamically active fields are nonobjective: they depend on the frame of the observer. This frame dependence is a serious challenge if one wants to conform to the basic requirement of objectivity for transport barriers that we formulated in §1.2.

In Chapter 9, we will discuss how barriers to dynamically active transport can also be characterized and located in an observer-independent way. This in turn makes the experimental visualization of these barriers via material tracers feasible in practice.

1.8 Coherent Sets, Coherence Clusters and Coherent States

We close this Introduction by mentioning three notions of coherence that we will *not* be discussing in further detail here. They are beyond the scope of this book because their primary focus is not transport barriers.

The first such notion is that of a *coherent set*: an equivalence class of trajectories that stay closer to each other during their evolution than to other trajectories. Specifically, finite-time coherent sets comprise trajectories that disperse slower than others over a given finite time interval of interest (see, e.g., Froyland et al., 2010; Froyland, 2013; Bollt and Santitissadeekorn, 2013). This view of coherence focuses on sets enclosed by closed transport barriers as opposed to the barriers themselves. A coherent set is defined as a region of initial conditions that continues to have significant overlap with its deterministically advected final position even if a small random perturbation (or diffusion) is added to its originally deterministic advection.

https://doi.org/10.1017/9781009225199.002 Published online by Cambridge University Press

Available approaches seek to locate such coherent sets based on properties of the transfer operator (or Frobenius–Perron operator), which maps passive, scalar-valued functions defined over initial positions of fluid particles into the evaluations of the same functions over later positions of those particles.[2] More recent reformulations and extensions of the transfer operator approach are given by Froyland (2015); Froyland and Kwok (2017); Froyland et al. (2020) and the references cited therein.

The transfer operator can be approximated by its finite-dimensional discretization **P**, constructed as a matrix of transition probabilities within a finite partition of the flow domain. The second (left) singular vector of **P** is then expected to characterize a dominant coherent set after appropriate thresholding. Further coherent structures are expected to appear from the appropriate thresholding of higher singular vectors of **P**. The number of singular vectors to consider and the applied thresholding are typically determined empirically by inspection of the unprocessed singular vectors of **P**. Transport barriers are then inferred by implication as boundaries of the coherent sets obtained in this fashion.

A notion related to coherence sets is the coherence cluster, which comprises a group of trajectories that are more similar to each other than to other groups. Methods for identifying clusters in data sets were originally developed in the computer science literature (see, e.g., Everitt et al., 2011). Using such ideas, Froyland and Padberg-Gehle (2015) use fuzzy C-means clustering to identify clusters of initial conditions that remain close to the same virtual cluster center. Hadjighasem et al. (2016) seek clusters as sets of trajectories that maintain small relative distances to each other. They note that clusters can be identified by applying the technique of spectral clustering to the eigenvectors of matrix (the graph Laplacian) that encodes the average distances among trajectories. Schlueter-Kuck and Dabiri (2017) use the same clustering approach but with a different distance function. Banisch and Koltai (2017) apply the spectral clustering approach in the transfer operator framework.

These spectral methods targeting coherent sets or coherence clusters are different in spirit from the geometric approaches we will be surveying in this book. First, we will focus on barriers minimizing advective, diffusive, stochastic or active transport of a physical quantity without requiring the barriers to be closed. In contrast, barrier surfaces inferred from spectral methods have to be closed. As a consequence, spectral methods are inapplicable to some of the most frequently observed barriers, such as jets and fronts in the ocean, just because they happen to be open. Second, spectral approaches require empirical user input in the selection of the relevant singular vectors, which is not the case for the geometric methods we will survey. To illustrate the differences between probabilistic and geometric approaches on a specific flow, Fig. 1.10 shows a comparison of coherent sets obtained from transfer operator methods with transport barriers obtained from LCS methods in a two-dimensional, spatially double-periodic turbulence simulation.

A third frequently used notion of coherence in contemporary fluid mechanics is that of an exact coherent state[3] (or ECS), which refers to an exact, nonlinear solution of the Navier–Stokes equation in canonical shear flows, such as plane Couette, Poiseuille and pipe flows (Waleffe, 1998, Graham and Floryan, 2021). These solutions display simple spatiotemporal

[2] Technically speaking, the transfer operator is the pushforward operation carried out by the flow on measurable functions.

[3] Exact coherent states are sometimes called exact coherent structures, causing occasional confusion with the coherent structures discussed in §1.3 and elsewhere in this book.

https://doi.org/10.1017/9781009225199.002 Published online by Cambridge University Press

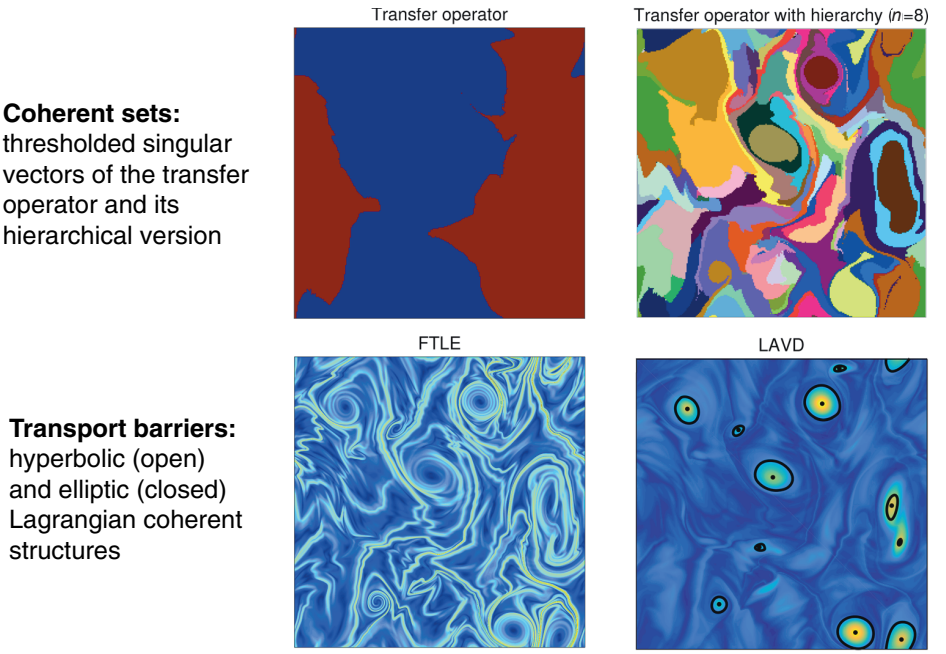

Figure 1.10 Coherent sets obtained as thresholded second singular vectors of the transfer operator approach and from its hierarchical version, which applies the original algorithm again within each coherent set recursively n times (Ma and Bollt, 2013). Also shown are transport barriers obtained as Lagrangian coherent structures (LCSs) from the finite-time Lyapunov exponent (FTLE) and from the Lagrangian-averaged vorticity deviation (LAVD). The FTLE highlights hyperbolic LCSs and the LAVD highlights elliptic LCSs, as we will discuss in Chapter 5. Adapted from Hadjighasem et al. (2017).

behavior that is nevertheless reminiscent of coherent features of turbulent solutions that arise recurrently. This is because ECSs are of saddle type and hence other time-evolving Navier–Stokes velocity fields may, from time to time, pass by ECSs. These passing solutions will mimic some of the spatial features of the approached ECSs for a while before departing. Examples of ECSs include steady-state, time-periodic or traveling-wave solutions, such as the one shown in Fig. 1.11.

Features of ECSs have traditionally been assessed via heuristic visualization tools, such as instantaneous level surfaces of a velocity component (see, e.g., Fig. 1.11). This visualization approach is useful in a direct comparison between the features of ECS velocity fields and those of turbulent velocity fields. However, the surfaces obtained in this fashion tend to differ from experimentally observable material coherent structures even in two-dimensional steady flows (see §3.7.1 for examples). In summary, while the transport barriers and coherent structures we discuss here are special material surfaces of a given velocity field, the ECSs are special velocity fields themselves. Their particular transport barriers and coherent structures can nevertheless be studied by the methods discussed in the upcoming chapters of this book.

https://doi.org/10.1017/9781009225199.002 Published online by Cambridge University Press

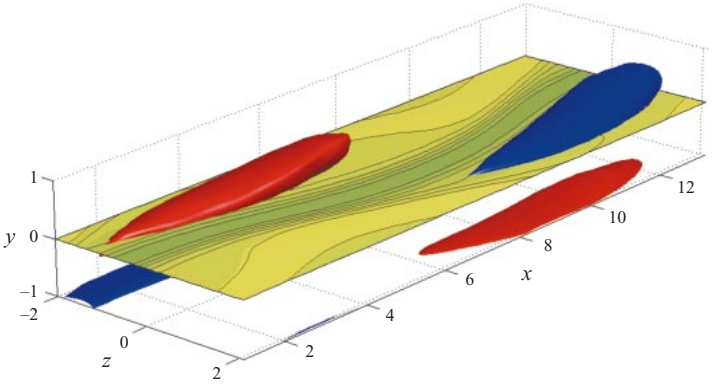

Figure 1.11 Exact coherent state (ECS) in a channel flow moving in the positive x-direction of a channel at Reynolds number $Re = 376$. Level curves of the streamwise velocity in the $y = 0$ plane are shown together with isosurfaces of the streamwise vorticity. Positive vorticity is shown in blue, negative in red. Adapted from Waleffe (2001).

https://doi.org/10.1017/9781009225199.002 Published online by Cambridge University Press