The Transformative Role of Artificial Intelligence in Fluid Dynamics: A Critical Literature Review

Summary See: https://circularastronomy.com/2025/09/26/artificial-intelligence-in-fluid-dynamics-a-literature-review-of-applications-challenges-and-future-directions/

I. Introduction: A New Paradigm for an Enduring Challenge

1.1 The Computational Bottleneck of Fluid Dynamics

Fluid dynamics is a foundational discipline in science and engineering, governing phenomena from atmospheric circulation and blood flow to the aerodynamics of an aircraft wing.[1] For centuries, the field has been governed by a set of complex, nonlinear partial differential equations known as the Navier-Stokes equations, which describe the motion of viscous fluids.[3] While these equations are the cornerstone of modern fluid mechanics, solving them at scale for real-world problems remains a formidable, and often intractable, challenge. The primary obstacle is the computational cost required to resolve the wide range of length and time scales, particularly in turbulent flows. The computational expense of a Direct Numerical Simulation (DNS), which resolves all features of the flow, scales roughly as

Re3, where Re is the Reynolds number.[2] This means that a mere tenfold increase in the Reynolds number can lead to a thousandfold increase in computational cost, rendering high-fidelity simulations for large-scale systems like airplane design or climate prediction computationally unfeasible with current technology.[2]

To circumvent this computational bottleneck, engineers and scientists have historically relied on approximations and models that trade accuracy for efficiency. These include methods like Reynolds-Averaged Navier-Stokes (RANS) and Large-Eddy Simulation (LES), which are widely used in industrial applications due to their relative tractability.[3] However, these models often rely on empirical relations and simplifying assumptions, which can reduce their accuracy in complex flow regimes and limit their universality.[4] The inability to devise a single, universal model for all turbulence flows is a persistent challenge that has long defined the field.[5] As a result, a fundamental trade-off has existed between accuracy and computational tractability, a dilemma that has historically limited the scope of what is possible in both scientific inquiry and engineering design.[2]

The emergence of artificial intelligence (AI) and machine learning (ML) presents a new and revolutionary approach to this enduring challenge. As computational tools have evolved from pen-and-paper theory to computer-aided modeling, the advent of AI represents the latest chapter in this progression, driven by the proliferation of extensive datasets from advanced sensors and large-scale simulations.[1] This convergence of computational demands and abundant data has created a fertile ground for AI to not only accelerate existing methods but also to fundamentally transform how fluid dynamics problems are approached.

1.2 The Emergence of Data-Driven Approaches

For many decades, the study of fluid dynamics has been supported by three foundational pillars: theoretical analysis, experimental measurement, and computational fluid dynamics (CFD).[6] In recent years, data-driven fluid dynamics has solidified its place as a powerful fourth pillar, leveraging the unprecedented volumes of data from simulations and experiments to extract insights and develop new models.[1] This new methodological framework provides a robust information-processing capability that can augment and even transform current lines of research and industrial applications.[7]

The integration of AI techniques into this fourth pillar signifies a major shift in methodology. Rather than being mere enhancements to existing computational tools, machine learning, deep learning (DL), and reinforcement learning (RL) are being applied as powerful frameworks to solve problems that are not amenable to traditional analytical or numerical treatment.[5] These AI methods are distinct in their ability to learn complex, nonlinear relationships directly from data, enabling them to make rapid predictions and identify hidden patterns.[9] The initial applications of this new toolkit focused on creating computationally inexpensive surrogate models, while more recent work has moved towards a deeper integration with the underlying physics. This evolution demonstrates a maturing understanding of the problem space, moving from AI as a superficial augment to AI as a deeply integrated component of physical modeling.[2]

II. Evolution of AI in Fluid Dynamics: From Augmentation to Integration

2.1 The Dawn of Computational Fluid Dynamics (CFD)

The historical trajectory of fluid dynamics is marked by a continuous evolution of its tools and methodologies. The field's theoretical foundation was laid by Claude-Louis Navier and George Gabriel Stokes in 1822 with the formulation of their eponymous equations.[3] Decades later, the groundwork for numerical solutions was established by Lewis Fry Richardson in 1928, whose pioneering work on numerical weather prediction introduced the finite difference method.[3] The true revolution in CFD, however, was catalyzed by the advent of digital computers in the mid-20th century. This technological leap made it possible to tackle increasingly complex problems, leading to the development of early CFD software and the introduction of crucial turbulence models, such as the k-ε model in the 1980s.[3] These advancements enabled more accurate simulations of real-world flows and expanded the reach of CFD into critical engineering and design problems.[3]

2.2 The Rise of Data-Driven Surrogates

Early applications of machine learning in fluid dynamics were largely utilitarian, focused on developing surrogate models to accelerate computationally expensive tasks. A surrogate model is a fast, approximate replica of a high-fidelity simulation or physical experiment.[11] These models, which can be based on techniques like regression or classification, learn the relationship between a set of input parameters and an output variable (e.g., lift, drag, or temperature).[10] This approach allows for the rapid exploration of vast design spaces, a task that would be prohibitively time-consuming with traditional CFD simulations.[9] Engineers could use these surrogates to quickly identify promising designs, which would then be validated with a small number of expensive, full-fidelity simulations. Techniques such as Gaussian Processes (Kriging) and Radial Basis Functions (RBFs) were employed to interpolate or approximate system responses from scattered data points, effectively balancing the need for accuracy with the computational budget.[11] This initial phase of AI integration was characterized by its agnosticism to the underlying physics, relying solely on input-output data pairs to learn the system behavior.[11]

2.3 The Hybrid and Physics-Informed Revolution

A major paradigm shift has occurred as the field moved beyond purely data-driven, "black-box" models to more sophisticated hybrid approaches.[2] This evolution was driven by the recognition that while pure ML models can be extremely efficient, they often struggle with fundamental physical constraints, such as conservation of momentum, and may not generalize well to unseen conditions.[2] A significant portion of the literature and ongoing development now focuses on methods that combine the best aspects of traditional physics-based solvers with the speed of machine learning.[2]

One of the most prominent examples of this revolution is the development of Physics-Informed Neural Networks (PINNs).[12] Unlike conventional neural networks that learn solely from large datasets, PINNs embed the governing physical laws directly into their loss function.[12] During training, the network's output is continuously checked against the expectations of the physical equations, and it learns by minimizing the "residual," or the amount by which its solution fails to satisfy those equations.[12] This approach promises a viable alternative to classical numerical methods for solving partial differential equations (PDEs), with the potential to leverage future hardware like GPUs and quantum computers more effectively.[13]

A complementary hybrid approach involves using AI to improve specific components of traditional solvers. For instance, ML can be used to discover improved spatial discretizations on a coarse grid, effectively making an under-resolved simulation as accurate as a traditional solver with a much finer mesh.[2] This strategy, exemplified by the work of Kochkov et al. in

PNAS (2021) [2], allows for significant computational speedups while preserving the stability and predictable generalization properties that are often lacking in purely data-driven models.[2] This shift reflects a more mature understanding of AI's role, repositioning it not as a replacement for traditional methods, but as a powerful tool for their augmentation and enhancement.[2]

III. Core Research Themes and Seminal Contributions

The integration of AI has created several distinct and vibrant research themes within fluid dynamics. Each theme addresses a specific set of challenges, and together, they illustrate the breadth of AI's transformative potential. The following table provides a summary of landmark studies that have significantly shaped the field.

Study Title	Authors & Year	Core Finding
Machine learning–accelerated computational fluid dynamics	Kochkov et al., 2021	Hybrid ML models can improve accuracy and speed (40-80x speedup) on coarse grids while maintaining stability and generalizing to new Reynolds numbers.[2]
Machine Learning Strategies for Upgrading Turbulence Modelling	Billard et al., 2025	A neural network-based correction to a RANS model successfully mitigated shortcomings for separated channel flow, significantly improving the prediction of the separation region.[16]
Discovering new solutions to century-old problems in fluid dynamics	DeepMind et al. (Brown, NYU, Stanford), 2025	Used high-precision Physics-Informed Neural Networks (PINNs) to systematically discover new families of unstable singularities ("blow-ups") in the Navier-Stokes equations, a long-standing mathematical problem.[12]
Large Language Model Driven Development of Turbulence Models	[Authors not listed], 2025	Demonstrated that a large language model can act as a collaborator in scientific

		discovery, proposing and refining physically interpretable turbulence models that outperform baselines.[17]

3.1 AI for Turbulence Modeling and Simulation Acceleration

The accurate modeling of turbulence remains one of the greatest unsolved problems in classical physics.[5] The high-fidelity simulation of turbulent flows, particularly for industrial applications, is often computationally prohibitive.[2] AI is now being used to address this by either augmenting traditional turbulence models or accelerating the simulation process itself.

One approach is to use AI to enhance existing models like RANS or LES. For example, a neural network can be trained on high-fidelity DNS data to compute a correction to the turbulent stress of a RANS model, thereby mitigating its shortcomings in complex flow regimes.[16] The work by Billard et al. (2025) demonstrated that an NN-enhanced RANS model could significantly improve the prediction of flow separation, even when tested on cases with different geometries or Reynolds numbers than those it was trained on.[16] A more recent and promising development involves the use of large language models (LLMs) to reason about and synthesize new turbulence models.[17] Instead of a "black-box" model, this approach uses an LLM in a closed-loop, iterative workflow to propose, refine, and reason about physically interpretable turbulence models that outperform baselines.[17]

A second major thrust is the use of AI for simulation acceleration. The work of Kochkov et al. (2021) is a seminal contribution in this area. Their research introduced a method for learning a numerical solver that can achieve the same accuracy as a traditional solver on a fine grid but by using a much coarser mesh.[2] By replacing the components of the traditional solver most affected by the loss of resolution with learned alternatives, their method achieved an impressive 40- to 80-fold computational speedup.[2] This hybrid approach maintains stability during long simulations and demonstrates robust generalization to novel forcing functions and Reynolds numbers, a critical feature often lacking in purely data-driven methods.[2] This work exemplifies how scientific computing can strategically leverage AI and hardware accelerators to improve simulations without sacrificing accuracy or generalizability.[2]

3.2 Physics-Informed Neural Networks (PINNs) as PDE Solvers

Physics-Informed Neural Networks (PINNs) are a class of models that have demonstrated success not only in engineering applications but also in tackling foundational scientific problems. Their core strength lies in their ability to embed physical laws directly into their loss function, which allows them to solve partial differential equations without requiring massive amounts of labeled data.[12] Instead, the network learns by minimizing its "residual"—the amount by which its solution fails to satisfy the governing equations, such as the compressible Euler equations or the Navier-Stokes equations.[12] This framework offers a compelling alternative to classical discretization methods, with potential advantages for implementation on emerging hardware like GPUs and quantum computers.[13]

A landmark achievement that highlights the profound potential of this approach is the work by Google DeepMind and its collaborators (2025).[12] They used high-precision PINNs to systematically discover an entirely new family of mathematical "blow-ups," or singularities, in some of the most complex equations that describe fluid motion.[12] Finding a singularity in the Navier-Stokes equations is a Millennium Prize Problem, and this breakthrough demonstrates that AI can be used as a new instrument for foundational scientific discovery.[12] By embedding mathematical insights into the PINN's training, the researchers were able to capture unstable singularities that had long eluded conventional methods.[12] This work signifies a new era of "computer-assisted mathematics," where AI can tackle long-standing challenges that have resisted centuries of human effort.[12]

3.3 Machine Learning for Flow Control and Optimization

Altering the natural dynamics of fluid flows is desirable for a vast range of engineering applications, from reducing aircraft fuel consumption to preventing structural damage from flow-induced oscillations.[19] AI is proving to be an exceptionally powerful tool for both active and passive flow control. Passive flow control involves fixed geometric modifications, such as wing shape optimization, while active flow control requires energy input and real-time actuation.[8]

Machine learning offers two primary paradigms for active flow control: model-free and gradient-based methods.[8] Model-free approaches, such as Reinforcement Learning (RL), treat the control problem as a black-box optimization.[8] An agent learns the optimal control strategy through a trial-and-error process, interacting with the fluid environment to maximize a reward signal, such as drag reduction.[8] These methods are promising for problems that are not easily amenable to analytical treatment.[8] In contrast, gradient-based methods, such as the deep learning partial differential equation augmentation method (DPM), use adjoints of the governing equations to compute the end-to-end sensitivities needed for optimization.[19] A study comparing these approaches found that the DPM-based controller was significantly more effective and computationally less intensive to train than its DRL-based counterpart, indicating the value of incorporating physics into the optimization process.[19] Examples of successful applications include reducing drag and stabilizing vortex shedding in flows over a cylinder.[8]

3.4 AI for Data-Driven Fluid Mechanics Fundamentals

Beyond direct simulation and control, AI is also enabling new ways to analyze and understand fluid dynamics from a fundamental perspective. One key application is reduced-order modeling (ROM), which aims to capture the essential dynamics of a complex system in a low-dimensional representation.[6] Nonlinear machine learning methods can achieve superb data compression, creating a compact representation of a flow field that can be used for faster simulations and control.[6] For example, a Gaussian Mixture Variational Autoencoder (GMVAE) can encode high-dimensional flow data into a low-dimensional latent space that is both physically meaningful and globally consistent.[20]

AI is also being used for real-time prediction and forecasting, a process often referred to as "nowcasting".[1] This involves leveraging the ability of AI to extract dynamic information from time-sequence data, a task that goes beyond static image processing.[5] Examples of this include using AI for short-term storm forecasting in Africa, predicting space weather, and modeling blood flow for medical diagnostics.[1] These applications demonstrate the power of AI to synthesize vast, dynamic datasets and provide rapid, actionable insights that would be impossible with traditional methods.[9]

IV. Key Debates and Controversies

The rapid growth of AI in fluid dynamics has been accompanied by a number of technical and philosophical debates that are central to the field's maturation. These discussions address fundamental questions about model robustness, interpretability, and the role of data.

Debate	Conflicting Viewpoints

Generalization vs. Specificity	**Hybrid Models:** Proponents argue that models that incorporate physics can generalize to different forcing functions and Reynolds numbers outside of the training data.[2] This is because the underlying physical laws provide a robust framework that extends beyond the training distribution.	Pure ML Models: Detractors note that purely data-driven "black-box" models, which lack physics constraints, often fail to generalize. They perform poorly on unseen data and cannot enforce fundamental laws like conservation of momentum, as their "logic" is entirely dependent on the training data and they have no inherent ability to extrapolate.2
The Black Box Problem & Interpretability	**Interpretability:** This viewpoint asserts that models must be physically interpretable to be trustworthy for critical engineering applications.[5] A lack of understanding of the underlying physical reasons for a prediction is a major hurdle to adoption, especially in safety-critical domains. This has spurred the development of Explainable AI (XAI) methods to link model predictions to physical features.[4]	Performance: This perspective prioritizes speed and accuracy over physical interpretation.21 For some problems, the utility of the model's output (e.g., reduced drag) is more important than understanding the inner workings of the complex neural network that produced it.21 The complexity of deep learning can make interpretability difficult, and forcing it can sometimes reduce model performance.5
PINNs vs. Traditional CFD	**PINNs:** Proponents argue that PINNs are a viable alternative to classical	Traditional CFD: Detractors

	methods, especially on future hardware, as they can directly solve PDEs.[13] The discovery of fluid singularities with PINNs demonstrates their unique ability to solve problems that have long resisted traditional methods.[12]	argue that PINNs may fail in complex physical phenomena like turbulence compared to traditional numerical methods.5 They note that PINNs can take much longer to "converge" and that the calculated weights are only valid for that specific domain and boundary conditions, lacking the concept of generalization or inference that is critical for engineering use cases.21

4.1 Generalization vs. Specificity

The capacity of an AI model to perform well on data it has not been trained on—a property known as generalization—is a central point of contention.[21] The provided research highlights a clear divide between "pure" machine learning models and "hybrid" approaches.[2] Pure ML models that aim to replace an entire simulation can be extremely efficient but often exhibit poor generalization because they do not explicitly enforce physical constraints like conservation of mass and momentum.[2] Their performance outside the specific distribution of their training data is often unpredictable and unreliable.[21] This represents a significant limitation for engineering applications, where models must perform robustly under a wide range of operating conditions.[2] In contrast, hybrid approaches that strategically integrate AI within a traditional physics-based framework are shown to be more stable and demonstrate robust generalization properties, even to different Reynolds numbers and forcing functions.[2] This suggests that a more nuanced, integrated approach is required to build models that are both fast and universally applicable.

4.2 The Black Box Problem and the Need for Interpretability

A significant hurdle to the widespread adoption of AI-augmented CFD is the "black box" nature of many deep learning models.[5] In critical engineering applications, a prediction's physical basis is as important as the prediction itself, as engineers need to understand

why a model behaves in a certain way to ensure safety and reliability.[5] This lack of interpretability creates a trust deficit, preventing the models from moving beyond proofs of concept to become trusted industrial tools.[11] To address this, the field is seeing a growing focus on Explainable AI (XAI).[4] Methods such as additive-feature-attribution techniques are being developed to link the input features of a model to its predictions, providing a physically meaningful interpretation of the relationships it has learned from data.[4] Furthermore, the recent demonstration that a large language model can develop physically interpretable turbulence models with clear reasoning suggests a path forward where AI is not just a black box but a collaborative partner in scientific discovery.[17]

4.3 Data Requirements and Algorithmic Challenges

The success of data-driven methods is fundamentally tied to the availability and quality of data, but this presents a paradox in fluid dynamics.[5] For many biomedical and engineering problems, the data is not "massive," which poses limitations for training data-hungry models.[5] However, when dealing with turbulence, the sheer scale of high-fidelity simulations can generate so much data that issues of storage, retrieval, and post-processing become significant challenges.[5] The data paradox also underscores a lack of standardization. The absence of high-fidelity, open databases for the scientific community hinders knowledge sharing and the development of common benchmarks for model evaluation.[5] This highlights the need for a strategic approach to data, focusing not just on obtaining "more data," but on curating "smarter engineering data" that is relevant to the problem at hand.[10]

V. Gaps in the Literature and Future Research Directions

Despite the significant progress, several critical gaps remain in the literature, which serve as a roadmap for future research.

5.1 Bridging the Gap from Theory to Practice

Many of the AI models developed for fluid dynamics remain at the "proof-of-concept" stage.[16] While they have demonstrated remarkable capabilities on canonical problems, there is a persistent challenge in translating these promising results into robust, industrially scalable tools that can handle the complex geometries and operational variability of real-world applications.[10] Future research must focus on the operational challenges of integrating AI models into existing engineering workflows, including issues of database management, feature engineering, and the coordination of different engineering groups.[10] There is a critical need for more research focused on the scalability and reliability of these models in a diverse range of conditions.

5.2 The Need for Standardized Datasets and Benchmarks

A recurring theme in the provided literature is the scarcity of high-fidelity, open data, which limits the development of robust AI models and hinders progress in the field.[5] To accelerate research and foster collaboration, a concerted effort is needed to create and maintain publicly available databases that are representative of a wide range of flow scenarios, including complex geometries and multi-directional forces.[5] The creation of such standardized datasets would allow for fair and reproducible comparisons between different AI methodologies, establishing clear benchmarks for model performance and generalizability.[5]

5.3 Advancing Physics-Constrained and Explainable AI

The debates around generalization and interpretability point to a clear direction for future work. Research should move beyond purely data-driven "black-box" models and focus on developing more advanced hybrid frameworks that are not only accurate and efficient but also physically consistent and interpretable.[6] This includes the continued development of Physics-Informed Neural Networks (PINNs) to handle the complexities of turbulence and other nonlinear phenomena.[5] Furthermore, the wider application of Explainable AI (XAI) methods is essential to make AI models more trustworthy and to provide deeper insights into the underlying physics of fluid flows.[4] The goal is to develop models that can provide accurate data-derived physics-based equations, rather than relying on empirical relations.[5]

5.4 The Role of Foundation Models and LLMs

The recent breakthrough in using a large language model to reason about and propose novel turbulence models suggests a new and profound frontier for AI in fluid dynamics.[17] This work indicates that AI is no longer limited to simply augmenting human capabilities but can act as an "equal partner" in scientific discovery, capable of performing complex, long-chain reasoning and hypothesis generation.[17] Future research should explore how these foundation models can be leveraged for other grand challenges in fluid dynamics, such as creating self-learning simulations or autonomously discovering new physical laws. This represents a potential paradigm shift, moving the field from computer-assisted design to computer-assisted scientific discovery.

VI. Conclusion

The integration of artificial intelligence into fluid dynamics marks a transformative new era for a field that has long been constrained by computational limitations. This literature review has shown that AI is not a fleeting trend but a fundamental shift in methodology, evolving from early, utilitarian applications in surrogate modeling to a deeply integrated, physics-informed approach that is now tackling some of the discipline's most enduring challenges. The field has progressed from using AI as a tool to accelerate existing methods to leveraging it as a novel instrument for foundational scientific discovery, as evidenced by the breakthrough work on fluid singularities in the Navier-Stokes equations.

Despite these remarkable advances, the field is at a transitional stage, grappling with critical debates concerning generalization, interpretability, and the operational challenges of moving from proof-of-concept models to scalable, industrial tools. The analysis highlights a clear consensus: future progress depends on a strategic move toward hybrid, physically-informed, and explainable AI models. The development of standardized, open datasets and the exploration of new AI paradigms, such as the use of large language models for scientific discovery, are essential to accelerating this progress. Ultimately, AI offers a new framework for overcoming the computational trade-offs of the past and opens up unprecedented opportunities for innovation in science and engineering. This review establishes a critical need for new research that bridges the gap between theoretical AI advancements and practical application, a challenge that is at the heart of the research undertaken in this thesis.

Works cited

1. AI in Fluid Dynamics Research at Leeds, accessed on September 21, 2025,

https://fluids.leeds.ac.uk/ai-in-fluid-dynamics-research-at-leeds/
2. Machine learning–accelerated computational fluid dynamics - PNAS, accessed on September 21, 2025, https://www.pnas.org/doi/10.1073/pnas.2101784118
3. The Evolution of Computational Fluid Dynamics: From Theory to SuperCFD - Simularge, accessed on September 21, 2025, https://www.simularge.com/blog/the-evolution-of-computational-fluid-dynamics-from-theory-to-supercfd
4. Additive-feature-attribution methods: a review on explainable artificial intelligence for fluid dynamics and heat transfer - arXiv, accessed on September 21, 2025, https://arxiv.org/html/2409.11992v1
5. Can Artificial Intelligence Accelerate Fluid Mechanics Research?, accessed on September 21, 2025, https://www.mdpi.com/2311-5521/8/7/212
6. Machine learning in fluid dynamics: A critical assessment - arXiv, accessed on September 21, 2025, https://arxiv.org/html/2508.13430v2
7. [1905.11075] Machine Learning for Fluid Mechanics - arXiv, accessed on September 21, 2025, https://arxiv.org/abs/1905.11075
8. Comparative analysis of machine learning methods for active flow ..., accessed on September 21, 2025, https://www.cambridge.org/core/journals/journal-of-fluid-mechanics/article/comparative-analysis-of-machine-learning-methods-for-active-flow-control/DF06699ACFFCCD1B5778ED63DFDADCFF
9. Applying Machine Learning in CFD to Accelerate Simulation | Neural Concept, accessed on September 21, 2025, https://www.neuralconcept.com/post/applying-machine-learning-in-cfd-to-accelerate-simulation
10. 4 Myths about AI in CFD - Simcenter - Siemens Digital Industries Software Blogs, accessed on September 21, 2025, https://blogs.sw.siemens.com/simcenter/4-myths-about-ai-in-cfd/
11. (PDF) AI-Augmented Computational Fluid Dynamics: Surrogate ..., accessed on September 21, 2025, https://www.researchgate.net/publication/392623231_AI-Augmented_Computational_Fluid_Dynamics_Surrogate_Models_for_Real-Time_Simulation
12. Discovering new solutions to century-old problems in fluid dynamics ..., accessed on September 21, 2025, https://deepmind.google/discover/blog/discovering-new-solutions-to-century-old-problems-in-fluid-dynamics/
13. Adopting Computational Fluid Dynamics Concepts for Physics ..., accessed on September 21, 2025, https://arc.aiaa.org/doi/10.2514/6.2025-0269
14. Machine learning–accelerated computational fluid dynamics - Scite, accessed on September 21, 2025, https://scite.ai/reports/machine-learning-accelerated-computational-fluid-dynamics-zRNm8VK5
15. Dmitrii Kochkov - Google Scholar, accessed on September 21, 2025, https://scholar.google.fr/citations?user=MSRmkssAAAAJ&hl=ja
16. Machine Learning Strategies for Upgrading Turbulence Modelling in ..., accessed

on September 21, 2025, https://arc.aiaa.org/doi/10.2514/6.2025-2048
17. Large Language Model Driven Development of Turbulence Models, accessed on September 21, 2025, https://arxiv.org/abs/2505.01681
18. AI Just Solved a 100 Year Old Million Dollar Science Mystery - YouTube, accessed on September 21, 2025, https://www.youtube.com/watch?v=i2uUFpHNsL8
19. Adjoint-based machine learning for active flow control | Phys. Rev. Fluids, accessed on September 21, 2025, https://link.aps.org/doi/10.1103/PhysRevFluids.9.013901
20. [2502.02605] Physically Interpretable Representation and Controlled Generation for Turbulence Data - arXiv, accessed on September 21, 2025, https://arxiv.org/abs/2502.02605
21. What are your opinions on AI for CFD? - Reddit, accessed on September 21, 2025, https://www.reddit.com/r/CFD/comments/1karur1/what_are_your_opinions_on_ai_for_cfd/
22. [2409.11992] Additive-feature-attribution methods: a review on explainable artificial intelligence for fluid dynamics and heat transfer - arXiv, accessed on September 21, 2025, https://arxiv.org/abs/2409.11992
23. A Pioneering Neural Network Method for Efficient and Robust Fluid Simulation - arXiv, accessed on September 21, 2025, https://arxiv.org/html/2412.10748v3