

A physics-oriented review of neural-network-based modeling for bluff-body flows

Downloaded from: https://research.chalmers.se, 2026-11-08 16:10 UTC

Citation for the original published paper (version of record):
Tang, X., Zhang, J., Zhu, H. (2026). A physics-oriented review of neural-network-based modeling for bluff-body flows. Physics of Fluids, 38(10). http://dx.doi.org/10.1063/5.0351690

N.B. When citing this work, cite the original published paper.

research.chalmers.se offers the possibility of retrieving research publications produced at Chalmers University of Technology. It covers all kind of research output: articles, dissertations, conference papers, reports etc. since 2004. research.chalmers.se is administrated and maintained by Chalmers Library

(article starts on next page)

Physics of Fluids

REVIEW ARTICLE | OCTOBER 05 2026

A physics-oriented review of neural-network-based modeling for bluff-body flows

Xin Tang (汤鑫) ; Jingnan Zhang (张竞楠) ; Heng Zhu (朱恒)

Check for updates

Physics of Fluids 38, 101302 (2026)
https://doi.org/10.1063/5.0351690

View Online
Export Citation

Articles You May Be Interested In

A review of experiments on stationary bluff body wakes
Physics of Fluids (January 2022)

Review on piezoelectric wind-excited vortex-induced vibration energy harvesting technologies
J. Renewable Sustainable Energy (September 2026)

A survey of flow-based energy harvesters for powering sustainable wireless sensor nodes
J. Renewable Sustainable Energy (January 2025)

13 October 2026 05:32:48

AIP Publishing
Physics of Fluids
Special Topics Open for Submissions
Learn More

AIP Publishing

Physics of Fluids REVIEW pubs.aip.org/aip/pof

A physics-oriented review of neural-network-based modeling for bluff-body flows

Cite as: Phys. Fluids **38**, 101302 (2026); doi: 10.1063/5.0351690
Submitted: 8 July 2026 · Accepted: 17 September 2026 ·
Published Online: 5 October 2026

View Online Export Citation CrossMark

Xin Tang (汤鑫),[1,2] Jingnan Zhang (张竞楠),[3] and Heng Zhu (朱恒)[4,a)]

AFFILIATIONS

[1]College of Science, Mathematics and Technology, Wenzhou-Kean University, Wenzhou, Zhejiang 325060, People's Republic of China
[2]The Dorothy and George Hennings College of Science, Mathematics and Technology, Kean University, Union, New Jersey 07083, USA
[3]Division of Glycoscience, Department of Chemistry, KTH Royal Institute of Technology, AlbaNova University Center, Roslagstullbacken 21, 114 21 Stockholm, Sweden
[4]Department of Mechanical Engineering, Division of Marine Technology, Chalmers University of Technology, Gothenburg 41296, Sweden

[a)]Author to whom correspondence should be addressed: heng.zhu@chalmers.se

ABSTRACT

Bluff-body flows combine early separation, coherent wake organization, vortex shedding, and strong sensitivity to geometry, Reynolds number, inflow condition, and three-dimensional transition. These features make them both canonical fluid-mechanics benchmarks and demanding targets for neural-network-based modeling. This review provides a representative, physics-oriented synthesis of neural-network approaches for bluff-body flows. Instead of organizing the field by architecture or prediction task, the review asks which wake mechanisms current models represent and which validation evidence makes those claims trustworthy. It examines how current models represent separation, wake topology, pressure-wake coupling, coherent shedding, turbulence statistics, closure behavior, and flow-control mechanisms. We discuss data sources, flow representations, model classes, validation practices, and practical uses in surrogate prediction, reconstruction, monitoring, control, and digital-twin-type workflows. The review emphasizes that low numerical error by itself does not establish physical learning, and argues for validation based on topology, spectra, force response, uncertainty, generalization, and interpretable mechanism evidence.

© 2026 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/). https://doi.org/10.1063/5.0351690

13 October 2026 05:32:48

I. INTRODUCTION

Bluff-body flows remain one of the most canonical and persistently challenging problems in fluid mechanics. Despite the apparent geometric simplicity of typical configurations, such as circular and noncircular cylinders, prisms, and other block-like bodies immersed in crossflow, the resulting dynamics involve boundary-layer separation,[1–3] free-shear-layer instability,[4–6] unsteady wake formation,[7,8] and a broad range of coherent and turbulent structures.[9–11] These flows have therefore long served as benchmark systems for theoretical analysis, experimental investigation, and numerical simulation, from the classical studies to modern investigations of wake instability, three-dimensional transition, and turbulence.[12–18] Their enduring importance lies in the richness of the underlying physics and in the coexistence of robust large-scale organization with strong local variability, which makes predictive modeling intrinsically nontrivial.

From an engineering perspective, bluff-body flows are central to a wide range of practical problems, including pressure-drag-dominated aerodynamics,[19,20] vortex-induced vibration of marine and offshore structures,[21,22] aeroacoustics associated with unsteady separation,[23–26] and wind loading on buildings and industrial installations.[27–29] In many of these applications, models must predict global quantities such as drag, lift, and shedding frequency together with local flow features such as separation behavior, recirculation structure, and wake development. However, as the objective shifts from single-condition analysis to repeated prediction over broad parameter spaces, traditional approaches face clear limitations.[13,14,18,27,30,31] Experimental methods provide valuable physical realism, but are

© Author(s) 2026

constrained by trade-offs among spatial resolution, temporal resolution, accessible measurement volume, and cost.[12,32–34] High-fidelity numerical simulations offer detailed flow information, yet direct numerical simulation (DNS) is prohibitively expensive in most practically relevant regimes, while large-eddy simulation (LES) and unsteady Reynolds-averaged Navier–Stokes (URANS) approaches remain limited by computational cost and model dependence.[35–37] These difficulties become especially pronounced in contexts such as design optimization, uncertainty quantification, flow monitoring, and control, where many-query or near-real-time prediction is required.[33,37,38]

Against this background, neural-network-based methods have attracted growing attention because they can approximate high-dimensional nonlinear mappings[39] and can therefore support reduced-order modeling, fast surrogate prediction, field reconstruction from sparse measurements, and flow-aware control.[38,40] Their appeal in bluff-body flow problems is not accidental. Bluff-body wakes often exhibit recurring coherent structures, dominant frequencies, and partially low-dimensional spatiotemporal organization, especially in canonical laminar or weakly three-dimensional regimes such as periodic vortex shedding.[12] These features make such flows promising candidates for data-driven compression and learning. At the same time, bluff-body flows also expose fundamental limits of learning-based approaches, because the apparent regularity of the wake can break down with changes in Reynolds number (Re), geometry, inflow conditions, and three-dimensional instability.[41] In this sense, bluff-body flows provide a useful application area for neural networks and a demanding testbed for examining what such models truly learn from complex unsteady flows.[38,40,42]

This combination of learnability and regime sensitivity makes bluff-body wakes particularly informative for assessing machine-learning models. Coherent shedding and partially low-dimensional dynamics make prediction and reconstruction feasible, yet the same organization can make narrow interpolation look more physically meaningful than it is.[38,43–45] The learnability of canonical bluff-body wakes is also consistent with the manifold hypothesis: although flow snapshots are high-dimensional, coherent wake dynamics may occupy a substantially lower-dimensional state space, which motivates both linear reduced-order methods and nonlinear latent representations.[46–52] This low-dimensional structure, however, can change with transition, three-dimensionality, turbulence, and geometry, limiting transfer across wake regimes. This review therefore evaluates models by the wake mechanisms and physical relationships they preserve: separation and recirculation, shedding dynamics, pressure-force coupling, three-dimensional transition, turbulent transport, closure behavior, and control response, instead of organizing the evidence primarily around architecture labels or scalar errors.

Recent surveys have organized the field around experimental workflows,[53] operator learning and scientific simulation,[54] turbulence modeling, or control. This review instead synthesizes the bluff-body literature around wake mechanisms, validity domains, and the strength of physical evidence. Accordingly, this scope covers three-dimensional and turbulent wakes, learned closures, sparse reconstruction, control, and practical workflows, while keeping canonical cylinder shedding as a diagnostic baseline. In this review, physical learning refers to the preservation of wake-relevant relationships among inputs, states, and outputs; it does not imply that the network has recovered the Navier–Stokes equations or an underlying causal mechanism. Evidence becomes stronger as validation progresses from in-support field accuracy to topology, dynamics and spectra, turbulence or closure behavior, held-out regime transfer, and interpretable mechanism-level response.

The review first identifies the wake features that support or limit learnability, then assesses data sources, model formulations, and evidence quality. It subsequently examines how strongly current models support mechanism-level claims across different wake regimes and practical uses. The resulting framework is intended for prediction, reconstruction, monitoring, closure, control, and digital-twin-type applications within explicitly stated validity domains.

In this context, the review advances three connected arguments: wake mechanisms should lead the assessment; current evidence is strongest in canonical, data-rich coherent regimes; and trustworthy use requires validation of wake topology, spectra, force response, uncertainty, regime transfer, and mechanism-relevant behavior, in addition to field accuracy.

II. PHYSICAL CHARACTERISTICS OF BLUFF-BODY FLOWS

This section briefly outlines the physical characteristics of bluff-body flows. The aim is not to provide a comprehensive survey of bluff-body wake physics, but to establish the key mechanisms that dominate these flows, including separation, wake development, coherent shedding, three-dimensional instability, and regime sensitivity. These mechanisms provide the physical basis for the later discussion of neural-network formulations and, more importantly, for the critical assessment in Sec. IV of how well existing neural-network models capture the underlying wake physics. Figure 1 summarizes the wake-physics elements used throughout this review.

A. Separation and canonical wake organization

The defining feature of bluff-body flows is early flow separation. On rounded bodies, adverse-pressure-gradient separation and boundary-layer state strongly influence where the flow separates; on sharp-edged bodies, geometry can impose separation directly, as illustrated in Figs. 1(a) and 1(f). In both cases, a separated wake forms behind the body and pressure drag becomes a dominant contribution to the total aerodynamic force.[68,69] This basic picture has been established in the classical circular-cylinder literature and remains central to modern reviews of bluff-body wakes.[12,70] Recent studies have emphasized that even this seemingly canonical picture is often modified by geometric irregularity, controlled laminar–turbulent transition, finite-body effects, and external constraints such as blockage and wall proximity, all of which can reshape the near-wake structure and alter the apparent universality of separated-wake organization.[12,61–63]

Once separation occurs, the near wake contains detached shear layers, recirculation and vortex-formation regions, and a progressively recovering far wake, as illustrated in Figs. 1(a) and 1(b); its topology depends on geometry, Re, and confinement.[11,12,17,46,61] Recent reviews have summarized how blockage and confinement modify wake development for circular cylinders and broader bluff-body classes.[61,71] Separation location, recirculation length, base pressure, wake width, shear-layer thickness, vortex-formation length, and force spectra are therefore more informative validation targets than image similarity alone.

13 October 2026 05:32:48

© Author(s) 2026

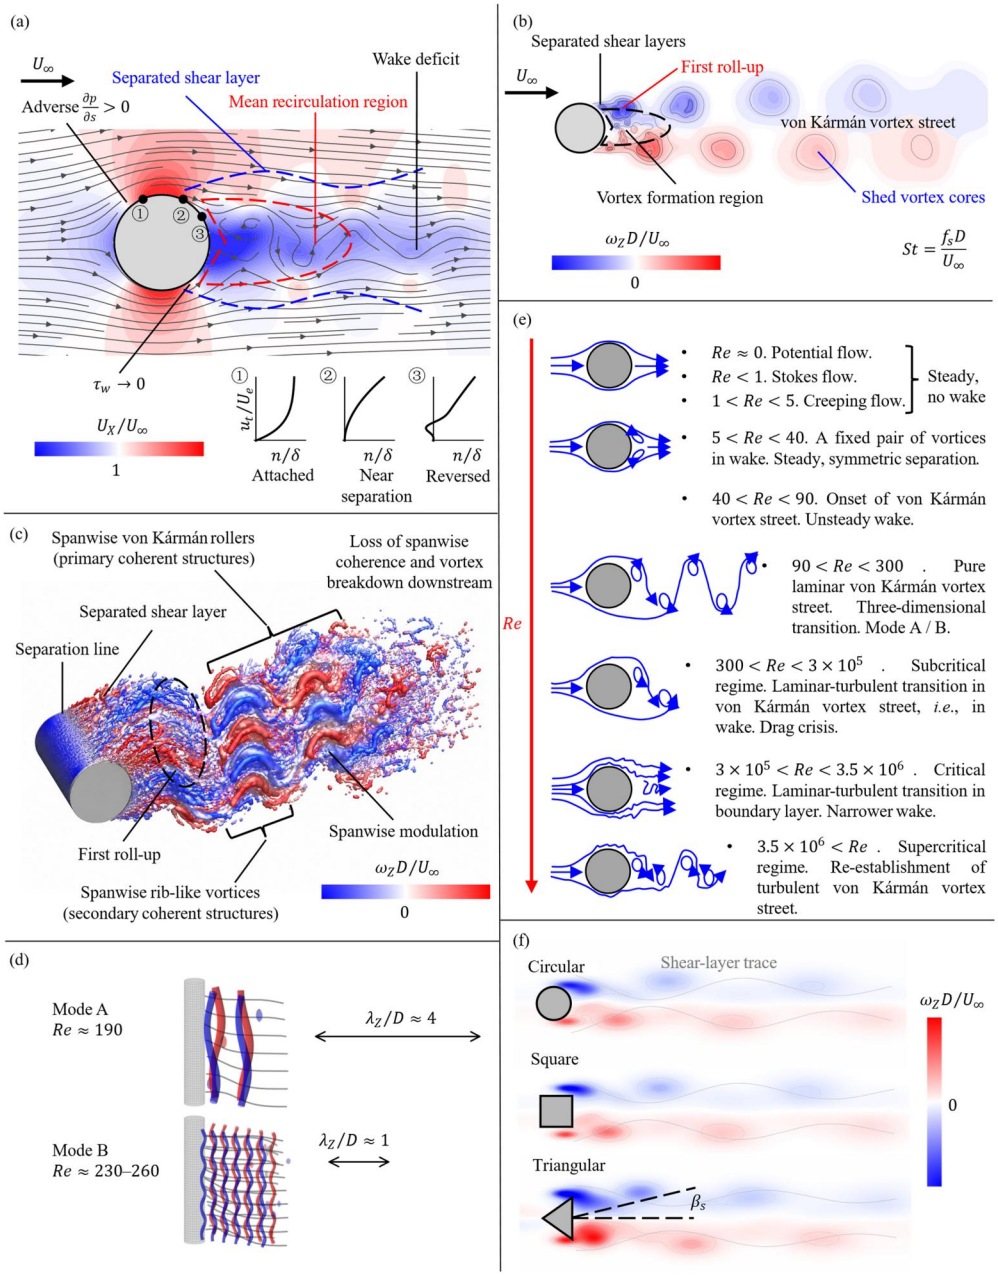


FIG. 1. Flow-physics features of bluff-body wakes relevant to neural-network modeling. (a) Instantaneous two-dimensional near-cylinder flow around a circular cylinder at $Re \sim 10^4$, showing boundary-layer development, flow separation, separated-shear-layer roll-up, and the associated wake deficit. The velocity profiles below the contour illustrate the progression from an attached boundary layer to near-separation and locally reversed flow.[1–6] (b) Instantaneous planar $Q > 0$ wake structures for a circular-cylinder configuration at $Re \sim 10^4$. The contours identify vortical regions in the formation region and downstream von Kármán-type wake, with red and blue indicating the sign of the spanwise vorticity ω_Z.[4–8,17] (c) Schematic three-dimensional $Q = Q_0$ iso-surfaces in the wake of a circular cylinder at $Re \sim 10^4$, illustrating spanwise rollers, streamwise rib-like connections, smaller formation-region structures, and downstream loss of coherence.[12,17,26,55–59] (d) Canonical spanwise instability modes in circular-cylinder wakes: Mode A occurs near $Re \approx 190$ with a characteristic spanwise wavelength $\lambda_Z/D \approx 4$, whereas Mode B occurs near $Re \approx 230$–260 with $\lambda_Z/D \approx 1$.[46,47,55–58,60] (e) *Re*-dependent wake regimes for a smooth circular cylinder, from creeping or attached flow, steady separation, two-dimensional vortex shedding, and three-dimensional transition to subcritical shear-layer transition, critical drag-crisis behavior, and supercritical/transcritical shedding. The indicated *Re* bands are approximate and can shift with surface roughness, freestream turbulence, blockage, aspect ratio, and end conditions.[4,12,17,61–63] (f) Effect of bluff-body cross-section geometry on signed spanwise vorticity and wake organization. Circular, square, and triangular cylinders produce different separation locations, near-wake formation lengths, wake widths, and vortex-efflux directions; for the triangular case, β_s denotes the vortex efflux angle. The color map shows $\omega_Z D/U_\infty$. Thin gray curves indicate schematic shear-layer or wake traces only and do not represent an additional scalar field.[12,17,64–67]

13 October 2026 05:32:48

© Author(s) 2026

The detailed separation process differs substantially between rounded and sharp-edged bluff bodies. For a circular cylinder, the separation location is not fixed geometrically and depends on *Re* and boundary-layer state, whereas for square and rectangular cylinders, separation is often imposed directly by the sharp leading edges; side-surface reattachment and trailing-edge shedding can then depend on aspect ratio, incidence, corner radius, and confinement.[12] In such cases, the downstream evolution of the free shear layers, including possible reattachment and re-separation, becomes especially important in determining the wake structure. This distinction is well documented for rectangular-cylinder flows, where the balance between fixed-edge separation, side-surface reattachment, and wake redevelopment strongly affects the overall topology.[11,12,17] Thus, although bluff-body wakes share a common separated-wake framework, they are not physically identical systems. Their organization is shaped jointly by separation mechanism, body geometry, and downstream shear-layer evolution. The geometry dependence illustrated in Fig. 1(f) is therefore not a cosmetic variation of the same wake image, but a change in the separation constraint imposed on the learning problem.

B. Vortex shedding and dominant coherent wake dynamics

Alternating vortex shedding emerges when the two separated shear layers interact in the near wake and roll up into large-scale vortical structures, producing the classical Bénard–von Kármán vortex street that organizes the unsteady wake of many canonical bluff bodies. This mechanism has been documented extensively for the circular-cylinder wake and remains the standard physical picture for the onset and maintenance of coherent wake oscillation.[60,70] This local roll-up should not be equated with the full frequency-selection mechanism, because alternating shedding is a global wake oscillation governed by the coupled interaction of the separated shear layers and the recirculating near wake.[46,60,70] Its dominant frequency is commonly expressed by the Strouhal number, $St = \frac{f_s D}{U_\infty}$, where f_s is shedding frequency, D is a characteristic body dimension, and U_∞ is freestream velocity.[46,60,72,73] For circular-cylinder wakes, several empirical St–Re correlations have been proposed over restricted Re ranges, including relations of the form $St = St^* + \frac{m}{Re}$, where St^* and m are fitted coefficients, as well as expansions in $\frac{1}{\sqrt{Re}}$.[74,75] Their coefficients and validity ranges vary with the shedding regime and experimental conditions. The resulting St depends on geometry, Re, blockage, and boundary conditions and should therefore be treated as a regime-dependent wake descriptor rather than a universal constant or a complete test of the frequency-selection mechanism. In canonical cylinder wakes, vortex shedding is also the principal source of periodic transverse loading and thus directly underlies the fluctuating lift that is central to vortex-induced vibration and related unsteady force phenomena.[73,76] Recent studies of nonidealized cylinder configurations further show that twist, inclination, tapering, and finite-span effects can modify the spatial coherence, spanwise organization, and modal content of the nominal von Kármán mode without necessarily eliminating vortex shedding altogether.[77–79] The coherent wake pattern in Fig. 1(b) explains why cylinder wakes are attractive for sequence models and reduced-order neural representations, but it also indicates why frequency, phase, and force spectra must be reported rather than only pointwise field error.

This organized dynamics is most evident in isolated-body flows under steady inflow, but it does not persist unchanged as the flow becomes more complex. In the circular-cylinder wake, the well-known transition from nominally two-dimensional periodic shedding to three-dimensional instability, including Mode A and later Mode B, marks the beginning of a progressive loss of simple wake coherence.[47,80] Recent studies on sharp-edged and multi-body bluff-body systems likewise suggested that the onset of vortex shedding, hysteresis between wake states, and mode competition can depend sensitively on cross-sectional geometry and body interaction, reinforcing that coherent shedding is a robust but not unique wake state.[66,67] More broadly, recent work continues to emphasize that practical bluff-body wakes often sit in regimes where the primary vortex street coexists with three-dimensionality, turbulence, or far-wake reorganization, so the classical coherent picture remains fundamental but not exhaustive.[81,82]

C. Three-dimensionality and breakdown of coherence

The nearly two-dimensional periodic picture applies only over a limited envelope. In low-disturbance, spanwise-periodic circular-cylinder flow, increasing *Re* leads first to Mode A and subsequently to the shorter-wavelength Mode B instability, as summarized in Figs. 1(c)–1(e).[55–58] These transition ranges are canonical guides rather than universal thresholds because roughness, freestream turbulence, blockage, aspect ratio, end conditions, and boundary-layer state can shift their onset.[83]

Three-dimensional transition changes the wake state qualitatively rather than merely adding spanwise variation to an otherwise two-dimensional vortex street. Spanwise modulation, streamwise rib-like vortices, secondary instability, modal competition, and intermittent loss of coherence become intrinsic to the wake dynamics. The dominant shedding mode may remain identifiable, but its spatial coherence and spectral content are no longer represented adequately by a single planar phase or two-dimensional field.

This change has direct implications for neural-network modeling. A model trained only on planar or weakly three-dimensional wakes may reproduce the two-dimensional attractor accurately while lacking the state variables required to represent spanwise instability, end effects, or turbulent transport. More complex bluff bodies further introduce reattachment, corner, wall, and finite-span effects that reorganize wake topology.[11,77–79,82,84–86] *Re*, geometry, and spanwise boundary conditions should therefore be treated as regime-defining variables rather than merely scalar inputs.

III. APPLIED NEURAL-NETWORK FORMULATIONS

A training dataset is a sample of physical regimes, not a generic image or time-series collection. Its intended validity domain should report *Re*; geometry, angle, spacing, and blockage; inflow turbulence and boundary-layer state; aspect ratio, span, walls, and end conditions; solver, turbulence model, grid, domain, sampling, and averaging; sensor layout and calibration; and fidelity, uncertainty, and provenance.[35,37,53,87–90] These variables determine which wake relationships a model can plausibly preserve and which changes constitute extrapolation.

Neural-network studies of bluff-body flows can be viewed from two complementary perspectives. Table I organizes representative studies by the wake mechanism or physical modeling claim they address, whereas Table II classifies them by neural-network

13 October 2026 05:32:48

© Author(s) 2026

TABLE I. Physics-oriented classification of neural-network studies for bluff-body flows. Evidence-scope symbols: (i) direct bluff-body wake evidence; (ii) closely related separated-flow or partial bluff-body evidence; (iii) broader fluid-mechanics or scientific-machine-learning methodological support. Abbreviations: LES, large-eddy simulation; RANS, Reynolds-averaged Navier–Stokes; SGS, subgrid-scale; PIV, particle image velocimetry; PINN, physics-informed neural network.

Mechanism	Context	Learning target	Explicit physics	Common limitation	Examples
Separation and near-wake organization	Circular cylinders, rectangular cylinders, multiple bluff bodies	Velocity/pressure fields in separated wakes; recirculation-zone structure; near-wake topology	Separated shear layers, recirculation, wake recovery	Separation onset, boundary-layer state, separation angle/location, base pressure, recirculation length, and shear-layer thickness are often not directly tested	Refs. 42 and 91–93 (i)
Vortex shedding and coherent wake dynamics	2D circular-cylinder wakes, moderate *Re*	Snapshot-to-snapshot wake prediction; phase evolution; reduced-order latent dynamics	Periodic shedding, coherent vortical structures, temporal wake evolution	Long-term phase drift and off-design generalization remain difficult	Refs. 43–45, 92, 94, and 95 (i)
Sparse-to-dense and non-time-resolved wake reconstruction	PIV-based cylinder wakes; limited sensors	Recovering dense or time-resolved velocity fields from sparse/non-time-resolved data	Wake periodicity, spatial coherence, temporal sequencing	Observability, pressure/force consistency, and nonunique reconstruction are often secondary	Refs. 96–102 (i)
Pressure-wake observability	Cylinder wakes with wall-pressure input	Inferring near-wake velocity fields from surface pressure	Wall pressure, vortex formation, unsteady loading	Uniqueness and robustness decrease in turbulent or multi-state wakes	Refs. 99 and 102 (i)
Three-dimensional and turbulent wake complexity	3D bluff-body geometries; high-Re cylinder wakes; turbulent-region identification	Reconstructing 3D or high-*Re* wake fields; identifying turbulent regions; representing spanwise or turbulent wake features	Three-dimensional geometry, spanwise organization, turbulent/nonturbulent interface, instantaneous/mean/phase-averaged fields	Direct 3D-wake evidence, high-*Re* reconstruction, and turbulence classification are often tested in separate regimes	Refs. 91, 98, 101, 103, and 104 (i/ii)
Turbulence closure and unresolved-scale effects	LES/RANS of bluff-body wakes	Learning SGS stresses or improving coarse-grid LES/RANS predictions	Mean flow, drag, resolved Reynolds stress, SGS stress	Learned closure may remain difficult to interpret physically; a posteriori solver-coupled behavior is essential	Refs. 105–108 (i/ii)
Flow control and wake stabilization	Cylinder wake with jets or blowing/suction	Learning actuation policies to suppress shedding and reduce drag	von Kármán vortex street, drag reduction, lift fluctuation control	Robustness to *Re*, sensor noise, time delay, actuator limits, actuation cost, and 2D-to-3D transfer must be tested	Refs. 109–116 (i)
Compressible/transonic bluff-body flow	High-Re transonic flow around cylinder	PINN-based prediction of compressible steady aerodynamic fields	Compressibility, boundary layer, wake, RANS/Euler constraints	Unsteady vortex shedding is not the central target	Ref. 41 (i)
Geometry dependence and multi-body interaction	Multiple bluff bodies, variable spacing/angle	Fast prediction across arrangements and geometries	Interference effects, geometry-conditioned wake fields	Mechanistic decomposition of interference modes is limited	Refs. 91, 93, 117, and 118 (i/ii)

© Author(s) 2026
13 October 2026 05:32:48

TABLE II. Neural-network-method classification. The evidence-scope symbols are the same as in Table I: (i) direct bluff-body wake evidence; (ii) closely related separated-flow or partial bluff-body evidence; (iii) broader fluid-mechanics or scientific-machine-learning methodological support. Abbreviations: CNN, convolutional neural network; RNN, recurrent neural network; DNN, deep neural network; FCNN, fully connected neural network; LSTM, long short-term memory; DMD, dynamic mode decomposition; SINDy, sparse identification of nonlinear dynamics; SDF, signed-distance function; PINN, physics-informed neural network; PDE, partial differential equation; DRL, deep reinforcement learning; PPO, proximal policy optimization; ANN, artificial neural network.

Method category	Typical architecture	Main input	Main output	Main bluff-body use	Validation practice	Examples
Snapshot-to-snapshot flow prediction	CNN, U-Net-type CNN, knowledge-enhanced CNN	Previous velocity/pressure fields; sometimes geometry/SDF/body-mask channels	Future velocity/pressure fields	Fast surrogate prediction of unsteady wake evolution	Comparison with CFD fields; error growth over recursive prediction; sometimes unseen *Re*	Refs. 42, 43, 92, and 95 (i)
Super-resolution, generative reconstruction, and resolution enhancement	CNN, encoder–decoder, adversarial/generative model, diffusion model	Low-resolution fields; sparse observations; PTV-like particle data; partial velocity/vorticity fields	High-resolution or reconstructed flow fields; possible ensembles of plausible fields	Spatial/spatiotemporal resolution enhancement and sparse reconstruction, including cylinder-wake tests	High-resolution reference comparison; spectra; phase alignment; statistics; uncertainty calibration and ensemble spread when probabilistic	Ref. 101 (i); Ref. 119 (i/iii); Refs. 120–122 (iii)
Sparse-sensor/PIV / sectional-field reconstruction	RNN/CNN-sequential, spatio-temporal DNN, graph/transformer models	Non-time-resolved PIV, sparse sensors, sectional fields, pressure probes, partial fields	Time-resolved or dense velocity fields	Experimental wake reconstruction	Robustness to sensor noise, sensor removal, layout changes, time delay, and unseen wake states	Refs. 96–99, 101, 102, 104, and 123 (i); Ref. 100 (ii)
Reduced-order temporal modeling	CNN autoencoder + LSTM; SINDy/latent low-dimensional models; DMD + CNN-LSTM variants	Flow snapshots, modal coefficients, latent variables	Future latent states and reconstructed fields	Learning low-dimensional vortex-shedding dynamics	Reconstruction error, temporal prediction accuracy, comparison across Re/cases	Refs. 44, 45, and 95 (i)
Pressure-to-field mapping	CNN encoder–decoder/pressure-input networks	Surface pressure histories	Near-wake velocity field	Wake inference from wall measurements	Comparison with numerical velocity fields	Ref. 99 (i)
Physics-informed neural networks	PINN, mixed-variable PINN, physics-informed DL	Coordinates, boundary conditions, sparse data, PDE residuals	Velocity, pressure, density, compressible variables	Physics-constrained field solution or reconstruction	Residual balance; boundary-condition sensitivity; pressure and force consistency; solver-coupled checks where possible	Refs. 41 and 108 (i); Refs. 124 and 125 (ii)
Knowledge-enhanced / conservation-	CNN with conservation/physics-informed loss	Flow fields, geometry representation, prior time step	Unsteady velocity/pressure fields	Embedding physical constraints into data-driven	Comparison with CFD; ablation of knowledge	Refs. 42 and 43 (i)

© Author(s) 2026
13 October 2026 05:32:48

TABLE II. (*Continued.*)

Method category	Typical architecture	Main input	Main output	Main bluff-body use	Validation practice	Examples
informed CNN Neural turbulence closure	DNN, FCNN, neural SGS/RANS closure models	Resolved LES/RANS variables, strain rate, velocity gradient	SGS stress or closure correction	prediction LES/RANS improvement for separated wakes	integration A priori error; a posteriori stability; mean wake, base pressure, drag, Reynolds stresses, and spectra	Refs. 105, 106, and 108 (i); Ref. 107 (ii)
Reinforcement-learning control	DRL, PPO-based ANN	Flow-state probes, pressure/velocity observations	Jet/blowing-suction actuation	Active wake control and drag reduction	Drag/lift reduction, wake visualization, robustness over *Re*	Refs. 109–116 (i)
Geometry-conditioned surrogate modeling	FCNN, fully convolutional NN, point-cloud/geometry-based DL	Shape description, SDF, spacing, angle of attack, geometry channels	Flow fields or aerodynamic coefficients	Rapid design-oriented prediction	Testing must distinguish interpolation within a trained shape family from extrapolation to new bodies or wake regimes	Refs. 91, 93, 117, and 118 (i/ii)

formulation and input–output structure. Together, the two tables distinguish what physical content is being modeled from how the learning problem is constructed.

A. Data sources and representations

1. Simulation and experimental datasets

Most neural-network studies of bluff-body flows rely on simulation data, especially for field prediction, reduced-order modeling, and closure learning. Numerical data provide complete access to velocity, pressure, vorticity, surface forces, and derived quantities at controlled *Re* and geometries. This makes simulation data particularly suitable for supervised learning, where the input–output mapping can be defined explicitly. For example, deep learning models have been trained to predict future velocity and pressure fields around a circular cylinder from previous snapshots,[43] to build reduced-order models for two-dimensional unsteady cylinder wakes,[45] and to infer near-wake velocity fields from instantaneous wall-pressure input.[99] In these cases, the availability of dense flow-field data allows the network to learn spatiotemporal correlations that would be difficult to obtain from sparse experiments alone.

Simulation data, however, are not neutral representations of the physical flow. The predicted wake can depend on the numerical method, grid and domain resolution, boundary conditions, turbulence treatment, sampling strategy, and averaging interval. These sensitivities are particularly relevant to neural-network training because the learned model may reproduce numerical or closure-specific features of the data generator rather than flow physics that is independent of that setup. For example, numerical studies of the circular-cylinder wake around $Re = 3900$ have shown that spanwise domain extent, mesh resolution, confinement, and averaging choices can materially affect predicted wake and turbulence statistics.[126,127] Similar sensitivity has also been reported across $Re = 400$–3900.[126] The fidelity and numerical provenance of simulation data should therefore be treated as part of the model's validity domain.

Meanwhile, experimental data provide a complementary route. PIV and volumetric velocimetry capture physical wake fields but carry calibration, outlier-detection, field-of-view, and spatial-resolution limits.[88–90] Pressure sensors, force measurements, and sparse probes are attractive for monitoring because they can be sampled at accessible locations, but they convert full-state prediction into an observability problem.[53,128,129] Experimental-data-driven models have been used to reconstruct time-resolved flow fields from non-time-resolved PIV and sparse high-frequency measurements,[96,97] and to recover wake fields from limited sensor information in high-*Re* cases.[98–100] These studies move neural-network modeling closer to monitoring and digital-twin applications, where only partial observations are available.[53,128] Nevertheless, experimental datasets also impose specific limits: the field of view may not include the full wake,[34,90] near-wall or near-body regions may be difficult to measure,[34,88,90] pressure and velocity may not be available simultaneously,[34,53] and temporal or spatial resolution may be limited by acquisition rate, seeding, calibration, or noise.[89,130,131] Consequently, experimental neural-network models often solve an observability problem instead of a complete flow-prediction problem, and they should not be judged only as interpolation exercises.

13 October 2026 05:32:48

© Author(s) 2026

More recent reconstruction studies further broaden this partial-observation setting. Jin *et al.*[102] combined non-time-resolved velocity fields with time-resolved surface-pressure sequences in a CAE-Transformer framework for time-resolved circular-cylinder wake reconstruction, directly connecting velocity-field reconstruction with pressure-sequence observability. Xie *et al.*[101] reconstructed laminar and turbulent bluff-body wakes from randomly distributed sparse data using a physics-informed generative framework, including laminar circular-cylinder, turbulent circular-cylinder, and turbulent square-cylinder datasets. These studies reinforce that experimental or sparse-input learning should be interpreted as an observability and validity-domain problem rather than as ordinary image completion.

A practical data-fidelity hierarchy follows from these limitations. DNS offers controlled high-fidelity fields in limited regimes;[32,34] LES gives access to turbulent separated-wake statistics while carrying grid, domain, and subgrid-model dependence;[35,36,132] and URANS/RANS supports engineering-scale datasets and many-query studies, with outputs conditioned by turbulence-model bias.[37,132,133] PIV/PTV grounds the model in laboratory measurements but is limited by field of view, resolution, near-wall access, calibration, and noise.[88–90,130] Pressure taps, force balances, and sparse probes are valuable for monitoring and digital twins, provided observability is tested.[97,99,128,129] Hybrid datasets can combine numerical and experimental strengths only when provenance, uncertainty, and consistency are reported.[53,87]

Multi-fidelity learning provides one way to address the trade-off between broad low-cost parameter coverage and limited high-fidelity data. It combines broad but biased or inexpensive low-fidelity coverage with limited high-fidelity data through shared or learned latent representations and an explicit discrepancy or calibration map.[134,135] Cylinder-flow studies illustrate both Gaussian-process correction of autoencoded low-fidelity states[134] and a frozen-latent-space bi-fidelity strategy for confined flow.[135] This construction changes how information is combined; it does not automatically remove solver or model bias, so cross-fidelity consistency and held-out high-fidelity validation remain necessary.

More generally, data fidelity and data quantity should be interpreted together with physical coverage and provenance. A large numerical dataset may still represent only a narrow range of Re, geometries, or numerical assumptions, while sparse experimental data may contain distinct information about sensor noise and observability. As one representative within-study example, Li *et al.*[98] reported L_2 errors of about 0.06–0.13 for high-Re cylinder reconstruction using 5–128 sensors, with larger errors near shear layers and the near wake. Such values should be interpreted within the specific sensor layout, target variable, normalization, grid, regime, and train-test definition rather than as cross-study benchmarks.

Reported quantitative performance also depends strongly on the learning task and metric. For example, Yun and Lee[99] reported correlation coefficients above 0.999 for pressure-to-wake reconstruction ($Re = 60$) and an average $R = 0.996$ for the CNN-D model at $Re = 300$, while Nemati Taher and Subaşı[91] reported a maximum relative error of 4.58% for three-dimensional bluff-body flow-field prediction. Together, these examples demonstrate that quantitative evidence is available across different bluff-body learning tasks, but the reported values should not be compared directly because the target variables, error definitions, normalization, flow regimes, geometries, and test protocols differ.

Within the direct bluff-body studies reviewed here, explicitly calibrated uncertainty intervals remain comparatively uncommon. This limits meaningful cross-study uncertainty benchmarking and motivates the uncertainty-reporting recommendations in Sec. VI C and Table III.

Against these data and fidelity background, current studies can be interpreted through four broad learning settings. The first is fully observed numerical learning, where dense fields are used for training and testing.[42,43,92,93,95] The second is partially observed reconstruction, where the model infers missing temporal or spatial information from non-time-resolved PIV, sparse sensors, wall-pressure input, or turbulent-region indicators.[96–100,103] The third is physics-constrained or closure-oriented learning, where residuals, conservation constraints, or learned closure terms are used to improve consistency or turbulence modeling.[105–107,124,125,136] The fourth is active-flow-control learning, where policies use wake measurements or states to select actuation for drag reduction, lift-fluctuation suppression, or stabilization.[109–113]

The same distinction can be expressed as an input–output taxonomy: coordinates to fields, fields to future fields, pressure to wake fields, sparse sensors to dense fields, geometry descriptors to flow responses or coefficients, resolved variables to closure terms, and probes or state variables to actuation policies. These mappings test different physical assumptions, so identical error metrics cannot validate them equally.

2. Representations of bluff-body flows

The representation of the flow determines what the neural network can learn. The most common representation is a Eulerian grid of velocity, pressure, or vorticity fields. Such snapshot-based representations are natural for CNNs, U-Nets, and convolutional autoencoders, because these architectures can exploit local spatial correlations. They have been widely used for flow-field prediction and reduced-order modeling.[43–45,92] In bluff-body wakes, this representation is effective because the dominant structures, such as shear layers, recirculation zones, and alternating vortices, appear as spatially organized patterns. The success of CNN-based models in cylinder-wake prediction and latent reduced-order modeling reflects, in part, the image-like regularity and low-dimensional organization of canonical wake fields.[43,44] Super-resolution should be treated more cautiously: unless a cited study explicitly includes a bluff-body wake case, it provides broader fluid-mechanics support for resolution enhancement instead of direct evidence for bluff-body wake learning.

A second representation uses temporal signals, such as surface pressure histories, probe velocities, lift and drag histories, or sparse sensor measurements. This representation is physically attractive because practical bluff-body monitoring often relies on wall pressure, force signals, or limited probes instead of full-field measurements.[30,46,73] Pressure-to-field learning is especially relevant because surface pressure contains information about separation, vortex formation, and unsteady loading.[94,99] For example, Yun and Lee[99] used instantaneous wall pressure to predict near-wake transverse velocity fields for laminar cylinder flows, while Jin *et al.*[102] combined non-time-resolved velocity fields with time-resolved surface-pressure sequences to reconstruct time-resolved velocity fields around a circular cylinder. Such formulations implicitly assume that the measured signal or reduced coordinates contain sufficient information about the wake state; this observability assumption should be tested with

13 October 2026 05:32:48

© Author(s) 2026

TABLE III. Recommended uncertainty-aware validation and reporting items for neural-network-based bluff-body-flow models.

Model type	Main uncertainty source	Recommended robustness tests	Recommended physical reporting	Minimum acceptable evidence
Snapshot or field predictor[42,43,92,95]	Phase drift, sparse training domain, solver bias	Long-horizon rollout; unseen *Re* interval; unseen geometry; grid/domain sensitivity	Strouhal number, phase error, recirculation length, mean wake, spectra, lift/drag statistics	Long rollout; Strouhal number; spectra; recirculation length; force statistics where available
Sparse-sensor or PIV reconstruction[96–102,123]	Sensor noise, limited field of view, nonunique reconstruction	Sensor-noise perturbation; sensor dropout; alternative layouts; independent PIV/CFD comparison	Wake topology, vortex-core location, spectra, Reynolds stresses, force-relevant structures	Sensor noise/dropout tests; independent PIV/CFD or withheld-field validation
Pressure-to-field model[88,99,128]	Pressure calibration, wall-pressure noise, observability limits	Pressure-noise perturbation; sensor removal; time-delay test; unseen wake state	Surface-pressure/wake consistency, shedding phase, near-wake velocity, force spectra	Pressure-wake consistency; time-delay sensitivity; unseen wake-state tests
Reduced-order temporal model[44,45,49,50,95]	Latent-space nonuniqueness, long-horizon instability	Long-horizon prediction; phase-drift test; unseen *Re* or wake regime	Modal energy, latent-variable interpretability, Strouhal number, force spectra	Phase drift; modal energy; long-horizon behavior; *Re* or regime holdout
PINN or physics-informed model[41,108,124,125,136,142]	Residual imbalance, boundary-condition uncertainty, unmodeled turbulence	Residual-weighting sensitivity; boundary-condition perturbation; independent CFD/experiment comparison	Conservation residuals, pressure field, wake topology, force response	Residual balance; boundary-condition sensitivity; pressure/force consistency; independent validation
Neural closure or RANS/LES correction[105–108,132,133]	Model-form uncertainty, grid dependence, coupled-solver instability	A priori and a posteriori tests; grid transfer; solver-coupled stability test	Mean wake recovery, base pressure, Reynolds stresses, turbulent kinetic energy (TKE), spectra, drag	Both a priori and a posteriori solver-coupled tests
Reinforcement-learning control[109–116,146]	Sensor noise, actuator saturation, time delay, policy overfitting	Unseen *Re*; 2D-to-3D transfer; noisy sensors; actuator limits; delayed feedback; partial observability; sparse pressure sensing; high-*Re* transfer	Drag/lift reduction, actuation cost, wake visualization, force spectra, robustness envelope	Actuation cost; robustness envelope; sensor/actuator constraints; force/spectral consequences
Geometry-conditioned surrogate[91,93,117,118]	Shape-space sparsity, interpolation/extrapolation ambiguity	Hold out body family; hold out spacing/angle; test finite-span or wall-mounted cases	Geometry-specific wake topology, pressure recovery, forces, uncertainty versus training-shape distance	Clear interpolation/extrapolation distinction for unseen shapes or wake regimes

noise, sensor removal, time delay, and unseen wake states.[128,129] The assumption is plausible for coherent periodic wakes,[46,47] but becomes more restrictive when the wake is turbulent, three-dimensional, or strongly affected by external disturbances.[55–58]

A third representation is based on latent variables or modal coefficients. Autoencoders, CNN-LSTM models, and DMD-assisted neural networks seek to compress high-dimensional wake fields into lower-dimensional coordinates and then learn their temporal evolution.[45,49,95] This approach is closely connected to reduced-order modeling and modal analysis in fluid mechanics.[48,50,137,138] For canonical bluff-body wakes, latent-variable learning is physically reasonable because vortex shedding often produces a dominant oscillatory structure, and a large fraction of the coherent energy may be represented by a small number of modes.[46–48] Data-driven equation discovery and coordinate discovery can help connect reduced coordinates to governing dynamics, but only when the discovered variables and equations are checked against the flow physics.[139,140] However, the physical meaning of the learned latent coordinates is not guaranteed; latent variables can be phase-like without being mechanistic.[140] Unless the latent variables are compared with known modal structures, shedding phase, force coefficients, or spectral content,[49,50,141] a reduced reconstruction error alone does not prove that the model has learned physically interpretable wake dynamics.

13 October 2026 05:32:48

© Author(s) 2026

A fourth representation introduces geometry explicitly. Geometry-conditioned models use signed-distance functions, body masks, coordinate fields, point-cloud representations, or other geometric descriptors to allow prediction across different bodies or arrangements.[42,91,93,117] This direction is important because many engineering bluff-body problems involve variations in shape, spacing, angle of attack, blockage, or multi-body interaction. Knowledge-enhanced CNNs and geometry-aware models have been proposed for fast prediction of unsteady flow around bluff bodies.[42,93] For example, Yu and Wu[42] developed a knowledge-enhanced CNN for near-real-time pixel-to-pixel simulation of unsteady bluff-body aerodynamics. Extending this idea to irregular three-dimensional domains, Nemati Taher and Subaşı[91] used a modified PointNet architecture to map unstructured-grid point-cloud coordinates around right-rhombic prism-shaped bluff bodies to three velocity components and pressure, avoiding interpolation onto structured grids. Yet geometry-conditioned learning remains challenging because the network must distinguish between universal wake features and geometry-specific mechanisms, such as fixed-edge separation in rectangular cylinders or variable separation location in circular cylinders.

B. Neural network formulations

Figure 2 organizes the neural-network formulations discussed in this Section according to their input–output structure and physical role.

1. Field prediction and wake reconstruction

Field prediction is one of the most direct applications of neural networks to bluff-body flows. In this formulation, the network receives one or several previous flow snapshots and predicts future velocity, pressure, or vorticity fields. The task is attractive because bluff-body wakes often possess strong temporal coherence, especially in two-dimensional or weakly three-dimensional regimes. CNN and U-Net-type architectures are commonly used because they can learn spatial correlations in Eulerian fields while preserving the image-like structure of the data.[43,92] For instance, Lee and You[43] showed that deep networks can predict future unsteady flow fields over a circular cylinder from previous fields and examined the effect of adding conservation-law information to the loss function; Abucide-Armas *et al.*[92] used a U-Net-type CNN to predict URANS-based cylinder-flow fields from previous samples and auxiliary spatial channels. These studies demonstrate that wake evolution can be learned over finite horizons, but they also highlight a recurring issue: recursive prediction can accumulate phase and amplitude errors when the model is advanced beyond the timescale represented in training. This formulation corresponds to the field-to-field mapping in Fig. 2(a), where previous wake fields are encoded into multiscale features and decoded into future velocity or pressure fields.

Wake reconstruction is a related but distinct task. In contrast to future-state prediction from fully observed past fields, reconstruction models infer missing spatial or temporal information from incomplete data. This includes time-resolved reconstruction from non-time-resolved PIV,[96,97] pressure-sequence-assisted reconstruction,[102] dense-field reconstruction from sparse sensors,[98] random-sparse-data wake reconstruction,[101] and pressure-input reconstruction of near-wake velocity fields.[99] Resolution enhancement from coarse fields is an important broader flow-learning task, but it should not be conflated with sparse-identification studies based on low-dimensionalized flow representations.[44] These reconstruction problems are particularly relevant to experiments, because full three-dimensional, time-resolved velocity, and pressure measurements are difficult to obtain in separated wakes. Their success relies on wake coherence: if the dominant shedding phase and spatial structure can be inferred from a small number of observations, the missing field can be reconstructed with reasonable accuracy.

Resolution enhancement and sparse reconstruction are increasingly connected to generative modeling. Super-resolution studies show that learned models can restore flow structures that are damped by coarse or sparse inputs.[120,122] Recent generative and diffusion-based studies make this direction more explicit: Oommen *et al.*[119] used generative modeling for turbulent-flow super-resolution, forecasting, and sparse reconstruction, including a cylinder-wake case with sparse PTV-like inputs, while Amoros-Trepat *et al.*[121] developed diffusion-based sparse flow-field reconstruction methods that generate high-fidelity samples and preserve smaller-scale spectral content. For bluff-body wakes, these methods are promising because sparse observations may not determine a unique instantaneous three-dimensional state; credibility therefore depends on distributional validation such as phase, spectra, pressure/force consistency, ensemble spread, and uncertainty calibration.[87] The probabilistic reconstruction logic in Fig. 2(f) is especially relevant when sparse observations admit multiple physically plausible wake states.

Within bluff-body wakes, Xie *et al.*[101] provide a more directly targeted example by reconstructing global turbulent wakes from randomly distributed sparse measurements and by testing both cylinder and square-cylinder wake datasets. This kind of evidence is important because it combines sparse sensing, turbulent wake complexity, and geometry-dependent bluff-body flow structure.

The CAE-Transformer reconstruction of Jin *et al.*[102] illustrates the same trend from another direction: a surface-pressure sequence is used to infer latent velocity-field coefficients, showing how pressure observability and sequence learning can be combined. Recent graph- and transformer-based reconstruction models extend this sparse-observation view to more flexible spatial representations. For example, Zhang *et al.*[123] introduced a graph-transformer reconstruction network that represents discrete grid points as graph nodes and uses transformer-style attention within message passing to reconstruct pressure and velocity fields from very sparse pressure sensors. Their sensor-placement analysis suggests that locating sensors in the core vortex-shedding region can be more important than simply increasing the number of sensors, which directly supports sensor-aware validation for bluff-body monitoring while remaining case-dependent with respect to *Re*, geometry, objective, and sensor layout. Figure 2(e) illustrates why graph and mesh-aware models are attractive for sparse sensors, irregular grids, and geometry-conditioned bluff-body configurations.

The maturity of emerging architectures is uneven. CNN and recurrent models have substantial bluff-body evidence for canonical prediction and reconstruction,[42,43,45,94,96] whereas neural operators, diffusion models, generative models, graph neural networks, and transformer-based reconstruction are still a mixture of direct wake studies and broader methodological support.[54,119–122] Sensor-graph and transformer reconstruction is promising for sparse measurements,

13 October 2026 05:32:48

© Author(s) 2026

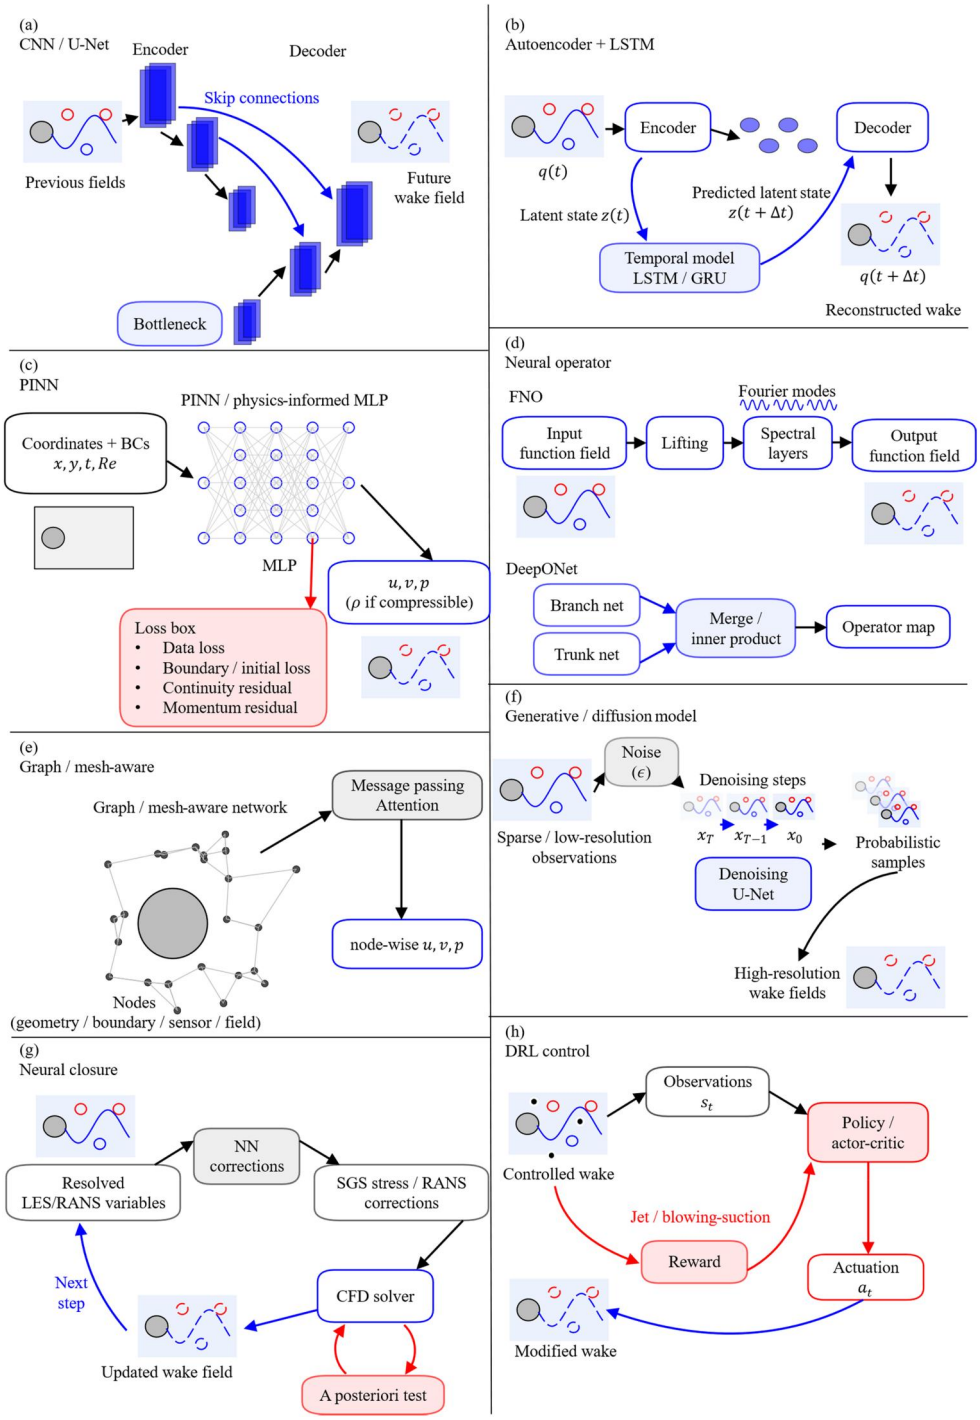


FIG. 2. Representative neural-network architectures used for bluff-body-flow modeling. (a) CNN/U-Net maps wake fields through multiscale convolutional features and skip connections.[42,43,92,93,95] (b) Autoencoder–LSTM models temporal wake evolution in a latent state.[43,45,94,95] (c) PINN combines data, boundary, continuity, and momentum residual losses for physics-constrained inference.[41,124,125,136,142] (d) Neural operators represent function-to-function mappings, here illustrated by FNO and DeepONet blocks.[54,142,143] (e) Graph and mesh-aware networks operate on nodes, edges, and message passing around bluff-body geometry.[84–86,123,128,129] (f) Generative/diffusion models combine sparse observations, denoising steps, and probabilistic samples for reconstruction or super-resolution.[87,119,121] (g) Neural closure models couple learned corrections with a CFD solver and require *a posteriori* testing.[132,133,144,145] (h) DRL control closes the loop between wake observations, policy, actuation, and reward.[146–148] Blue denotes flow or learned mappings, red denotes residual/control/validation paths, and gray/black denotes geometry, sensors, and structural links.

13 October 2026 05:32:48

© Author(s) 2026

but its evidence should be read against the tested sensor layout and wake regime.[123,128,129] Their promise should therefore be read against the evidence scope of the specific flow case, input–output map, and validation target.

Reduced-order neural models occupy an intermediate position between field prediction and reconstruction. CNN autoencoders, LSTM networks, and DMD-enhanced deep models first compress the wake field into a latent or modal representation and then learn its temporal evolution.[45,49,95] This type of model is physically appealing because it mirrors the organization of many bluff-body wakes around dominant coherent modes.[46–48] However, its applicability decreases as the wake becomes broadband, three-dimensional, or strongly turbulent.[55–58] In such regimes, the reduced-order representation may still reconstruct dominant large-scale features but fail to capture smaller-scale turbulence, intermittent events, or regime transitions.[50,137,138] The latent-state version of this idea is shown in Fig. 2(b), where an autoencoder compresses the wake and an LSTM/GRU-type temporal model advances the latent coordinates.

This is further discussed in Sec. IV: field prediction and reconstruction models primarily exploit correlation structure; correlation with a vortex street is not the same as learning the local separation process or the global wake mode that organizes shedding.[46,47] They may reproduce a vortex street, recirculation zone, or velocity fluctuation field with low numerical error, but this does not by itself show that they have learned the mechanism of separation, instability, or vortex formation.[49,50,141] A wake-prediction network may be excellent within a fixed *Re* range and geometry, yet fail when the separation process, shedding mode, or three-dimensionality changes.[55–58] Thus, prediction accuracy needs to be evaluated together with physical diagnostics, such as shedding frequency, phase stability, wake length, force coefficients, and spectral content.

2. Physics-informed and closure-oriented formulations

PINNs differ from pure regression because governing-equation and boundary residuals enter the objective.[136,142] For incompressible bluff-body flow, these constraints are typically derived from the continuity and momentum equations. Accordingly, a generic PINN objective can be written as Eq. (1), where $\mathcal{L}_{data}$ and $\mathcal{L}_{BC}$ penalize mismatch with available observations and boundary conditions, respectively, while $\mathcal{L}_{continuity}$ and $\mathcal{L}_{momentum}$ penalize residuals of the incompressible continuity and momentum equations; λ_d, λ_b, λ_c, and λ_m denote the corresponding loss weights. The latter residuals are constructed from the incompressible continuity equation and the Navier–Stokes momentum equations. Derivatives are usually evaluated by automatic differentiation. Residual weights, pressure identifiability, sampling, and boundary-condition enforcement materially affect the solution; satisfying these residuals does not by itself resolve transition, inflow uncertainty, wall treatment, or unresolved turbulence.[41,124,125,136] These residual terms therefore act as physical constraints and regularizers; their satisfaction is not evidence that the model has recovered the underlying instability or separation mechanism. The relevant evidence remains wake topology, pressure and force coupling, spectral behavior, robustness to boundary conditions, and transfer across physically withheld regimes[55–58,132,133,142]

$$\mathcal{L} = \lambda_d \mathcal{L}_{data} + \lambda_b \mathcal{L}_{BC} + \lambda_c \mathcal{L}_{continuity} + \lambda_m \mathcal{L}_{momentum}. \quad (1)$$

Knowledge-enhanced CNNs represent a more pragmatic variant of this idea. Without solving the governing equations through residual minimization, they incorporate physical knowledge through input design, loss functions, or constraints on the predicted field.[42,43] These methods are attractive for engineering applications because they retain the computational efficiency of CNN-based prediction while improving physical consistency. Such formulations have been used to improve unsteady bluff-body field prediction by adding geometry information, conservation-aware losses, or auxiliary physical channels to otherwise data-driven CNN-style predictors.[42,43] In this context, the central question is not whether the model is purely physics-based, but whether the embedded knowledge improves off-design behavior, temporal stability, and physically meaningful quantities.

Closure-oriented networks modify unresolved terms inside LES or RANS rather than directly replacing the full field map.[105–107,132,133] The representative unresolved stresses can be written as Eq. (2), where τ_{ij}^{SGS} denotes the subgrid-scale stress in LES, τ_{ij}^{R} denotes the Reynolds stress in RANS, the overbar indicates filtering or averaging as appropriate, and the prime denotes velocity fluctuations. Learning may target the stress itself, an eddy-viscosity or anisotropy correction, a model-form discrepancy, or an equivalent forcing. Credibility requires *a posteriori* solver-coupled tests of wake recovery, base pressure, drag and lift, Reynolds stresses, spectra, and numerical stability; *a priori* stress regression alone is insufficient[105,106,132,133,144,145]

$$\tau_{ij}^{SGS} = \overline{u_i u_j} - \overline{u_i}\,\overline{u_j}, \qquad \tau_{ij}^{R} = -\rho \overline{u_i' u_j'}. \quad (2)$$

More generally, physics constraints provide stronger evidence only when they improve the wake quantities associated with the mechanism they are intended to represent; residual satisfaction or offline prediction accuracy alone does not establish mechanism-level learning.

3. Reinforcement learning and flow control

The DRL setting in Fig. 2(h) closes the loop between wake observations, a policy network, actuation, reward, and the modified wake; this loop makes control validation fundamentally different from field-prediction validation. Reinforcement learning has been used to discover active-control strategies for bluff-body wakes, especially for the circular-cylinder wake with synthetic jets or blowing/suction actuation. In these studies, the network does not primarily predict the flow field. Instead, it learns a policy that maps flow observations, probe signals, or state variables to control actions. For example, Rabault *et al.*[109] showed that artificial neural networks trained through deep reinforcement learning can discover active-control strategies for a cylinder wake; Tang *et al.*[110] extended this idea by training a robust controller over a range of *Re*; Li and Zhang[111] further combined reinforcement-learning-based control of confined cylinder wakes with stability analysis. These examples connect control performance to wake response, robustness, and stability instead of treating reward reduction alone as sufficient evidence of physical learning.

More recent control studies have moved the field toward partial observability, pressure-based sensing, and higher-*Re* turbulent conditions. Xia *et al.*[114] studied bluff-body drag reduction under partial measurements and showed how history-augmented reinforcement learning can alleviate the loss of full wake observability. Wang *et al.*[115] used dynamic features built from sparse surface-pressure sensing for

13 October 2026 05:32:48

© Author(s) 2026

circular-cylinder wake control, directly connecting pressure observability with closed-loop actuation. Chen *et al.*[116] extended DRL-based bluff-body control to a high-*Re* turbulent cylinder-flow setting, using surface probes and zero-net-mass-flux jets. These studies make clear that the evidence standard for DRL control must include not only reward reduction, but also sensing constraints, actuation cost, *Re* robustness, and interpretable wake modification.

The physical target in these studies is usually suppression or modification of the von Kármán vortex street, reduction of drag, reduction of lift fluctuations, or stabilization of the wake.[31,46,149] This makes reinforcement learning attractive for bluff-body flows because the control objective is directly linked to wake dynamics.[109–111] However, policy success does not by itself establish mechanistic transparency. A DRL controller may reduce drag effectively, but the learned policy may be difficult to interpret in terms of classical flow-control concepts such as separation delay, base-pressure recovery, shear-layer manipulation, or global-mode stabilization.[146,148–151] Some studies address this issue by examining controlled wake structures, actuation histories, or stability properties,[111–113] but interpretability remains a general limitation unless learned actuation can be related to separation delay, base-pressure recovery, shear-layer forcing, or global-mode modification.[146–149]

Another challenge is robustness. Bluff-body control policies trained in two-dimensional laminar simulations may not transfer directly to turbulent, three-dimensional, noisy, or experimentally constrained flows.[110,112,147] The sensor arrangement, actuation authority, *Re* range, and reward definition strongly shape the learned strategy.[109,110,147] Therefore, DRL-based control should be evaluated not only by drag reduction in the training environment, but also by robustness to *Re* variation, measurement noise, actuation limits, three-dimensionality, and geometry changes.[110,112,113,147] From a physics-oriented perspective, reinforcement learning is most valuable when it reveals control-relevant wake mechanisms, instead of only producing a numerical policy.[146,148,149]

C. Validation practices and current limitations

Field-wise metrics such as mean-squared error, relative error, and pointwise error are useful first-level diagnostics, but they need to be evaluated together with wake-relevant quantities. For bluff-body flows, the appropriate validation target depends on the physical claim made by the model.[38,43,45] Figure 3(a) orders evidence from field accuracy (L1), through wake topology (L2), dynamics and loading (L3), turbulence or closure behavior (L4), physically defined generalization (L5), and mechanism-relevant response (L6). The hierarchy is cumulative: a low field error cannot substitute for missing topology, spectra, forces, or robustness tests.

The hierarchy in Fig. 3 implies that the appropriate validation evidence depends on both the physical claim and the model role. For unsteady wake prediction, spectral validation should include the dominant shedding peak, harmonics, broadband content, and consistency with lift and drag spectra, alongside mean flow, recirculation length, Reynolds stresses, and modal energy.[43,45,46,69,73,98] A −5/3 slope is relevant only for sufficiently turbulent wakes with credible scale separation and sampling; it is not a universal target for laminar or transitional cylinder wakes.[69] Sparse reconstruction additionally

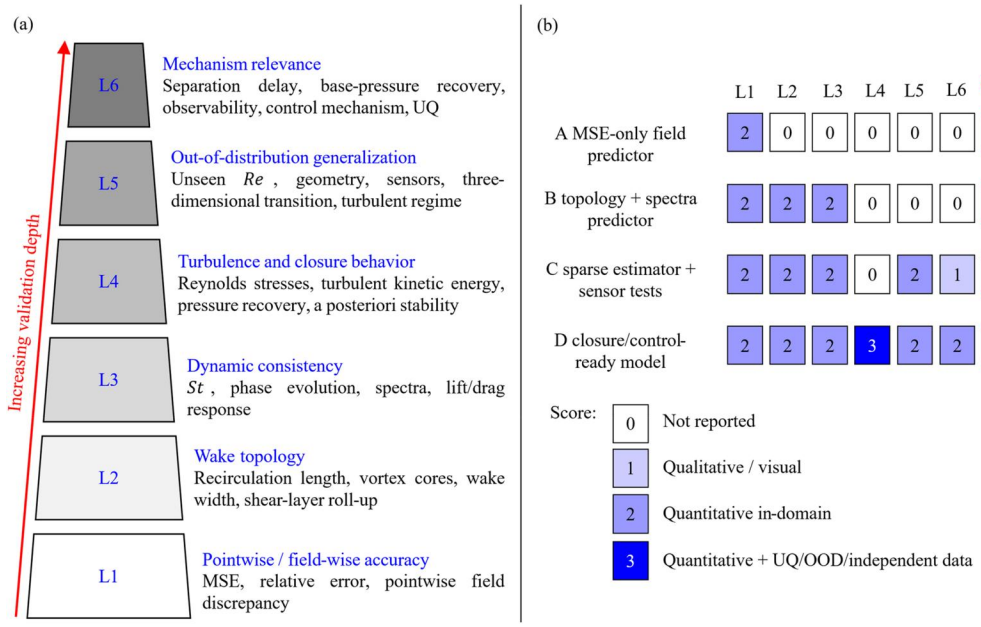


FIG. 3. Validation hierarchy and pseudo-quantitative evidence profiles for neural-network-based bluff-body-flow models. (a) Validation levels progress from pointwise or field-wise accuracy to wake topology, dynamic consistency, turbulence or closure behavior, out-of-distribution generalization, and mechanism relevance. (b) Illustrative validation profiles use ordinal reporting scores s_i: 0, not reported; 1, qualitative or visual evidence; 2, quantitative validation within the sampled regime; and 3, quantitative validation with uncertainty quantification, out-of-distribution testing, or independent data. The evidence class E_k denotes the highest consecutive validation level for which $s_1, \ldots, s_k \geq 2$. Scores are not averaged; high-field accuracy cannot compensate for missing wake-topology, dynamics, turbulence, generalization, mechanism-level, or other physics diagnostics.

13 October 2026 05:32:48

© Author(s) 2026

requires observability, sensor-noise, and sensor-removal tests; closures require *a posteriori* coupling; and control requires actuation-cost, force-response, and robustness checks.

Interpretation of these validation results also depends on a reproducible training and test protocol. Studies should therefore report the data split by physical case, preprocessing and normalization, architecture and parameter count, optimizer, learning rate and schedule, loss terms and weights, batch size, number of training epochs/iterations or stopping criteria, random seeds or repeat statistics, software, precision, and hardware. For time-dependent wakes, the forecast horizon and any temporal overlap between training and test data should also be explicit. These details are particularly important when apparent generalization may depend on how physical regimes are partitioned between training and testing.

Generalization beyond the training distribution is a difficult problem more broadly in machine learning (ML), where performance can depend strongly on benchmark construction and model-selection strategy, while formal results in specific settings illustrate fundamental limits to unrestricted generalization.[152,153] In bluff-body flows, this issue is more usefully formulated in terms of physically defined regime changes. Generalization can therefore be assessed at three levels, as listed below.

(1) In-domain testing: new states are sampled within a represented physical regime.
(2) Regime transfer: complete *Re* intervals, geometry families, blockage conditions, or finite-span configurations are withheld, while the governing phenomena remain represented somewhere in training.
(3) Unseen-physics testing: mechanisms absent from training are introduced, such as confinement-induced asymmetric wake states, end-effect or tip-vortex systems, compressibility-dependent shear-layer instability, or other regime-changing phenomena.[154–157] Such cases probe extrapolation to a changed physical state space rather than ordinary interpolation or parameter transfer and should be regarded as valuable stress tests rather than a reasonable minimum requirement for every surrogate.

Two limitations are especially important when interpreting such tests: changes in the physical state space and loss of trajectory predictability. The first arises when three-dimensionality or turbulence changes the relevant flow state itself. Many successful prediction and reduced-order studies are based on two-dimensional or quasi-two-dimensional cylinder wakes, where coherent vortex shedding dominates.[43,45,95] These cases are valuable benchmarks, but they do not fully represent spanwise instabilities, broadband turbulence, mode competition, intermittent structures, or turbulent transport. High-*Re* reconstruction, turbulent-region detection, and neural turbulence-closure models begin to address this gap,[98,103,105,106] but how neural networks handle three-dimensional wake transition remains only partly resolved. The second concerns temporal predictability. Long-horizon validation should separate trajectory or phase accuracy, statistical accuracy, and physical plausibility. Before the predictability horizon, phase and trajectory errors remain meaningful validation measures. Beyond this horizon, dynamical instability progressively limits deterministic trajectory agreement, so long-term statistical behavior becomes a more appropriate target.[158] For chaotic or turbulent wakes, validation should therefore emphasize mean and RMS fields, spectra, force distributions, Reynolds stresses, modal energy, topology distributions, and other statistical invariants.[44,45,47,95,129,141] A model may lose instantaneous phase while preserving a credible statistical wake, or retain smooth fields while drifting to an unphysical attractor.

Taken together, these considerations explain why model classes should not be ranked by a single scalar metric. Figure 3(b) therefore records the highest consecutively supported evidence level rather than averaging heterogeneous validation scores. An estimator that survives sensor-noise and regime-transfer tests may thus provide stronger evidence for its intended claim than a lower-error predictor that has been validated only by field-wise accuracy.

IV. PHYSICAL CONTENT LEARNED BY NEURAL NETWORKS IN BLUFF-BODY FLOWS

Sections I and III separated the physical mechanisms of bluff-body wakes from the neural-network formulations used to model them. This section connects the two by asking which physical content is supported by current evidence. While Tables I and II provide two complementary maps of the literature, Sec. IV uses these classifications not to repeat the taxonomy, but to assess the strength of mechanism-level evidence behind the main modeling claims. The assessment is organized around five mechanisms: near-wake organization and separation, coherent shedding and reduced-order dynamics, pressure-wake coupling and sparse observability, three-dimensional and turbulent regime transfer, and physics-constrained closure/control formulations. The validation hierarchy in Fig. 3 provides the evidence standard, while Fig. 4 summarizes the resulting qualitative evidence map across architecture families and physical targets.

The synthesis in Fig. 4 indicates that current evidence is most mature for coherent shedding and near-wake reproduction in canonical regimes, while separation causality, three-dimensional transition, turbulent transport, closure robustness, control robustness, and geometry-dependent regime transfer remain weak to emerging.

A. Wake organization and near-body separation

A useful starting point is to distinguish wake-image reproduction from representation of the mechanisms that create the wake. Convolutional and encoder–decoder models can reproduce bluff-body or URANS-based flow fields with small pointwise error when training and test data share geometry, grid, and numerical method.[42,92] Cylinder-wake forecasting studies show the same strength within prescribed *Re* ranges and wake regimes.[43,95] This performance matters because separated shear layers, recirculation regions, and alternating vortices are organized flow structures instead of random image patterns.[46,70] This still does not show that the model has learned which variables control the separation location.[1–3] Nor does it show that the model will predict changes caused by shear-layer instability or drag-crisis behavior.[4–6]

For rounded bodies, this distinction is especially important because separation is not imposed by a fixed sharp edge. On a circular cylinder, separation depends on *Re*, shear-layer instability, and boundary-layer transition.[4,62] Blockage, confinement, and boundary-layer tripping add additional sensitivities that can reshape the near

13 October 2026 05:32:48

© Author(s) 2026

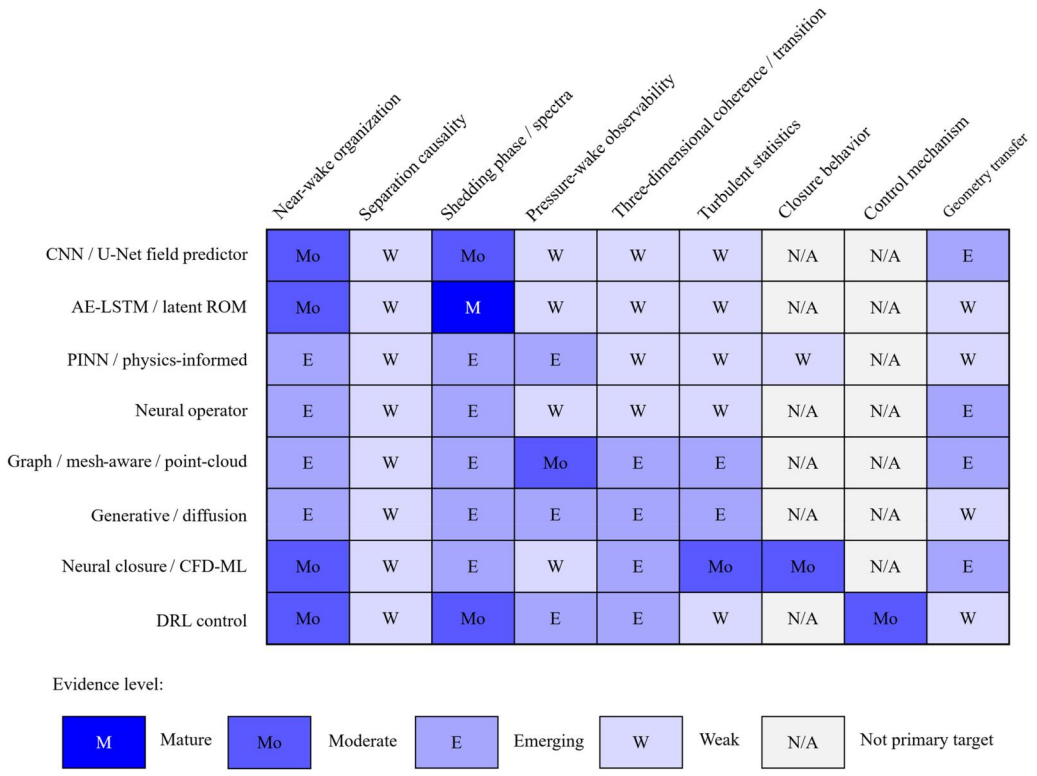


FIG. 4. Qualitative evidence map linking neural-network architectures to bluff-body-flow physics. Rows correspond to the architecture families shown in Fig. 2, whereas columns denote physical contents or validation targets. Cell colors indicate qualitative evidence strength rather than quantitative scores. Mature (M) denotes multiple direct bluff-body wake studies with quantitative physical diagnostics and at least limited robustness or regime-transfer evidence. Moderate (Mo) denotes direct bluff-body evidence with quantitative validation, but mainly within canonical or in-domain regimes. Emerging (E) denotes a small number of direct cases or evidence transferred from closely related separated-flow problems. Weak (W) denotes cases where the architecture has been used or proposed but mechanism-level bluff-body evidence remains limited. N/A indicates that the physical target is not a primary purpose of that model family. The grading is intended as a synthesis of evidence scope, validation depth, and physical relevance, not as a bibliometric count of papers.

wake.[61,63] If such variables are absent from the input representation, the network can only infer them indirectly from the sampled fields. A field predictor may therefore reproduce the separation pattern present in a dataset while remaining conditional on that dataset.[42,43,95] Its physical credibility becomes weaker when the intended task is extrapolation across blockage, body arrangement, or geometry.[61,93,117]

For sharp-edged and multi-body bluff bodies, geometry-conditioned learning moves closer to physical modeling because the network must respond to fixed-edge separation, body spacing, angle, or shape descriptors. Three-dimensional bluff-body prediction with geometric inputs, fast prediction around multiple bluff bodies, and consistent shape descriptions show that body geometry can be represented explicitly instead of treated as an implicit feature of one grid.[91,93,117] Brahmachary and Thuerey[118] give a sharper example: their differentiable physics-assisted neural network combines a coarse differentiable solver with a neural correction and tests unsteady wakes from arbitrary bodies, including out-of-distribution shapes. Such formulations make the physical target clearer because the learned component is asked to correct local boundary-layer and wake behavior while the solver carries part of the transport structure.

The main lesson is that current neural-network models provide stronger evidence for reproducing near-wake organization than for inferring separation causality. They can reproduce visible consequences of separation, including recirculation zones, shear-layer roll-up, coherent vortex streets, and mean wake recovery, when those features are sufficiently sampled in the training regime.[42,43,95] Low-dimensional and spatiotemporal reconstruction studies support the same point for organized wake evolution.[44,97] More direct evidence of physically meaningful prediction would require diagnostics such as separation angle, base pressure, recirculation length, shear-layer thickness, reattachment behavior, and wake recovery, instead of only field-wise metrics.[30,46,73]

Representative studies illustrate this hierarchy. Cylinder-wake predictors and URANS-based field models show how geometry or auxiliary field information can improve wake-field reproduction, but their physical interpretation still depends on whether recirculation length, base pressure, force spectra, and off-design cases are reported.[43,92] Differentiable-physics-assisted and geometry-conditioned formulations move closer to mechanism-aware surrogate modeling when a solver or explicit geometry representation carries part of the transport and boundary information,[91,118] but generality still rests on the breadth of out-of-distribution shape and regime tests.

13 October 2026 05:32:48

© Author(s) 2026

B. Vortex shedding, phase, and reduced-order dynamics

1. Frequency, phase, and spectral consistency

The strongest current evidence concerns coherent shedding in canonical two-dimensional cylinder wakes, where a dominant shedding frequency and strong phase organization make finite-horizon field prediction and reduced-order evolution comparatively tractable.[43,45,46,73] Such results show that the organized consequences of shedding can be reproduced within a tested regime, but they do not establish that the frequency-selection mechanism itself has been identified.

Mechanism-level evidence therefore combines the dominant spectral peak with harmonics, broadband content, phase stability, modal energy, lift and drag spectra, and sensitivity to *Re*, boundary conditions, end effects, and multi-body mode competition.[30,46,60,66,67,72,73]

2. Latent dynamics and long-horizon stability

Low-dimensional representations are effective when coherent coordinates organize the wake.[46–52] Linear or nonlinear latent models can track a canonical shedding cycle, but recursive prediction may accumulate phase error over long horizons even when the reconstructed wake retains a plausible dominant frequency or statistical structure.[45,95,129] Long-horizon performance should therefore be assessed through stability and statistical fidelity as well as reconstruction or pointwise trajectory error.

For latent-dynamics models, the general evidence hierarchy can be specialized further. A useful progression is: (1) representation, whether the latent coordinates reconstruct the sampled wake states with acceptable fidelity;[44,45,51] (2) in-manifold dynamics, whether local neighborhoods, phase evolution, and trajectories are preserved within the sampled regime;[44,45,95] (3) same-regime extrapolation, whether latent dynamics remain credible beyond the sampled support without changing the underlying wake mechanism;[45,52,95] (4) cross-regime or cross-geometry transfer, whether useful latent organization persists when the physical state manifold itself changes;[51,55–58,64,65] and (5) physical interpretability, whether latent coordinates can be related to quantities such as shedding phase, modal amplitude or energy, force response, or other identifiable wake mechanisms.[44,50,51,140] These levels are complementary to the general validation hierarchy in Fig. 3 rather than a replacement for it: accurate reconstruction is necessary evidence for a latent representation, but it does not establish preservation of latent dynamics or physical meaning.

The validity of this low-dimensional picture narrows quickly once the wake departs from the two-dimensional circular-cylinder archetype, because the relevant state manifold can change with geometry, transition, and turbulence. Twisted, inclined, and tapered cylinders introduce spanwise modulation, geometry-dependent vibration, and wake variation.[77–79] Elliptical, rectangular, and rounded-square cylinders show that body shape changes instability routes and coherent structures.[64,65,85,86] Neural networks trained mainly on two-dimensional circular-cylinder wakes can therefore learn an important archetype, but that archetype should not be treated as bluff-body wake physics in general.

C. Pressure-wake coupling and sparse reconstruction

Sparse and pressure-assisted reconstruction relies on wake observability. Recurrent, spatiotemporal, and pressure-sequence models have reconstructed cylinder wakes from incomplete measurements.[96,97,102] Representative quantitative reconstruction errors are discussed in Sec. III A 1; their interpretation remains conditional on *Re*, sensor layout, normalization, and target field rather than defining a universal benchmark.

The key question is therefore not only whether a dense field can be reconstructed, but whether the inferred state remains physically consistent with the available observations. PIV- and sensor-based cylinder studies show that neural networks can recover dense or time-resolved fields when coherent structures provide strong spatiotemporal constraints.[96–98] Mo and Magri[100] further demonstrated reconstruction from sparse and noisy measurements using a physics-constrained convolutional neural network. Stronger physical evidence requires preservation of the relationships among measured pressure or velocity, unobserved wake regions, dominant coherent dynamics, spectra, Reynolds stresses, and force-relevant structures.[46,73,76]

This observability weakens as the wake becomes less coherent. Similar pressure or probe signatures may correspond to different instantaneous three-dimensional wake states,[55–58] so credible reconstruction in these regimes requires uncertainty quantification, independent validation, spectral and force consistency, and sensor-robustness tests.[87,128,129]

D. Three-dimensionality, turbulence, and regime transfer

1. From two-dimensional benchmarks to three-dimensional wakes

The success of many neural-network models on two-dimensional cylinder wakes needs to be evaluated in light of the special physical structure of this benchmark. In this regime, the wake is dominated by coherent shedding, strong phase organization, and relatively low-dimensional temporal evolution.[46,47] These features make the problem favorable for snapshot-based, recurrent, CNN, autoencoder, and sequence models.[43,45,95] The benchmark is therefore valuable, but it is also physically restricted.

Three-dimensional transition changes the learning problem instead of merely adding one spatial coordinate. Once spanwise instability appears, the wake is no longer described by a single planar shedding phase. The relevant state may include spanwise wavelength, modal content, streamwise vortical structures, end effects, and intermittent loss of coherence.[55–58] Inclination, tapering, finite-span effects, and wall-mounted configurations add further spanwise modulation and geometry-dependent vortical structures.[77–79,84] A model trained on planar snapshots may therefore reproduce a two-dimensional attractor while remaining blind to the variables that control three-dimensional wake organization.

Recent sectional-to-volumetric reconstruction studies make this challenge more concrete. Matsuo *et al.*[104] reconstructed the three-dimensional wake of a square cylinder from a small number of sectional flow fields using coupled two-dimensional and three-dimensional convolutional networks. This shows that spanwise wake structure can be inferred when sectional observations contain sufficient information, but it remains a reconstruction task within a

13 October 2026 05:32:48

© Author(s) 2026

specific simulated wake, sectional-plane arrangement, spanwise extent, boundary condition, and training regime. It is therefore emerging evidence for three-dimensional wake estimation rather than proof of general transfer across three-dimensional transition routes.

2. Generalization across Re, geometry, and boundary conditions

Current literature coverage is uneven across regimes. Evidence is strongest for finite-horizon prediction and reduced-order temporal learning in organized two-dimensional cylinder wakes.[43,45,95] Sparse reconstruction in coherent wakes is moderate to strong when observability is high.[96–98] High-*Re* turbulent reconstruction, random-sparse turbulent-wake reconstruction, and turbulent-region detection are emerging directions.[98,101,103] Direct evidence that networks learn transition routes such as Mode A/B, spanwise wavelength selection, oblique shedding, or geometry-dependent three-dimensional modes remains weak to emerging.

Regime transfer is difficult because bluff-body wakes are sensitive to *Re*, body shape, arrangement, inflow, confinement, and wake state.[12,17] Blockage can change the wake by constraining the surrounding flow.[61] Elliptical and rounded-square cylinders show that body shape changes the instability route.[64,65] Tandem or side-by-side sharp-edged cylinders add spacing-dependent mode competition.[66,67] Machine-learning models that interpolate within a prescribed parameter space may therefore appear robust while remaining fragile outside that space. Geometry-conditioned surrogates for multiple bluff bodies and consistent shape descriptions are promising because they expose the model to systematic geometric variation.[93,117] Even these models should be tested against physically distinct regimes, not merely withheld samples from the same design distribution.

Consistent with the validation framework introduced in Sec. III C, three-dimensional and turbulent wakes require a broader evidence base than visual field agreement alone. Spanwise coherence, modal content, and force spectra are particularly relevant to three-dimensional wake organization,[55–58] while turbulent wakes additionally require statistics related to transport and loading, including Reynolds stresses, turbulent kinetic energy, pressure recovery, spectra, and force statistics.[35,36,69] For learned closures, these quantities must ultimately be assessed *a posteriori* in the coupled LES or RANS solver, because accurate offline stress prediction does not guarantee stable or physically correct wake behavior.[105,106]

Recent volumetric, turbulent, generative, diffusion, pressure-assisted, and graph-based reconstructions extend the evidence beyond planar cylinder benchmarks.[101,102,104,119,121,123] However, robust transfer across three-dimensional transition, turbulent transport, and geometry-dependent wake regimes remains limited, as does robustness to changes in sensing configuration. *Re* and wake regime should therefore be treated as validation variables rather than merely as input labels.[45,55–58,85,86,98]

E. Physics constraints, closures, and control as tests of physical learning

1. Constraint-based models and learned closures

Physics-informed and closure-oriented models change the question from which correlations are learned to which constraints are enforced. PINNs embed governing-equation residuals or conservation laws in the training objective.[136] Bluff-body applications have used PINN-type formulations for transonic cylinder flow, velocity-pressure inference, and reconstruction or prediction from limited data.[41,124,125] A residual constraint improves consistency with the specified equations, but it is not a complete physical explanation. The wake state also depends on boundary conditions, separation history, and boundary-layer transition.[61–63,70] Three-dimensional instability and wake transition add constraints that a residual loss alone does not identify.[55–58] Zhang *et al.*[108] move this discussion closer to turbulence closure by training PINNs with unclosed RANS equations and sparse cylinder-flow data, using the inferred Reynolds-stress forcing to improve mean-flow and turbulence-statistics prediction in forward settings. This is stronger mechanism-level evidence than residual minimization alone because the learned closure is judged through its effect on solver-coupled wake prediction.

The physical content of a PINN is strongest when the embedded equations correspond to the unresolved part of the problem and when the data are sufficient to identify the relevant state. For transonic or compressible cylinder flow, equation constraints are central because pressure-density coupling is part of the target solution.[41] For low-dimensional laminar wakes, residual constraints can regularize sparse-data inference and reduce unphysical fields.[125,136] For high-*Re* separated wakes, unresolved turbulence introduces closure questions that cannot be eliminated by adding an incompressible residual alone.[35,69]

Closure-oriented neural networks are more directly tied to turbulence modeling. Instead of replacing the flow solver, they learn subgrid stresses, eddy-viscosity corrections, or closure terms to improve LES or RANS prediction of separated wakes.[105–107] This is physically meaningful because conventional turbulence models remain uncertain when pressure drag, separation, recirculation, and Reynolds stresses interact.[35,36,69] The credibility of a learned closure depends on *a posteriori* tests in a coupled solver rather than on *a priori* stress regression alone.[105,106] Relevant diagnostics include mean wake recovery, drag, pressure distribution, Reynolds stresses, spectra, and numerical stability.

2. Control policies and mechanism-level interpretation

Reinforcement-learning control provides a different test of whether a learned policy can exploit control-relevant wake information. In cylinder-wake control, the network represents or optimizes an actuation policy, and success is commonly measured by drag reduction, lift-fluctuation suppression, or wake stabilization.[109–111] Chatzimanolakis *et al.*[112] tested whether policies trained in planar cylinder flows can transfer to three-dimensional cylinder flows. Such results are more informative than a black-box reward improvement when the learned action can be related to separation delay, wake stabilization, or base-pressure recovery, subject to actuation energy, cost, and hardware constraints.[19,24,109,110]

Recent studies make this robustness issue more concrete. Xia *et al.*[114] addressed partial observability for bluff-body drag reduction by using body-mounted pressure measurements and history-augmented RL states. Wang *et al.*[115] showed that dynamic features extracted from sparse surface-pressure signals can improve circular-cylinder wake control and help bridge the gap between sparse sensing and closed-loop actuation. Chen *et al.*[116] extended this line of work to

13 October 2026 05:32:48

© Author(s) 2026

high-*Re* turbulent cylinder control, where surface probes, zero-net-mass-flux jets, *Re* variation, and spectral/mechanism analysis become part of the validation envelope. Déda *et al.*[159] broadened the feedback setting by using neural networks for flow analysis and control tasks including stabilization, equilibrium estimation, and sensor placement. The same physical standard applies: a learned controller or surrogate should be evaluated by wake structure, forces, stability properties, sensing assumptions, actuation cost, and robustness rather than by reward alone.

In summary, current models provide convincing evidence for reproducing several wake characteristics within canonical regimes, particularly coherent shedding and near-wake organization. Evidence remains weaker for physically meaningful transfer across separation mechanisms, three-dimensional transition, turbulent transport, learned closures, and control conditions. Model credibility therefore depends not only on predictive accuracy, but also on which wake relationship is being claimed and which physical regime changes have actually been tested.

V. SIGNIFICANCE AND PRACTICAL USES

A. Fast surrogate prediction for many-query problems

Neural-network modeling first changes the cost structure of repeated bluff-body prediction. Design-space exploration, uncertainty quantification, optimization, monitoring, and control can require many evaluations, while experiments and high-fidelity simulations remain expensive.[33,37,38] A trained surrogate can make these repeated queries feasible when the task is interpolation within a stated design or operating envelope.[42,91,92] This is the practical motivation behind many field-prediction, reduced-order, neural-operator, and geometry-conditioned studies: they exchange one expensive data-generation stage for fast evaluation during screening, monitoring, or preliminary design.[54,93,95]

This computational advantage, however, should be assessed against an appropriate baseline.[160] Comparisons should use a conventional solver or reduced-order method at comparable accuracy and should report data-generation and training cost, preprocessing, inference hardware and batch size, memory, target error, and the number of queries over which offline cost is amortized. A speed-up relative only to the high-fidelity label generator can overstate deployment value, especially when lower-fidelity CFD or classical reduced-order baselines already meet the required accuracy.

That computational advantage does not expand the physics domain by itself. A surrogate trained on a prescribed *Re* range, body family, inflow condition, and numerical model can support rapid screening of wake topology or loads.[42,91,92] Its credibility still depends on training-data provenance, uncertainty reporting, and validation on withheld physical regimes.[53,87] For bluff-body surrogates, domain validity means stating which geometries, *Re*, blockage levels, solver assumptions, and wake regimes are covered by the model, and which remain outside its intended use.[35,37,87]

Within these validity constraints, surrogate modeling is relevant to a broad range of bluff-body applications, including wind loads on buildings, offshore topsides and cylinders, ship airwakes, multiple-body design, monitoring, and flow control.[13,14,27–29,31,93,109–111] Most published ML evidence remains proof-of-concept or workflow demonstration: canonical geometries, simulated sensors, limited operating envelopes, and offline inference. Operational deployment additionally requires calibrated uncertainty, latency on specified hardware, data and sensor governance, failure handling, maintainability, and validation against the actual environment; this evidence remains comparatively limited.

B. Flow control and drag/noise/vibration mitigation

Active flow control is a second application scenario. DRL studies use pressure, velocity, force, or state information to choose actuation, with drag reduction, lift-fluctuation suppression, or wake stabilization as typical rewards.[109–111] The physical value of a policy increases when its action resembles separation delay, base-pressure recovery, shear-layer manipulation, or global-mode stabilization.[31,146,149,150] This interpretation matters because a control policy that only improves a scalar reward can hide whether the actuation is exploiting a known wake mechanism, overfitting a numerical environment, or using a strategy that would be infeasible in hardware.[110,112,113,147]

Recent DRL studies make the deployment constraints more explicit. Partial-measurement control of a square bluff body shows that performance can degrade when wake-state information is replaced by body-mounted sensors, unless the policy is designed to account for history and energy cost.[114] Sparse surface-pressure sensing for circular-cylinder control shows that pressure-based dynamic features can recover useful control information without full-field observation.[115] High-*Re* turbulent control demonstrates that practical relevance also requires testing sensor placement, actuation authority, *Re* variation, spectra, and interpretable wake modification.[116] These examples support the review's central claim that control success should be interpreted through physical wake response and robustness, not through reward reduction alone.

Reward improvement is insufficient as a deployment standard. Actuation cost, actuator authority, bandwidth, sensor noise, time delay, *Re* robustness, two-to-three-dimensional transfer, and experimental feasibility must accompany drag or lift reductions.[146,147,150] Existing robustness and experimental-transfer studies move in this direction, but the validation envelope remains central.[110,112,113] A policy that performs well in a two-dimensional laminar simulation should be treated as a control hypothesis, not as a deployable controller, until robustness, actuation constraints, and force or noise consequences have been tested.[110,112,147]

The most useful control studies therefore report both engineering and physical consequences. Engineering consequences include net power, actuator duty cycle, hardware limits, sensor availability, and tolerance to time delay.[146,147,150] Physical consequences include changes in pressure distribution, recirculation, vortex shedding, force spectra, and wake symmetry.[30,46,73,149] These diagnostics help distinguish a policy that has discovered a plausible flow-control mechanism from one that exploits numerical details or a narrowly tuned reward function.[111,148,149]

C. Feedback to experimental and computational fluid mechanics

Neural estimators can turn sparse measurements into wake-state, sensor-informativeness, closure-error, or control-relevant estimates. Figure 5 therefore uses a conservative deployment sequence: accompany every reconstructed state with uncertainty or OOD indicators; check recirculation, *St*, spectra, force statistics, Reynolds

13 October 2026 05:32:48

© Author(s) 2026

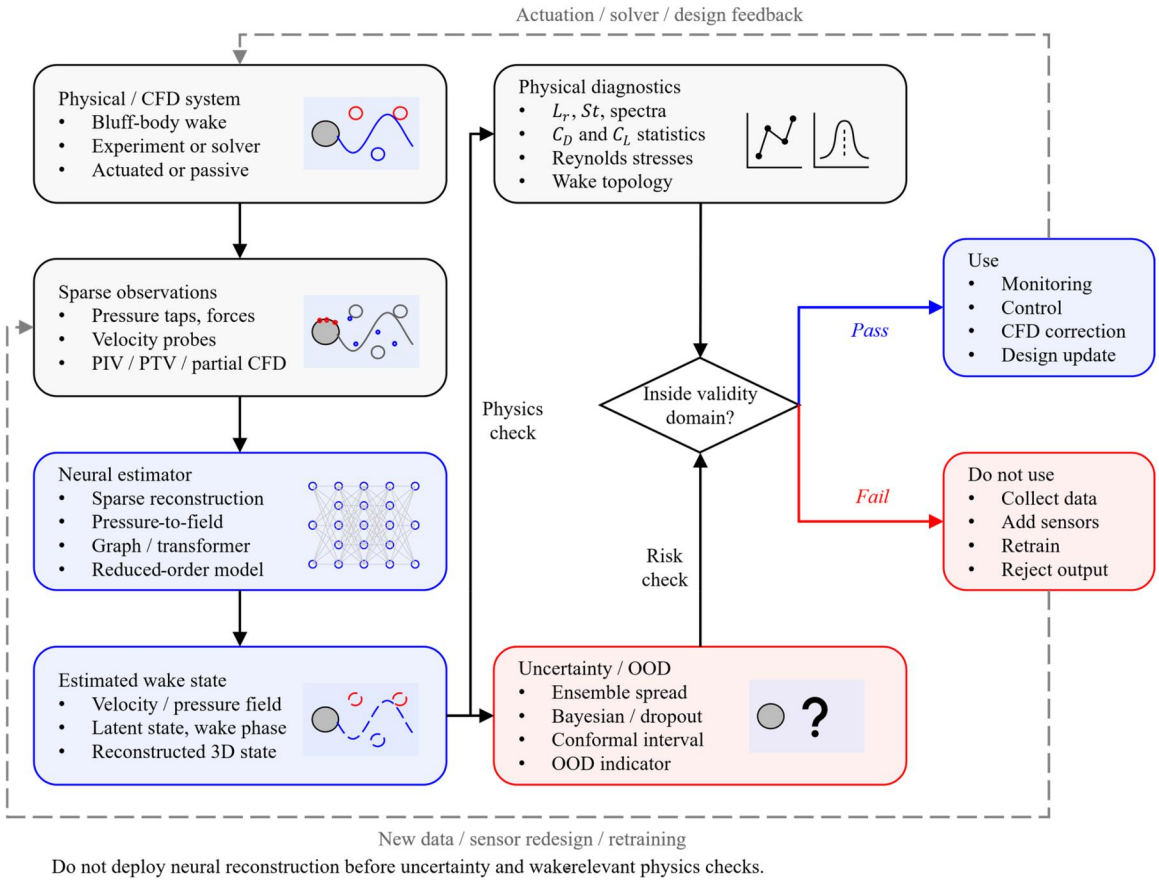


FIG. 5. Trustworthy sparse-observation workflow for bluff-body wake estimation and deployment.

stresses, and topology; and reject or reacquire data when the validity-domain test fails. This separates an informative proof-of-concept from an output suitable for monitoring, control, or CFD correction.

1. Experimental fluid mechanics

Sparse-sensor monitoring and experimental reconstruction form a third application scenario. Pressure-to-field, PIV-based, and limited-sensor studies show how partial observations can estimate wake states.[96–99] Their practical contribution is observability analysis: which pressure taps, PIV planes, force signals, or probes contain enough information to infer a target wake quantity.[128,129] This is particularly relevant for digital-twin-type workflows, where available measurements are often sparse, noisy, asynchronous, or limited to accessible surfaces, while the desired state may include near-wake velocity, vortex phase, loading, or a warning indicator for regime change.[53,87]

ML-based reconstruction should be treated as an estimator with explicit observability limits, not as a direct measurement of unobserved flow structures.[128,129] Independent PIV, pressure, force, spectra, calibration, and uncertainty checks are therefore part of the method, especially in digital-twin-type workflows.[34,88,89] A reconstruction that preserves coherent shedding in a laminar cylinder wake is informative,[96,97,99] but the same estimator may become nonunique in turbulent, three-dimensional, or multi-state wakes.[55–58,87] Sensor calibration, dropout tests, time-delay sensitivity, and independent validation define the practical credibility of this scenario.[87,128]

This distinction is important for experimental programs because reconstruction can otherwise look more precise than the measurements allow.[88,89,131] A pressure-to-velocity or sparse-probe-to-field model may supply a dense visual field, but that visual field is conditional on the training data, sensor layout, calibration, and assumed wake regime.[96–99] For reporting, the practical question is therefore not merely whether the reconstructed image is plausible, but which unmeasured quantities are actually observable from the available signals.[53,128,129]

2. Computational fluid mechanics

Hybrid CFD–ML closure or correction is the fourth application scenario. Neural networks can accelerate parametric studies, provide learned corrections, or supply turbulence-closure terms inside solver-based workflows.[105–107] This role is credible when the solver retains conservation, boundary conditions, transport structure, and interpretability, while the learned component supplies fast maps, closure corrections, or error compensation.[38,136,142] For separated wakes, this

13 October 2026 05:32:48

© Author(s) 2026

division is more defensible than a pure black-box replacement because pressure recovery, recirculation, Reynolds stresses, and vortex shedding remain constrained by the numerical model and its boundary conditions.[35–37,132,133]

A complementary hybrid strategy is to accelerate expensive internal solver operations rather than replace the flow solution itself.[38,142] In segregated incompressible CFD, pressure–velocity coupling requires repeated solution of a pressure Poisson equation, and the associated pressure-projection step can constitute a major computational bottleneck.[161] Recent studies have therefore used data-driven multigrid methods and neural-network surrogates to accelerate pressure projection, provide pressure updates or improved initial estimates, and reduce pressure-correction iterations while retaining the conventional CFD solver around the learned component.[161–164] This represents a distinct form of CFD–ML acceleration in which the learned model targets a computational bottleneck rather than the complete wake field.

The standard for hybrid CFD–ML is solver-coupled behavior: mean forces, base pressure, recirculation length, pressure recovery, Reynolds stresses, spectra, stability, grid/domain transfer, and robustness after the learned component is embedded in the computation.[105–107] *A priori* stress or coefficient accuracy is useful during model development, but practical credibility requires showing what happens when the learned component interacts with the solver, the grid, and the wake dynamics it is meant to correct.

Across these applications, practical value depends on the decision being supported: screening, monitoring, solver correction, or control.[53,132,146] Each decision requires its own evidence standard, so application claims should be tied to validity domain, observability, solver coupling, or actuation-aware robustness rather than to a single generic accuracy number.[87,147,149]

VI. FUTURE PERSPECTIVES

A. From task-specific success to physical generalization

Future benchmarks can distinguish among three distinct goals: in-support consistency, transfer to held-out physical regimes, and extrapolation to genuinely new physics. Useful regime-transfer tests hold out complete *Re* intervals, blockage levels, geometry families, or finite-span configurations; unseen-physics stress tests may add a new instability route, end-effect vortex system, confinement-induced asymmetry, compressibility mechanism, or drag-crisis state.[55–58,61,64–67,77–79,85,86] Reporting which category is tested prevents ordinary interpolation from being called OOD generalization and avoids treating failure when the relevant physics is absent from the training data as a universal minimum-standard failure.

Figure 6 places these three levels of generalization in the context of current research. Current studies are concentrated on in-domain prediction and reduced-order modeling of canonical wakes, while emerging work addresses sparse observation, geometry conditioning, graph/transformer reconstruction, generative reconstruction, and three-dimensional estimation. The major gap zone is physical transfer across three-dimensional transition, high-*Re* turbulent transport, finite-span or wall-confined wakes, experimental actuation, and uncertainty-aware deployment.

Achieving such physical generalization also requires neural-network formulations to engage more explicitly with the mathematical structure of fluid mechanics. Physics-informed machine learning provides one framework for combining data with governing

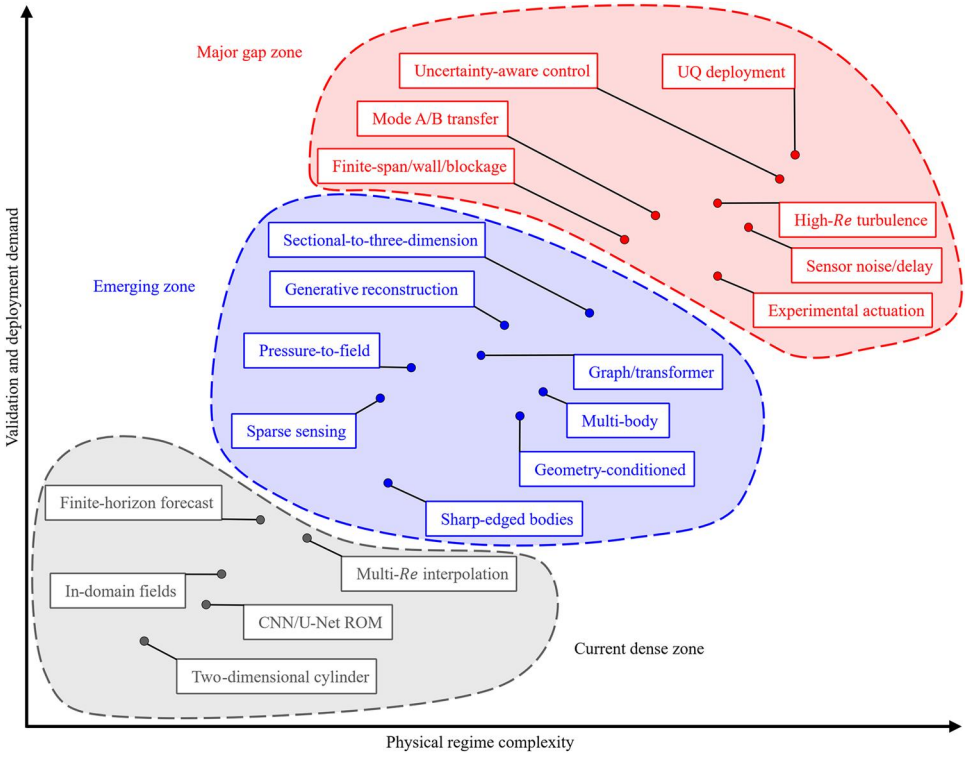


FIG. 6. Conceptual roadmap from interpolation-oriented benchmarks to physical-regime generalization. The horizontal direction denotes increasing departure from canonical two-dimensional cylinder-wake benchmarks, while the vertical direction denotes increasing validation and deployment requirements. The labels are conceptual anchors rather than quantitative coordinates.

13 October 2026 05:32:48

© Author(s) 2026

equations and constraints.[142] Neural operators provide another framework by learning mappings between function spaces for simulation and design.[54] PINNs and neural operators enter the modeling problem through equations, residuals, or operator priors.[136,143] Learned closures enter at the level of unresolved stresses or model-form correction.[105,106] The central challenge is not simply adding equations to a loss function, but choosing whether physics should enter as an input representation, a constraint, an operator prior, a closure term, a differentiable solver, or a validation diagnostic. For bluff-body flows, this choice should be guided by the mechanism being modeled: a pressure-to-field network, a learned SGS closure, and a reinforcement-learning controller encode and use physics differently.

B. Experimental, computational, and closure-oriented futures

Future progress also depends on closer integration with experiments. Vinuesa *et al.*[53] argue that machine learning can augment measurement fidelity, improve experimental design and digital-twin models, and enable real-time estimation and control in experimental fluid mechanics. For bluff-body applications, this means moving beyond idealized full-field training data toward pressure taps and force balances.[99] It also means using limited PIV and sparse or noisy probes as first-class data sources.[96–98] For control studies, actuator-constrained feedback should be tested as a first-class experimental condition.[109,110] The most valuable future studies will report reconstruction error and quantify observability, uncertainty, sensor robustness, and the physical consequences of missing information.

On the computational side, hybrid CFD–ML is a promising route when the division of work between the solver and the learned component is explicit. Rather than replacing the governing solver wholesale, learned models can provide surrogate mappings, model-form corrections, or unresolved closure terms while the numerical framework retains conservation, boundary conditions, and transport structure. Solver-internal acceleration, including learned assistance for pressure–velocity coupling and pressure Poisson solves, provides another hybrid route in which machine learning targets a computational bottleneck while the numerical solver retains the governing transport structure.[162] Closure-oriented approaches are particularly relevant to separated bluff-body wakes, where learned SGS or RANS corrections must remain stable when coupled back to the solver rather than merely reproduce offline training targets.[105–107,132,133] *A posteriori* assessment should therefore include drag, base pressure, wake recovery, Reynolds stresses, spectra, numerical stability, and robustness to grid, operating-condition, or model changes.

C. Uncertainty, out-of-distribution generalization, and sensor-aware validation

A major requirement for future bluff-body-flow learning is uncertainty-aware validation. Here, uncertainty includes measurement noise in PIV, pressure, and force signals;[34,88,89] numerical uncertainty associated with grids, domains, boundary conditions, and turbulence models;[35,132,133] epistemic uncertainty caused by sparse coverage of *Re*, geometry, blockage, or inflow conditions;[87] and aleatoric variability associated with turbulent wake fluctuations.[87] Each type should map to a validation action, such as calibration curves for probabilistic outputs, sensor perturbation for measurement uncertainty, and held-out regimes for epistemic uncertainty.[87] A deterministic neural-network prediction may therefore be visually accurate in an interpolation-oriented test while becoming unreliable when the wake regime, geometry, sensor layout, or observation process changes.

Consistent with the generalization framework in Sec. VI A, OOD claims should state whether the test concerns in-support variation, transfer across a held-out physical regime, or genuinely unseen physics. For sensor-based models, changes in the observation process should be distinguished from changes in the underlying wake regime: withholding a sensor layout, introducing sensor dropout, or perturbing measurement noise tests robustness to observability, whereas withholding *Re*, geometry, spacing, blockage, inflow, finite-span conditions, transition, or turbulence regime tests physical generalization.[55–58,85–87,128,129] Random snapshot splits alone provide neither type of evidence.

Table III is intended to align uncertainty sources, robustness tests, wake-relevant diagnostics, and minimum evidence thresholds so that validation claims are comparable across prediction, reconstruction, closure, control, and geometry-conditioned surrogate tasks.

Uncertainty tools should be tied to specific bluff-body validation purposes. Ensemble spread can be checked against reconstruction error; calibration curves test whether predicted intervals are reliable; sensor perturbations quantify measurement uncertainty; regime holdouts expose epistemic uncertainty; and generative or diffusion samples test distributional plausibility in sparse reconstruction.[87,119,121] Measurement-noise and calibration checks are especially important when the target field is inferred from PIV, pressure, or sparse sensors.[34,88,89] These checks matter for design, monitoring, control, and digital-twin workflows because model outputs may influence decisions.[53,87] Future studies should therefore report prediction intervals, calibration quality, failure cases, and the physical conditions where the model should not be used.

Progress will increasingly depend on explicit reporting of domain validity, data provenance, uncertainty, sensor assumptions, failure modes, and physical diagnostics.[53,87] This requirement is consistent with review-level perspectives on physics-informed machine learning and neural operators,[54,142] experimental fluid mechanics,[34,53,88] turbulence-model learning,[132,133,144,145] modal/reduced-order analysis,[50,137] and DRL-based control.[146–148] For bluff-body flows, trustworthy learning means specifying what has been learned, where it is valid, how it fails, and how the learned representation improves experiment, simulation, or control.

VII. CONCLUSIONS

This review leads to five take-home messages. First, the strongest current evidence for neural-network modeling is found in canonical, data-rich, and organized wake regimes. Second, physical learning should be treated as an evidence-based concept and requires validation of wake-relevant physical relationships, not pointwise or field-wise accuracy alone. Third, three-dimensional transition, turbulent transport, geometry-dependent regime transfer, and sparse or noisy observations remain difficult tests. Fourth, hybrid CFD–ML, experimental reconstruction, and control workflows become credible when paired with physical constraints, uncertainty, observability checks, fair computational baselines, and regime-aware validation. Fifth, future benchmarks should be regime-aware, sensor-aware, uncertainty-aware, and force-, spectra-, and topology-aware.

13 October 2026 05:32:48

© Author(s) 2026

A model can be credible inside a sampled coherent wake regime and unreliable under untested changes in Reynolds number, geometry, three-dimensionality, turbulence, sensors, or actuation. Such tests should distinguish interpolation within a represented regime, transfer to physically held-out regimes, and extrapolation into genuinely unseen physics. The field should therefore report training domain, data fidelity, observability assumptions, uncertainty, and wake-relevant diagnostics as part of every strong claim.

Neural networks are most valuable when they clarify what information controls a wake prediction, reconstruction, closure, or policy. They are weakest when treated as architecture-level replacements for physical reasoning. The review's central message is conditional: learned models can advance bluff-body science and engineering when their validity limits, failure modes, deployment constraints, and physical evidence are made explicit.

ACKNOWLEDGMENTS

The authors received no financial support for this work.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Xin Tang: Data curation (lead); Investigation (lead); Visualization (supporting); Writing – original draft (equal). **Jingnan Zhang:** Visualization (equal); Writing – original draft (supporting); Writing – review & editing (equal). **Heng Zhu:** Conceptualization (lead); Methodology (lead); Supervision (lead); Validation (lead); Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal).

DATA AVAILABILITY

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

APPENDIX: LITERATURE SELECTION AND CLASSIFICATION METHOD

The literature search was conducted as a narrative, physics-oriented search, not as a bibliometric or PRISMA-style systematic review. The search was last updated in August 2026 using Web of Science, Scopus, Google Scholar, ScienceDirect, SpringerLink, Cambridge Core, AIP Publishing, APS, and Taylor & Francis, supplemented by backward and forward citation tracking from canonical wake, review, turbulence-modeling, and scientific-ML sources.

Searches combined bluff-body and wake terms with neural network, physics-informed neural network, neural operator, reduced-order model, sparse sensing, pressure-to-field, multi-fidelity, hybrid CFD–ML, differentiable physics, industrial surrogate, learned LES/RANS closure, reinforcement learning, uncertainty, and OOD generalization. Studies were selected when the learning formulation could be connected to a wake mechanism, modeling task, or evidence requirement rather than counted by architecture keyword alone.

Included studies directly involved bluff-body wakes or closely related separated-flow configurations, including circular cylinders, rectangular or square cylinders, finite cylinders, wall-mounted bluff bodies, prisms, multiple bluff bodies, and generic bluff-body sections. Studies were included when they used neural-network-based or closely related learning methods for flow-field prediction, reconstruction, reduced-order modeling, physics-informed inference, closure modeling, geometry-conditioned surrogate prediction, or active flow control. Classical and review-level bluff-body references were included when they defined separation, coherent shedding, three-dimensional transition, wake topology, force response, pressure coupling, or turbulence mechanisms needed to interpret machine-learning results.

Broader scientific-machine-learning, neural-operator, sparse-reconstruction, uncertainty, and data-driven mechanics papers were treated as methodological support instead of as primary bluff-body evidence. Generic channel-flow, isotropic-turbulence, airfoil, or cavity-flow learning papers were not treated as direct bluff-body evidence unless they informed a method subsequently used for bluff-body wake modeling. Purely non-neural POD, DMD, Koopman, or SINDy studies were used only when they were coupled to, or directly interpreted, neural-network formulations. Pure CFD or experimental bluff-body studies were included only when they supported the physical mechanisms used to evaluate neural-network learning. Conference abstracts, inaccessible preprints, and non-peer-reviewed sources were avoided unless a specific methodological reason justified their use.

Tables I and II are representative classification tools, not exhaustive bibliographic inventories. Table I classifies studies by the physical mechanism or modeling claim they address: near-wake organization, coherent shedding, sparse observability, pressure-wake coupling, three-dimensional or turbulent wake complexity, closure, control, compressible or transonic flow, and geometry dependence. Table II classifies studies by neural-network formulation: snapshot prediction, sparse reconstruction, reduced-order temporal modeling, pressure-to-field mapping, PINNs, conservation-informed CNNs, closure learning, reinforcement learning, and geometry-conditioned surrogates.

For Fig. 4, the qualitative evidence level was determined from the combined strength of the available studies rather than from the number of publications alone. We considered whether a study provided direct bluff-body wake evidence, whether it reported quantitative physical diagnostics beyond field-wise error, whether it tested robustness, uncertainty, or out-of-distribution transfer, and whether the results were supported by independent numerical or experimental validation. This conservative grading is intended to indicate the credibility of mechanism-level claims, not the popularity of a neural-network architecture.

A study was assigned to more than one row only when it genuinely supported more than one interpretation. It was not placed in a row merely because it used a similar architecture; the physical claim, flow case, model input, model output, and validation target had to match the row definition. When a cited work came from broader fluid mechanics instead of a bluff-body-specific case, the table wording was kept explicit about that limitation. The representative examples in the tables were classified only when their model inputs, outputs, flow cases, and stated physical applications supported the corresponding category, rather than on the basis of architecture keywords alone.

13 October 2026 05:32:48

© Author(s) 2026

Because the purpose of this article is to interpret physical trends rather than to produce a complete bibliometric index, the review may not include every neural-network study involving separated flows. The classification should therefore be read as a synthesis of representative evidence. The main conclusions depend on whether current studies support mechanism-level claims about bluff-body-flow learning, not on the raw number of papers in each architecture category. The resulting taxonomy is also time-dependent: rapidly developing areas such as neural operators, uncertainty quantification, closure learning, and experimental digital twins will require periodic reassessment as more bluff-body-specific evidence becomes available.

REFERENCES

[1]R. T. Griffiths and C. Y. Ma, "Differential boundary-layer separation effects in the flow over a rotating cylinder," Aeronaut. J. **73**, 524–528 (1969).
[2]J. C. Westkaemper, "Turbulent boundary-layer separation ahead of cylinders," AIAA J. **6**, 1352–1355 (1968).
[3]M. B. Glauert and M. J. Lighthill, "The axisymmetric boundary layer on a long thin cylinder," Proc. R. Soc. London, Ser. A **230**, 188–203 (1955).
[4]S. P. Singh and S. Mittal, "Flow past a cylinder: Shear layer instability and drag crisis," Numer. Methods Fluids **47**, 75–98 (2005).
[5]J. Mihailovic and T. C. Corke, "Three-dimensional instability of the shear layer over a circular cylinder," Phys. Fluids **9**, 3250–3257 (1997).
[6]S. Yokota, K. Asai, and T. Nonomura, "Instability of separated shear layer around levitated freestream-aligned circular cylinder," Phys. Fluids **34**, 064104 (2022).
[7]G. H. Toebes, "The unsteady flow and wake near an oscillating cylinder," J. Basic Eng. **91**, 493–502 (1969).
[8]O. M. Griffin, "The unsteady wake of an oscillating cylinder at low Reynolds number," J. Appl. Mech. **38**, 729–738 (1971).
[9]S. Wornom, H. Ouvrard, M. V. Salvetti, B. Koobus, and A. Dervieux, "Variational multiscale large-eddy simulations of the flow past a circular cylinder: Reynolds number effects," Comput. Fluids **47**, 44–50 (2011).
[10]P. Baj and O. R. H. Buxton, "Interscale energy transfer in the merger of wakes of a multiscale array of rectangular cylinders," Phys. Rev. Fluids **2**, 114607 (2017).
[11]A. Chiarini, D. Gatti, A. Cimarelli, and M. Quadrio, "Structure of turbulence in the flow around a rectangular cylinder," J. Fluid Mech. **946**, A35 (2022).
[12]B. Forouzi Feshalami, S. He, F. Scarano, L. Gan, and C. Morton, "A review of experiments on stationary bluff body wakes," Phys. Fluids **34**, 011301 (2022).
[13]J. Wang, L. Geng, L. Ding, H. Zhu, and D. Yurchenko, "The state-of-the-art review on energy harvesting from flow-induced vibrations," Appl. Energy **267**, 114902 (2020).
[14]J. Wang, D. Fan, and K. Lin, "A review on flow-induced vibration of offshore circular cylinders," J. Hydrodyn. **32**, 415–440 (2020).
[15]T. Sarpkaya, "A critical review of the intrinsic nature of vortex-induced vibrations," J. Fluids Struct. **19**, 389–447 (2004).
[16]X. Wu, F. Ge, and Y. Hong, "A review of recent studies on vortex-induced vibrations of long slender cylinders," J. Fluids Struct. **28**, 292–308 (2012).
[17]J. F. Derakhshandeh and M. M. Alam, "A review of bluff body wakes," Ocean Eng. **182**, 475–488 (2019).
[18]Y. Wu, Z. Cheng, R. McConkey, F.-S. Lien, and E. Yee, "Modelling of flow-induced vibration of bluff bodies: A comprehensive survey and future prospects," Energies **15**, 8719 (2022).
[19]Y. Haffner, J. Borée, A. Spohn, and T. Castelain, "Mechanics of bluff body drag reduction during transient near-wake reversals," J. Fluid Mech. **894**, A14 (2020).
[20]S. Karimi, B. Mohammadikalakoo, and P. Schito, "Performance enhancement of single dielectric barrier discharge flow control actuators by means of rear linking tunnels on a reference bluff body using CFD," J. Wind Eng. Ind. Aerodyn. **209**, 104488 (2021).
[21]Z. Cheng, F.-S. Lien, E. Yee, and G. Chen, "Vortex-induced vibration of a circular cylinder with nonlinear restoring forces at low-Reynolds number," Ocean Eng. **266**, 113197 (2022).
[22]A. Du *et al.*, "Vortex-induced vibration of a rotating cylinder with dual splitter plates," Phys. Fluids **36**, 083623 (2024).
[23]H. Yao, L. Davidson, and Z. Chroneer, "Investigation of interior noise from generic side-view mirror using incompressible and compressible solvers of DES and LES," in *SAE Technical Paper Series* (SAE International, 2018), Vol. 1.
[24]R. Maryami, E. J. G. Arcondoulis, and Y. Liu, "Flow and aerodynamic noise control of a circular cylinder by local blowing," J. Fluid Mech. **980**, A56 (2024).
[25]K. Kusano, "Aeroacoustic simulation of bluff bodies with protrusions at moderate Reynolds number," Phys. Fluids **36**, 027146 (2024).
[26]G. Chen, B. Zang, and M. Azarpeyvand, "Numerical investigation on aerodynamic noise of flow past a cylinder with different spanwise lengths," Phys. Fluids **35**, 035128 (2023).
[27]M. S. Thordal, J. C. Bennetsen, and H. H. H. Koss, "Review for practical application of CFD for the determination of wind load on high-rise buildings," J. Wind Eng. Ind. Aerodyn. **186**, 155–168 (2019).
[28]M. H. Dao *et al.*, "Wind tunnel and CFD studies of wind loadings on topsides of offshore structures," Ocean Eng. **285**, 115310 (2023).
[29]K. Xu, X. Su, I. Jonsson, R. Bensow, and S. Krajnovic, "Numerical study of a generic ship's airwake for understanding bi-stability mechanism," J. Fluid Mech. **991**, A6 (2024).
[30]M. M. Alam, "Fluctuating forces on bluff bodies and their relationships with flow structures," Ocean Eng. **273**, 113870 (2023).
[31]F. Ren, X. Wen, and H. Tang, "Model-free closed-loop control of flow past a bluff body: Methods, applications, and emerging trends," Actuators **13**, 488 (2024).
[32]J. F. G. Schneiders, G. C. A. Caridi, A. Sciacchitano, and F. Scarano, "Large-scale volumetric pressure from tomographic PTV with HFSB tracers," Exp. Fluids **57**, 164 (2016).
[33]R. Yondo, E. Andrés, and E. Valero, "A review on design of experiments and surrogate models in aircraft real-time and many-query aerodynamic analyses," Prog. Aerosp. Sci. **96**, 23–61 (2018).
[34]S. Discetti and F. Coletti, "Volumetric velocimetry for fluid flows," Meas. Sci. Technol. **29**, 042001 (2018).
[35]H. Xiao and P. Cinnella, "Quantification of model uncertainty in RANS simulations: A review," Prog. Aerosp. Sci. **108**, 1–31 (2019).
[36]J. Holgate, A. Skillen, T. Craft, and A. Revell, "A review of embedded large eddy simulation for internal flows," Arch. Comput. Methods Eng. **26**, 865–882 (2019).
[37]A. J. Keane and I. I. Voutchkov, "Robust design optimization using surrogate models," J. Comput. Des. Eng. **7**, 44–55 (2020).
[38]S. L. Brunton, B. R. Noack, and P. Koumoutsakos, "Machine learning for fluid mechanics," Annu. Rev. Fluid Mech. **52**, 477–508 (2020).
[39]Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature **521**, 436–444 (2015).
[40]J. N. Kutz, "Deep learning in fluid dynamics," J. Fluid Mech. **814**, 1–4 (2017).
[41]X. Ren, P. Hu, H. Su, F. Zhang, and H. Yu, "Physics-informed neural networks for transonic flow around a cylinder with high Reynolds number," Phys. Fluids **36**, 036129 (2024).
[42]X. Yu and T. Wu, "Simulation of unsteady flow around bluff bodies using knowledge-enhanced convolutional neural network," J. Wind Eng. Ind. Aerodyn. **236**, 105405 (2023).
[43]S. Lee and D. You, "Data-driven prediction of unsteady flow over a circular cylinder using deep learning," J. Fluid Mech. **879**, 217–254 (2019).
[44]K. Fukami, T. Murata, K. Zhang, and K. Fukagata, "Sparse identification of nonlinear dynamics with low-dimensionalized flow representations," J. Fluid Mech. **926**, A10 (2021).
[45]K. Hasegawa, K. Fukami, T. Murata, and K. Fukagata, "CNN-LSTM based reduced order modeling of two-dimensional unsteady flows around a circular cylinder at different Reynolds numbers," Fluid Dyn. Res. **52**, 065501 (2020).
[46]C. H. K. Williamson, "Vortex dynamics in the cylinder wake," Annu. Rev. Fluid Mech. **28**, 477–539 (1996).
[47]B. R. Noack and H. Eckelmann, "A global stability analysis of the steady and periodic cylinder wake," J. Fluid Mech. **270**, 297–330 (1994).
[48]L. Sirovich, "Turbulence and the dynamics of coherent structures. I. Coherent structures," Q. Appl. Math. **45**, 561–571 (1987).

13 October 2026 05:32:48

© Author(s) 2026

[49]P. J. Schmid, "Dynamic mode decomposition of numerical and experimental data," J. Fluid Mech. **656**, 5–28 (2010).
[50]K. Taira *et al.*, "Modal analysis of fluid flows: An overview," AIAA J. **55**, 4013–4041 (2017).
[51]D. Kelshaw and L. Magri, "Proper latent decomposition," arXiv:2412.00785 (2024).
[52]G. Min and N. Jiang, "Data-driven identification and pressure fields prediction for parallel twin cylinders based on POD and DMD method," Phys. Fluids **36**, 023614 (2024).
[53]R. Vinuesa, S. L. Brunton, and B. McKeon, "The transformative potential of machine learning for experiments in fluid mechanics," Nat. Rev. Phys. **5**, 536–545 (2023).
[54]K. Azizzadenesheli *et al.*, "Neural operators for accelerating scientific simulations and design," Nat. Rev. Phys. **6**, 320–328 (2024).
[55]D. Barkley and R. D. Henderson, "Three-dimensional floquet stability analysis of the wake of a circular cylinder," J. Fluid Mech. **322**, 215–241 (1996).
[56]C. H. K. Williamson, "Three-dimensional wake transition," J. Fluid Mech. **328**, 345–407 (1996).
[57]T. Leweke and C. H. K. Williamson, "Three-dimensional instabilities in wake transition," Eur. J. Mech. - B/Fluids **17**, 571–586 (1998).
[58]S. Behara and S. Mittal, "Wake transition in flow past a circular cylinder," Phys. Fluids **22**, 114104 (2010).
[59]Y. Zhang, Y. Xiao, R. Liu, and H. Chen, "Aeroacoustic prediction based on large-eddy simulation and the Ffowcs Williams–Hawkings equation," Adv. Aerodyn. **4**, 19 (2022).
[60]C. H. K. Williamson, "Oblique and parallel modes of vortex shedding in the wake of a circular cylinder at low Reynolds numbers," J. Fluid Mech. **206**, 579–627 (1989).
[61]R. Mondal and M. M. Alam, "Blockage effect on wakes of various bluff bodies: A review of confined flow," Ocean Eng. **286**, 115592 (2023).
[62]N. A. Malamataris, I. E. Sarris, D. P. Pazis, and A. Liakos, "A comprehensive study of the onset of boundary layer separation in the unbounded flow around a circular cylinder," Phys. Fluids **34**, 103607 (2022).
[63]G. Chopra and S. Mittal, "The effect of trip wire on transition of boundary layer on a cylinder," Phys. Fluids **34**, 054103 (2022).
[64]M. C. Thompson, A. Radi, A. Rao, J. Sheridan, and K. Hourigan, "Low-Reynolds-number wakes of elliptical cylinders: From the circular cylinder to the normal flat plate," J. Fluid Mech. **751**, 570–600 (2014).
[65]D. Park and K.-S. Yang, "Flow instabilities in the wake of a rounded square cylinder," J. Fluid Mech. **793**, 915–932 (2016).
[66]M. Kouchakzad, A. Sohankar, and M. R. Rastan, "Onset of vortex shedding and hysteresis in flow over tandem sharp-edged cylinders of diverse cross sections," Phys. Fluids **36**, 013604 (2024).
[67]Q. Sun, H. Yuan, H. Liang, X. Zhou, and D. Zhan, "Numerical analysis of the flow over four side-by-side square cylinders with different gaps," Phys. Fluids **36**, 023601 (2024).
[68]J. A. Schetz and A. E. Fuhs, *Fundamentals of Fluid Mechanics* (John Wiley & Sons, New York, 1999), p. 935.
[69]L. Davidson, *Fluid Mechanics, Turbulent Flow and Turbulence Modeling* (Chalmers University of Technology, Gothenburg, Sweden, 2019).
[70]H. Oertel, "Wakes behind blunt bodies," Annu. Rev. Fluid Mech. **22**, 539–564 (1990).
[71]Q. D. Nguyen, W. Lu, L. Chan, A. Ooi, and C. Lei, "A state-of-the-art review of flows past confined circular cylinders," Phys. Fluids **35**, 071301 (2023).
[72]T. Leweke and M. Provansal, "The flow behind rings: Bluff body wakes without end effects," J. Fluid Mech. **288**, 265–310 (1995).
[73]C. Norberg, "Fluctuating lift on a circular cylinder: Review and new measurements," J. Fluids Struct. **17**, 57–96 (2003).
[74]U. Fey, M. König, and H. Eckelmann, "A new Strouhal–Reynolds-number relationship for the circular cylinder in the range $47 < \mathrm{Re} < 2 \times 10^5$," Phys. Fluids **10**, 1547–1549 (1998).
[75]C. H. K. Williamson and G. L. Brown, "A series in $1/\sqrt{\mathrm{Re}}$ to represent the Strouhal–Reynolds number relationship of the cylinder wake," J. Fluids Struct. **12**, 1073–1085 (1998).
[76]C. H. K. Williamson and R. Govardhan, "Vortex-induced vibrations," Annu. Rev. Fluid Mech. **36**, 413–455 (2004).
[77]Y. Yang, H. Li, and X. Jin, "Three-dimensional wake dynamics of a twisted cylinder," Phys. Fluids **36**, 073629 (2024).
[78]G. Li, W. Li, M. J. Janocha, G. Yin, and M. C. Ong, "Large eddy simulations of flow past an inclined circular cylinder: Insights into the three-dimensional effect," Phys. Fluids **35**, 115131 (2023).
[79]M. Verma and A. De, "Flow dynamics of the transversely oscillating tapered circular cylinder under vortex-induced vibrations at low Reynolds number," Phys. Fluids **36**, 103603 (2024).
[80]J. Hansson, S. Sasic, and H. Ström, "Low-frequency wake modulation governs back-side particle deposition on cylinders," J. Fluid Mech. **1029**, A48 (2026).
[81]B. Kumar and S. Mittal, "On the origin of the secondary vortex street," J. Fluid Mech. **711**, 641–666 (2012).
[82]H. Jiang, "On the absence of a secondary vortex street in three-dimensional and turbulent cylinder wakes," J. Fluid Mech. **993**, R1 (2024).
[83]H. Schlichting and K. Gersten, *Boundary-Layer Theory* (Springer, Berlin, Heidelberg, 2017).
[84]L.-R. Zhang *et al.*, "Vortex dynamics induced by a finite wall-mounted cylinder with various corner shapes," Phys. Fluids **36**, 075195 (2024).
[85]A. Mashhadi, A. Sohankar, and M. M. Moradmand, "Three-dimensional wake transition of rectangular cylinders and temporal prediction of flow patterns based on a machine learning algorithm," Phys. Fluids **36**, 094138 (2024).
[86]L. Qiao *et al.*, "Coherence mode and Floquet analysis on flow past a rectangular cylinder with small angle of attack," Phys. Fluids **36**, 123616 (2024).
[87]A. F. Psaros, X. Meng, Z. Zou, L. Guo, and G. E. Karniadakis, "Uncertainty quantification in scientific machine learning: Methods, metrics, and comparisons," J. Comput. Phys. **477**, 111902 (2023).
[88]R. J. Adrian, "Particle-imaging techniques for experimental fluid mechanics," Annu. Rev. Fluid Mech. **23**, 261–304 (1991).
[89]J. Westerweel and F. Scarano, "Universal outlier detection for PIV data," Exp. Fluids **39**, 1096–1100 (2005).
[90]F. Scarano, "Tomographic PIV: Principles and practice," Meas. Sci. Technol. **24**, 012001 (2013).
[91]F. Nemati Taher and A. Subaşı, "A fast three-dimensional flow field prediction around bluff bodies using deep learning," Phys. Fluids **36**, 025101 (2024).
[92]A. Abucide-Armas, K. Portal-Porras, U. Fernandez-Gamiz, E. Zulueta, and A. Teso-Fz-Betoño, "Convolutional neural network predictions for unsteady Reynolds-averaged Navier–Stokes-based numerical simulations," J. Mar. Sci. Eng. **11**, 239 (2023).
[93]P. Lin, C. You, Q. Wang, and G. Hu, "Deep learning-based fast prediction of flow field around multiple bluff bodies," Adv. Wind Eng. **1**, 100020 (2024).
[94]J.-Z. Peng, S. Chen, N. Aubry, Z. Chen, and W.-T. Wu, "Unsteady reduced-order model of flow over cylinders based on convolutional and deconvolutional neural network structure," Phys. Fluids **32**, 123609 (2020).
[95]S. Li, J. Yang, and P. Teng, "Data-driven prediction of cylinder-induced unsteady wake flow," Appl. Ocean Res. **150**, 104114 (2024).
[96]X. Jin, S. Laima, W.-L. Chen, and H. Li, "Time-resolved reconstruction of flow field around a circular cylinder by recurrent neural networks based on non-time-resolved particle image velocimetry measurements," Exp. Fluids **61**, 114 (2020).
[97]S. Laima, X. Zhou, X. Jin, D. Gao, and H. Li, "DeepTRNet: Time-resolved reconstruction of flow around a circular cylinder via spatiotemporal deep neural networks," Phys. Fluids **35**, 015118 (2023).
[98]R. Li *et al.*, "Deep learning reconstruction of high-Reynolds-number turbulent flow field around a cylinder based on limited sensors," Ocean Eng. **304**, 117857 (2024).
[99]J. Yun and J. Lee, "Prediction of near-wake velocity in laminar flow over a circular cylinder using neural networks with instantaneous wall pressure input," Appl. Sci. **13**, 6891 (2023).
[100]Y. Mo and L. Magri, "Reconstructing unsteady flows from sparse, noisy measurements with a physics-constrained convolutional neural network," Phys. Rev. Fluids **10**, 034901 (2025).
[101]P. Xie *et al.*, "A physics-informed deep learning model to reconstruct turbulent wake from random sparse data," Phys. Fluids **36**, 065145 (2024).
[102]X. Jin, M. Xu, Y. Yang, and W. Chen, "Spatiotemporal coupling deep neural network for time-resolved flow field reconstruction around a circular cylinder," Phys. Fluids **36**, 105191 (2024).

13 October 2026 05:32:48

© Author(s) 2026

[103]B. Li *et al.*, "Using machine learning to detect the turbulent region in flow past a circular cylinder," J. Fluid Mech. **905**, A10 (2020).
[104]M. Matsuo, K. Fukami, T. Nakamura, M. Morimoto, and K. Fukagata, "Reconstructing three-dimensional bluff body wake from sectional flow fields with convolutional neural networks," SN Comput. Sci. **5**, 306 (2024).
[105]J. Sirignano and J. F. MacArt, "Deep learning closure models for large-eddy simulation of flows around bluff bodies," J. Fluid Mech. **966**, A26 (2023).
[106]M. Kim, J. Park, and H. Choi, "Large eddy simulation of flow over a circular cylinder with a neural-network-based subgrid-scale model," J. Fluid Mech. **984**, A6 (2024).
[107]S. A. Shojaee, S. Goswami, C. F. Lange, and A. Hemmati, "A neural network approach to improve Reynolds-averaged Navier–Stokes modeling of bluff body wakes," AIP Adv. **15**, 025018 (2025).
[108]Z. Zhang *et al.*, "Turbulence closure in Reynolds-averaged Navier–Stokes and flow inference around a cylinder using physics-informed neural networks and sparse experimental data," J. Fluid Mech. **1034**, A16 (2026).
[109]J. Rabault, M. Kuchta, A. Jensen, U. Réglade, and N. Cerardi, "Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control," J. Fluid Mech. **865**, 281–302 (2019).
[110]H. Tang, J. Rabault, A. Kuhnle, Y. Wang, and T. Wang, "Robust active flow control over a range of Reynolds numbers using an artificial neural network trained through deep reinforcement learning," Phys. Fluids **32**, 053605 (2020).
[111]J. Li and M. Zhang, "Reinforcement-learning-based control of confined cylinder wakes with stability analyses," J. Fluid Mech. **932**, A44 (2022).
[112]M. Chatzimanolakis, P. Weber, and P. Koumoutsakos, "Learning in two dimensions and controlling in three: Generalizable drag reduction strategies for flows past circular cylinders through deep reinforcement learning," Phys. Rev. Fluids **9**, 043902 (2024).
[113]X. Dong, H. Hong, X. Deng, W. Zhong, and G. Hu, "Surrogate model-based deep reinforcement learning for experimental study of active flow control of circular cylinder," Phys. Fluids **35**, 105147 (2023).
[114]C. Xia, J. Zhang, E. C. Kerrigan, and G. Rigas, "Active flow control for bluff body drag reduction using reinforcement learning with partial measurements," J. Fluid Mech. **981**, A17 (2024).
[115]Q. Wang *et al.*, "Dynamic feature-based deep reinforcement learning for flow control of circular cylinder with sparse surface pressure sensing," J. Fluid Mech. **988**, A4 (2024).
[116]J. Chen, E. Ballini, and S. Micheletti, "Active flow control for bluff body under high Reynolds number turbulent flow conditions using deep reinforcement learning," Phys. Fluids **37**, 025111 (2025).
[117]K. Li, H. Li, S. Li, and Z. Chen, "Fully convolutional neural network prediction method for aerostatic performance of bluff bodies based on consistent shape description," Appl. Sci. **12**, 3147 (2022).
[118]S. Brahmachary and N. Thuerey, "Unsteady cylinder wakes from arbitrary bodies with differentiable physics-assisted neural network," Phys. Rev. E **109**, 055304 (2024).
[119]V. Oommen, S. Khodakarami, A. Bora, Z. Wang, and G. E. Karniadakis, "Learning turbulent flows with generative models for super resolution and sparse flow reconstruction," Nat. Commun. **17**, 3707 (2026).
[120]K. Fukami, K. Fukagata, and K. Taira, "Super-resolution reconstruction of turbulent flows with machine learning," J. Fluid Mech. **870**, 106–120 (2019).
[121]M. Amorós-Trepat, L. Medrano-Navarro, Q. Liu, L. Guastoni, and N. Thuerey, "Guiding diffusion models to reconstruct flow fields from sparse data," Phys. Fluids **38**, 015112 (2026).
[122]K. Fukami, K. Fukagata, and K. Taira, "Machine-learning-based spatio-temporal super resolution reconstruction of turbulent flows," J. Fluid Mech. **909**, A9 (2021).
[123]Y.-Q. Zhang *et al.*, "High-fidelity flow-field reconstruction from sparse sensors using a hybrid graph neural network-transformer architecture," Phys. Fluids **38**, 023610 (2026).
[124]J. H. Harmening, F.-J. Peitzmann, and O. el Moctar, "Effect of network architecture on physics-informed deep learning of the Reynolds-averaged turbulent flow field around cylinders without training data," Front. Phys. **12**, 1385381 (2024).
[125]Y. Dou, X. Han, and P. Lin, "Flow field reconstruction and prediction of the two-dimensional cylinder flow using data-driven physics-informed neural network combined with long short-term memory," Eng. Appl. Artif. Intell. **149**, 110547 (2025).
[126]H. Jiang and L. Cheng, "Large-eddy simulation of flow past a circular cylinder for Reynolds numbers 400 to 3900," Phys. Fluids **33**, 034119 (2021).
[127]A. Ooi *et al.*, "Turbulent flow over a cylinder confined in a channel at Re = 3900," Int. J. Heat Fluid Flow **96**, 108982 (2022).
[128]K. Manohar, B. W. Brunton, J. N. Kutz, and S. L. Brunton, "Data-driven sparse sensor placement for reconstruction: Demonstrating the benefits of exploiting known patterns," IEEE Control Syst. **38**, 63–86 (2018).
[129]J.-C. Loiseau, B. R. Noack, and S. L. Brunton, "Sparse reduced-order modelling: Sensor-based dynamics to full-state estimation," J. Fluid Mech. **844**, 459–490 (2018).
[130]R. J. Adrian, "Twenty years of particle image velocimetry," Exp. Fluids **39**, 159–169 (2005).
[131]B. Wieneke, "Volume self-calibration for 3D particle image velocimetry," Exp. Fluids **45**, 549–556 (2008).
[132]K. Duraisamy, G. Iaccarino, and H. Xiao, "Turbulence modeling in the age of data," Annu. Rev. Fluid Mech. **51**, 357–377 (2019).
[133]K. Duraisamy, "Perspectives on machine learning-augmented Reynolds-averaged and large eddy simulation models of turbulence," Phys. Rev. Fluids **6**, 050504 (2021).
[134]S. Mondal and S. Sarkar, "Multi-fidelity prediction of spatiotemporal fluid flow," Phys. Fluids **34**, 087112 (2022).
[135]W. Lu, W. K. Lam, and T. Zahtila, "An autoencoder-based bi-fidelity method utilising frozen latent spaces. Application to flow past a confined cylinder," Int. J. Heat Fluid Flow **118**, 110192 (2026).
[136]M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," J. Comput. Phys. **378**, 686–707 (2019).
[137]B. R. Noack, M. Morzyński, and G. Tadmor, *Reduced-Order Modelling for Flow Control* (Springer Vienna, 2011).
[138]J. N. Kutz, S. L. Brunton, B. W. Brunton, and J. L. Proctor, *Dynamic Mode Decomposition* (Society for Industrial and Applied Mathematics, Philadelphia, PA, 2016).
[139]S. L. Brunton, J. L. Proctor, and J. N. Kutz, "Discovering governing equations from data by sparse identification of nonlinear dynamical systems," Proc. Natl. Acad. Sci. U. S. A. **113**, 3932–3937 (2016).
[140]K. Champion, B. Lusch, J. N. Kutz, and S. L. Brunton, "Data-driven discovery of coordinates and governing equations," Proc. Natl. Acad. Sci. U. S. A. **116**, 22445–22451 (2019).
[141]C. W. Rowley, I. Mezic, S. Bagheri, P. Schlatter, and D. S. Henningson, "Spectral analysis of nonlinear flows," J. Fluid Mech. **641**, 115–127 (2009).
[142]G. E. Karniadakis *et al.*, "Physics-informed machine learning," Nat. Rev. Phys. **3**, 422–440 (2021).
[143]Z. Li *et al.*, "Fourier neural operator for parametric partial differential equations," arXiv:2010.08895 (2020).
[144]J. Ling, A. Kurzawski, and J. Templeton, "Reynolds averaged turbulence modelling using deep neural networks with embedded invariance," J. Fluid Mech. **807**, 155–166 (2016).
[145]J.-X. Wang, J.-L. Wu, and H. Xiao, "Physics-informed machine learning approach for reconstructing Reynolds stress modeling discrepancies based on DNS data," Phys. Rev. Fluids **2**, 034603 (2017).
[146]S. L. Brunton and B. R. Noack, "Closed-loop turbulence control: Progress and challenges," Appl. Mech. Rev. **67**, 050801 (2015).
[147]J. Viquerat, P. Meliga, A. Larcher, and E. Hachem, "A review on deep reinforcement learning for fluid mechanics: An update," Phys. Fluids **34**, 111301 (2022).
[148]T. Duriez, S. L. Brunton, and B. R. Noack, *Machine Learning Control: Taming Nonlinear Dynamics and Turbulence* (Springer International Publishing, Cham, 2017).
[149]H. Choi, W.-P. Jeon, and J. Kim, "Control of flow over a bluff body," Annu. Rev. Fluid Mech. **40**, 113–139 (2008).
[150]L. N. Cattafesta III and M. Sheplak, "Actuators for active flow control," Annu. Rev. Fluid Mech. **43**, 247–272 (2011).
[151]T. R. Bewley, "Flow control: New challenges for a new Renaissance," Prog. Aerosp. Sci. **37**, 21–58 (2001).
[152]I. Gulrajani and D. Lopez-Paz, "In search of lost domain generalization," arXiv:2007.01434 (2020).

13 October 2026 05:32:48

© Author(s) 2026

[153]Z. Xu, S. Jain, and M. Kankanhalli, "Hallucination is inevitable: An innate limitation of large language models," arXiv:2401.11817 (2025).
[154]K. Xue, Q. Li, and L. Zhao, "Compressibility effect on flow characteristics over a circular cylinder at Reynolds number of 3900," Phys. Fluids **36**, 085165 (2024).
[155]Q. Liu, S. Gong, G. Xie, H. Lu, and Z. Zuo, "Effects of free tip shape and aspect ratio on flow around a finite cylinder with two free tips at subcritical Reynolds numbers," Comput. Fluids **270**, 106147 (2024).
[156]H. Strandenes, B. Pettersen, H. I. Andersson, and M. Manhart, "Influence of spanwise no-slip boundary conditions on the flow around a cylinder," Comput. Fluids **156**, 48–57 (2017).
[157]W. Lu, D. Aljubaili, T. Zahtila, L. Chan, and A. Ooi, "Asymmetric wakes in flows past circular cylinders confined in channels," J. Fluid Mech. **958**, A8 (2023).
[158]D. Durstewitz, C. J. Hemmer, F. Hess, C. R. Doll, and L. Eisenmann, "Position: A dynamical systems perspective is needed to advance time series modeling," arXiv:2602.16864 (2026).
[159]T. Déda, W. Wolf, and S. T. M. Dawson, "Neural networks in feedback for flow analysis and control," Phys. Rev. Fluids **9**, 063904 (2024).
[160]N. McGreivy and A. Hakim, "Weak baselines and reporting biases lead to overoptimism in machine learning for fluid-related partial differential equations," Nat. Mach. Intell. **6**, 1256–1269 (2024).
[161]G. D. Weymouth, "Data-driven multi-grid solver for accelerated pressure projection," Comput. Fluids **246**, 105620 (2022).
[162]P. Sousa, A. Afonso, and C. V. Rodrigues, "Surrogate-based pressure–velocity coupling: Accelerating incompressible CFD flow solvers with machine learning," Comput. Fluids **318**, 107216 (2026).
[163]P. Sousa, C. V. Rodrigues, and A. Afonso, "Enhancing CFD solver with machine learning techniques," Comput. Methods Appl. Mech. Eng. **429**, 117133 (2024).
[164]P. Sousa, A. Afonso, and C. Veiga Rodrigues, "Application of machine learning to model the pressure Poisson equation for fluid flow on generic geometries," Neural Comput. Appl. **36**, 16581–16606 (2024).

13 October 2026 05:32:48

© Author(s) 2026