

Hybrid AI-Simulation for Aerodynamics Applications

Denis Gueyffier

ONERA – The French Aerospace Lab, Direction Scientifique Générale, Institut Polytechnique de Paris, Palaiseau, France

Abstract

Machine learning has moved from an exploratory topic to a structuring force in computational aerodynamics: GPU-native high-fidelity solvers, large open datasets, and early efforts towards foundation models for computational fluid dynamics (CFD). This paper presents a complementary route pursued at ONERA, the French Aerospace Lab: anchoring learning inside classical numerical methods rather than around them, so as to merge the accuracy and interpretability of physics-based solvers with the agility of modern AI. We review scientific machine learning (SciML) for aerodynamics at ONERA, spanning its production and next-generation solvers (elsA, CEDRE, SoNICS, CODA), European and national programmes, and open benchmark databases. We then focus on three research lines. The first embeds learned components at the discretization level of finite-volume solvers: data-driven flux discretizations for hyperbolic conservation laws, and adaptive gradient recovery with accuracy gains of 20–60% on coarse unstructured meshes for 2D Euler, while the finite-volume update, and hence conservation, is retained. The second builds guarantees into the scheme itself, on steady supersonic flow over obstacles: the network sets the gradient and the limiter only inside sets in which no inadmissible output exists, so admissibility holds for any network weights. There, the gain comes from the learned gradient rather than the learned limiter, with 38% lower error on an unseen geometry at a sixth of the price of a mesh refinement. The third keeps the exactly known part of the solution operator exact and learns only the remainder: equivariant boundary corrections around a fast integral core, then a differentiable solver for incompressible flow around bodies, whose predicted velocity is divergence-free to machine precision for any learned weights. There, a controlled study finds that the encoding of the geometry, not model capacity, governs generalization to shapes never seen in training. New results under review complete the picture with the allocation question of when a fixed simulation budget is best spent on finer grids, learned closures, or neural emulators. Physics-anchored approaches of this kind are data-frugal, more amenable to verification and certification, and complementary to foundation-model efforts. A roadmap for their integration into industrial aerospace simulation workflows concludes the paper.

Keywords: scientific machine learning; hybrid solvers; learned discretizations; equivariant neural operators; CFD

1. Introduction

Computational fluid dynamics (CFD) is undergoing a rapid transformation. GPU-native solvers have cut the cost of scale-resolving simulations by one to two orders of magnitude. This newly affordable high-fidelity capability has produced large open datasets for machine learning, first in automotive external aerodynamics [1] and more recently for high-lift aircraft configurations [2]. Neural operators, which learn mappings between function spaces rather than individual solutions (DeepONet [3],

Fourier neural operators [4]), provide the architectural basis for surrogates trained on such data. These surrogates predict surface and volume fields for unseen geometries in seconds with engineering-relevant accuracy [5], and transformer emulators of the AB-UPT line operate on industrial meshes of the order of one hundred million cells [6]. Quantitative estimates of what a foundation model for CFD would cost and require have also appeared [7].

These developments are reshaping expectations across the simulation community. Yet for aerospace applications they face specific headwinds. Certification-grade data are scarce and expensive; flight envelopes must be explored precisely where data are sparsest; and the aeronautical certification culture requires error control, reproducibility and traceability that black-box regression does not naturally provide. The CFD Vision 2030 study [8] already framed the challenge as one of predictive confidence across the envelope, beyond raw throughput.

ONERA's response, presented here, is not an alternative to the foundation-model route but a complement to it: anchor the learning in the physics and in the numerical method itself. Where the governing equations, Green's functions, conservation properties or symmetries are exactly known, they are computed or enforced exactly. Machine learning is confined to the components that classical numerics handles poorly (sub-grid discretization error, closure terms, boundary corrections) and is embedded inside the solver, not wrapped around it. The resulting hybrids are data-frugal, inherit the generalization properties of the underlying numerics, and expose learned components small enough to be analyzed, bounded and eventually certified.

The paper is organized as follows. Section 2 surveys scientific machine learning for aerodynamics at ONERA. Sections 3 to 5 present three research lines in physics-anchored learning: learned discretizations and gradient recovery for finite-volume methods; learned finite volumes whose guarantees hold by construction; and structurally exact hybrid solvers, which learn only the remainder of the solution operator, up to the encoding of geometry for generalization to unseen shapes. Section 6 gathers ongoing developments under review. Section 7 discusses how these approaches articulate with foundation-model efforts, and Section 8 outlines a roadmap.

2. Scientific Machine Learning for Aerodynamics at ONERA

2.1 The Advanced Numerical Simulation domain

ONERA's simulation activities are federated in the Advanced Numerical Simulation (SNA) scientific domain, which spans applied mathematics, numerical methods, high-performance computing, model reduction and uncertainty quantification across scientific departments, and which the author has directed since 2024. The departments develop and maintain the software that supports both research and French/European industrial programmes: the elsA CFD suite for aerodynamics [9], the CEDRE multiphysics platform for energetics and propulsion, structural and materials codes such as A-Set [10], and the next-generation CFD environments SoNICS (developed with Safran) and CODA (developed within the ONERA–DLR–Airbus partnership) [11]. This portfolio matters for SciML in two ways: it provides the industrial-grade numerical cores into which learned components must ultimately be integrated, and it generates the high-quality, provenance-controlled data on which trustworthy learning depends.

This portfolio serves an explicit strategy of hybridization between numerical simulation and artificial intelligence. High-fidelity simulation is essential for safety and performance certification, but its computational cost restricts the exploration of large design spaces, particularly for turbulent flows at high Reynolds numbers; ONERA therefore develops complementary avenues that significantly reduce cost while rigorously preserving physical realism. The ambition is threefold: machine-learning reduced-order models that reproduce complex flows much faster without losing simulation

fidelity; data assimilation that fuses sparse experimental measurements with fast simulations to reconstruct physically consistent flow fields; and hybrid AI-simulation (machine-learning accelerated PDE solvers that combine the mathematical guarantees of classical methods with the approximation power of neural networks). The remainder of this paper develops mainly the third pillar; the first two are surveyed in Section 2.2.

2.2 Machine learning for flow physics

ONERA teams apply machine learning to flow physics itself, a research line increasingly framed by the convergence between physics models and data that digital twins require: data correct and calibrate the models; models complete the data. Hybrid ML reduced-order models combine a truncated POD-Galerkin basis with autoencoder compression of the residual modes and a neural-ODE closure; they correct both truncation and reduced-trajectory errors and improve long-time dynamics by orders of magnitude over POD-Galerkin alone [12]. AI-learned wall laws outperform classical analytical wall models under pressure gradients and cut CPU cost with respect to wall-resolved RANS, with dynamic point-by-point switching between classical closures and learned laws [13]. Deep reinforcement learning has been applied to closed-loop flow control with an algorithm that selects a small set of optimally placed sensors [14]. Turbulence models augmented by machine learning, built by field inversion on massively separated flows [15, 16], anticipated the systematic effort now structured by the NEWMAC ANR project [17]. On the data-assimilation side, sparse experimental measurements are fused with fast simulations to reconstruct physically consistent flow fields: by nudging-based correction of unsteady models [18], by sequential assimilation feeding online low-dimensional models [19], and by mean-flow reconstruction with physics-informed neural networks augmented by a turbulence model [20, 21]. In each case the learning problem is posed inside a physical formulation (a control loop, a closure, a reduced model, an assimilation problem), not as unconstrained regression.

On the data side, ONERA released the CRM WBPN database (468 RANS solutions around NASA Common Research Model wing-body-pylon-nacelle configurations) together with a public regression challenge benchmarking neural and classical surrogates [22]. The effort is deliberately of the same family as the open datasets driving the foundation-model agenda [1, 2], at RANS fidelity and aircraft-relevant geometry: open, well-curated data are a precondition for measurable progress.

2.3 Positioning

Across these activities a consistent positioning emerges, shaped by aerospace constraints: data frugality (certification-grade data are orders of magnitude scarcer than in automotive or consumer domains), physics constraints as a substitute for data volume (conservation, invariances, known operators), integration into existing industrial solvers rather than replacement of them, and a certification perspective from the outset. The research lines below illustrate this positioning across the numerical spectrum, from grid-based finite-volume methods to integral-equation fast solvers.

3. Learning Inside the Discretization: Data-Driven Finite-Volume Components

A first line of research, developed in the doctoral work of G. de Romémont and colleagues, follows the learned-discretization paradigm introduced by Bar-Sinai et al. [23] and extended to accelerated CFD by Kochkov et al. [24]: instead of learning the solution, learn the discretization, that is, the small local operators (interpolants, flux reconstructions, gradient estimates) whose truncation error dominates coarse-mesh computations, while keeping the finite-volume update, and hence conservation, exact. This is a "solver-in-the-loop" approach in the sense of Um et al. [25]: the learned components are trained through the numerical scheme itself, a training mode that work showed to outperform supervised training of the same network on the same data. CNN-based

learned schemes handle structured grids; deep operator networks handle arbitrary unstructured topologies; physically consistent regularization is embedded in training.

3.1 Learned flux discretizations for hyperbolic conservation laws

In [26], compact neural networks are trained to provide data-driven discretizations for hyperbolic conservation laws in finite-volume form. The approach extends learned discretizations to solutions with shocks, coarse meshes and varying boundary conditions, a known weakness of early learned schemes, which tended to be trained and evaluated on periodic problems. The networks act strictly at the stencil level: the scheme remains conservative by construction, and the learned component replaces only the reconstruction step. On the test problems considered, the learned scheme maintains accuracy on meshes substantially coarser than those required by the baseline discretization, trading offline training cost for online resolution.

3.2 Adaptive gradient recovery on unstructured meshes

The follow-up work [27] targets a component closer to industrial practice: gradient reconstruction on unstructured meshes, a key accuracy bottleneck of second-order finite-volume solvers of the elsA/SoNICS class. A neural model performs data-driven adaptive gradient recovery for the compressible Euler equations on general unstructured meshes, and delivers accuracy improvements of 20–60% over standard reconstruction operators on the configurations tested, at a cost compatible with GPU batching of the local operations. Because the learned operator is local, independent of the mesh topology and fed with non-dimensionalized stencil data, it generalizes across meshes and flow conditions in a way that global surrogates cannot.

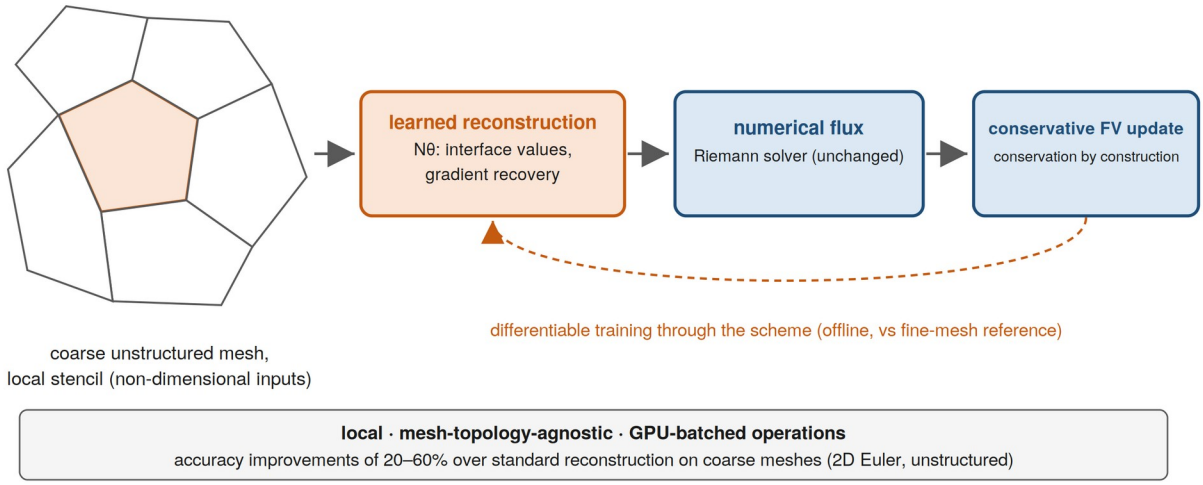


Figure 1 – Learning inside the finite-volume method: a local, mesh-agnostic learned reconstruction (orange) replaces the step where truncation error dominates, while flux evaluation and the conservative update (blue) are unchanged: conservation holds by construction.

4. Guarantees by Construction for Learned Finite Volumes

The second line closes the gap between the coarse-mesh promise of Section 3 and what an engineer can deploy: a learned scheme that cannot produce negative densities and cannot crash on a boundary condition it never saw. A preprint applies this programme to steady supersonic channel flow over obstacles, including the forward-facing step of Woodward and Colella [28]. The network sets the two local quantities every second-order finite-volume scheme rests on, the reconstructed

gradient and the limiter, but only inside sets built so that no inadmissible output exists: gradient weights bounded away from cancelling a neighbour, so that the least-squares system stays invertible, and a limiter capped by a value computed from the local flow, below which every choice keeps density and pressure positive and creates no new extremum. The rest of the solver is fixed and carries its own guarantees: an entropy-stable flux, a wall condition that makes the boundary entropy flux vanish by symmetry, and a time step from a closed-form positivity bound. Admissibility therefore holds for every value of the weights rather than as an outcome of training, and no negative density or pressure occurred in any computation reported.

Because the scheme is safe whatever the network does, the study can ask what the network actually contributes, and the answer is sharp. Of the two quantities the network sets, only one carries the gain: the learned gradient weights alone reproduce the full effect of the network to within a tenth of a percent, while the learned limiter alone moves the error ten times less. What learning adds to this solver is a better gradient, not a better limiter, so the safety imposed does not throttle what pays. Against its own guaranteed backbone, with cases and evaluation rules fixed before measuring, the method lowers the error by 38% on an unseen geometry and by 29% on an unseen obstacle topology; that backbone itself outperforms a classical smooth limiter on three cases of four. The limits are reported with the same candour: the learned weights are tuned to the regimes they were fitted on and do not carry to an unseen Mach number.

The trade is stated in the terms that matter industrially, accuracy per second. On the reference case, refining the mesh once improves the error fourfold and takes the run from 52 to 416 seconds; learning secures half of that improvement for 69 seconds, a sixth of the price, which is three times the accuracy per second spent. These results sharpen the thesis of this paper: the construction constrains the limiter, the gain comes from the gradient, and the guarantee costs nothing where learning pays. Extension towards three-dimensional configurations, the turbulent regime and the SoNICS/CODA solver generation is the next step.

5. Structurally Exact Hybrid Solvers: Learning Only the Remainder

The third line keeps the exactly known part of the solution operator exact and learns only the remainder. The template was first established in a controlled viscous setting with exactly known Green’s functions [29]: a fast core with proven error bounds computes the free-space physics exactly, while a compact learned operator, rotation-equivariant by construction and formulated on a second-kind boundary equation, supplies the only part that eludes closed-form treatment. The hybrid matches fully resolved boundary-integral solves at a fraction of the online cost, and its data efficiency exceeds that of end-to-end surrogates for a structural reason: learning is confined to the boundary-interaction physics, and far-field behaviour and conservation are inherited from the analytical core.

A new line of work transfers this template to the setting of ML-assisted design: incompressible flow around bodies the model has never seen [30]. One differentiable solver runs forward. Incompressibility is computed exactly, through a gauge-fixed vector potential, so the predicted velocity is divergence-free to machine precision for any value of the learned parameters; transport and the non-local physics are computed by the solver, never learned; learning is confined to a small, local closure (Figure 2). Emulation surrogates offer no such guarantee, whatever their accuracy.

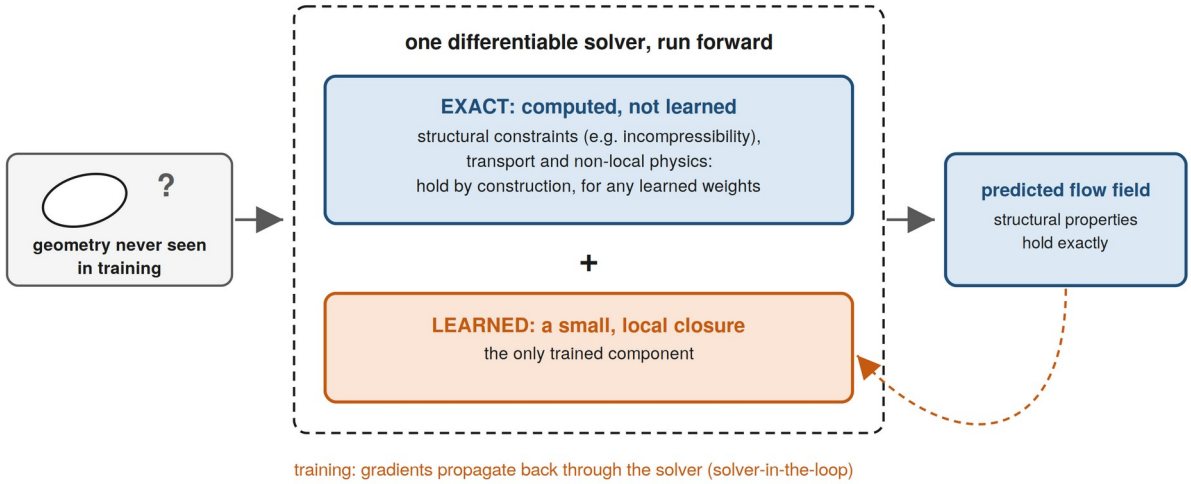


Figure 2 – The hybrid AI-simulation principle. One differentiable solver computes the exactly known physics by construction (blue) and adds a small, local learned closure, the only trained component (orange). Training propagates gradients back through the solver; structural properties hold for any value of the learned parameters, including on geometries never seen in training.

With the structure guaranteed, the study isolates the question intuition gets wrong: how should the geometry be encoded for a single trained model, used unchanged, to generalize to unseen shapes? In controlled, capacity-matched comparisons over five seeds, a signed-distance field read at the prediction point beats a global point-cloud baseline out of distribution; a hybrid encoder, which adds a global pose and shape code to that local reading, generalizes better still (0.27 versus 0.37 relative L2 error); and no tested encoder survives a change of topology. Negative results are reported alongside. The lesson matches the thesis of this paper: the encoding of the geometry, not model capacity, is the lever of out-of-distribution generalization.

6. Ongoing Developments

Several developments extend the published results above. As they are the object of publications currently under review or in preparation, only their direction is outlined here; methods and quantitative results are deferred to the dedicated papers.

A companion study, currently under review, addresses the resource question the hybrid route raises: at equal end-to-end cost, accounting for data generation, training and inference, when is a fixed simulation budget best spent on a finer grid for the classical solver, on a learned closure embedded inside it, or on a neural emulator in its place? First executed results support the physics-anchored thesis of this paper: structural exactness combined with small learned corrections generalizes across geometries in regimes where equal-budget emulation degrades.

The same solver-in-the-loop philosophy is applied to turbulence closures through an end-to-end differentiable pipeline built on the automatic differentiation of the CFD solver, with first results on a transonic turbine cascade (VKI LS-59) over a wide range of Mach numbers and incidences [31]. The hybrid template of Section 5 is likewise being extended towards learned multilevel fast-multipole pipelines and other kernels of aerospace interest, including wave-type problems.

7. Physics-Anchored Learning and Foundation Models: Complementary Routes

How do these hybrid approaches articulate with the foundation-model trajectory? The analysis of

Ashton, Brandstetter and Mishra [7] derives the first scaling laws coupling CFD data-generation cost with model-training cost. It concludes that large volumes of high-fidelity transient data are the credible route to a CFD foundation model, with data generation, not training, as the dominant cost. We read this analysis as strengthening the case for physics anchoring: whatever reduces the data requirement of a given accuracy target attacks the dominant cost term directly. Enforced symmetries, exactly computed operator components and formulations that preserve conservation are such data multipliers, as the equivariant boundary-correction results illustrate [29].

The two routes also serve different points of the engineering workflow, and each has a distinct research lineage. At the emulation end, neural-operator surrogates (DoMINO among them [5]) excel at interactive design-space exploration in regimes covered by data, where an evaluation in seconds matters more than a certified error bound. The UPT/AB-UPT transformer line has pushed emulation to industrial mesh sizes, with recent extensions from automotive to aerospace configurations [6, 32]. At the augmentation end, the differentiable physics school of Thuerey and co-workers established the training methodology that ONERA's solver-in-the-loop work builds on [25]. Physics-anchored hybrids target the complementary regime: extrapolation beyond the data envelope; analysis oriented towards certification, where the learned component must be small enough to be bounded and audited; and acceleration of the high-fidelity solvers that generate reference data in the first place, a loop in which better hybrid solvers cut the cost of the datasets foundation models need. Differentiable programming is the natural meeting point: embedding learned components in solvers and training through them requires the same software substrate as fine-tuning large pre-trained models against solver feedback. ONERA has drawn the software conclusion. Beyond HPC and GPU readiness, automatic differentiation is treated as a core architectural design constraint of the next-generation solver SoNICS [11, 33], so that gradient information flows natively between numerical and AI models, the precondition for actual AI–simulation hybridization rather than loose coupling.

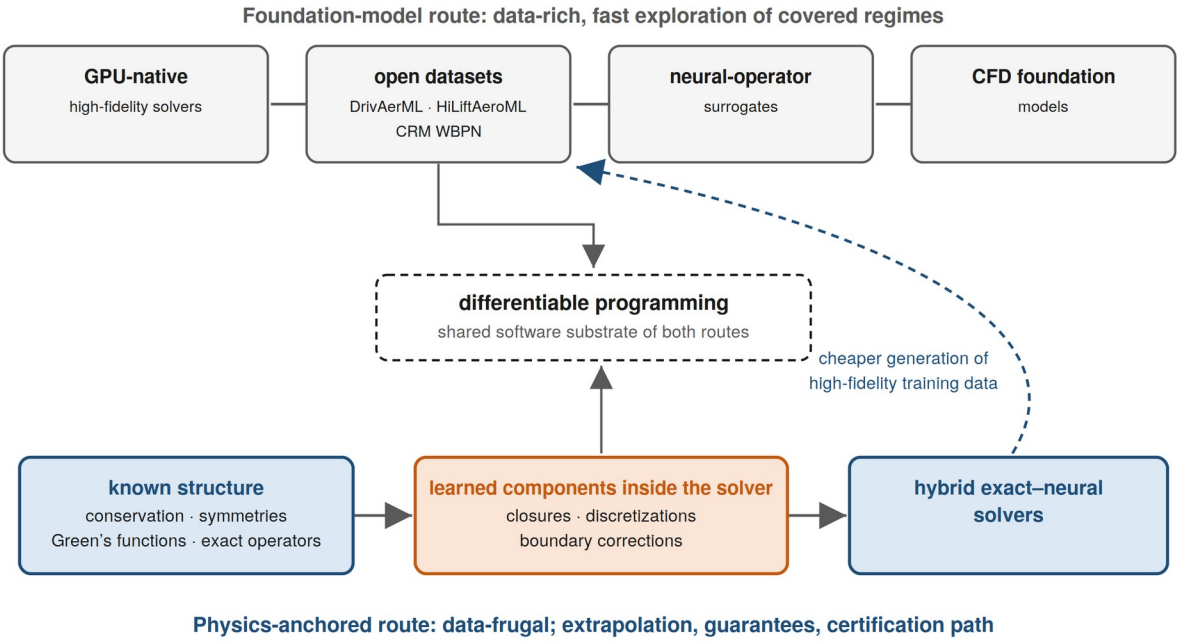


Figure 3 – Two complementary routes to AI-augmented simulation. The foundation-model route (top) is data-rich; the physics-anchored route (bottom) is data-frugal and oriented towards certification. Differentiable programming is their common substrate, and hybrid solvers feed back into the data pipeline by cutting the cost of high-fidelity data generation.

Finally, the institutional dimension matters. For a national aerospace laboratory, the ability to verify, requalify and eventually certify simulation capabilities, including their learned components, is a sovereignty issue as much as a scientific one. ONERA's contribution to trustworthy-AI frameworks [34] and its co-development of industrial solvers position it to help define what "certifiable SciML" should mean for aeronautics, in cooperation with the platform-scale efforts of industry rather than in competition with them.

8. Roadmap and Concluding Remarks

Four priorities structure the coming years: generalization of the hybrid exact–neural template across integral and grid-based methods, with structural exactness, equivariance and conservation as systematic design constraints; a shared software platform bridging ONERA's differentiable solvers with machine-learning frameworks and algorithm libraries, extended towards certifiable uncertainty quantification and digital-twin demonstrations; industrialization of learned components into the SoNICS and CODA ecosystem, with non-regression testing and error monitoring adapted to learned operators; and open evaluation, extending ONERA's database efforts so that physics-anchored and foundation-model approaches can be compared on common aerospace benchmarks [1, 2, 22].

The aerospace community will not choose between physics and learning; the question is where each is placed. ONERA's experience suggests that placing exactness at the core and learning at the margins, rather than the reverse, is a route that scarce data, hard guarantees and industrial codes can all accommodate. Our ambition for the coming years is to make that route industrial: hybrid solvers whose exact core, learned components and software substrate are engineered together, so that the resulting simulation capability can be verified, requalified and ultimately certified as a whole.

9. Contact Author Email Address

mailto: denis.gueyffier@onera.fr

10. Copyright Statement

The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The authors also confirm that they have obtained permission, from the copyright holder of any third party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.

11. References

[1] Ashton N, Mockett C, Fuchs M, Fliessbach L, Hetmann H, Knacke T, Schonwald N, Skaperdas V, Fotiadis G, Walle A, Hupertz B and Maddix D. DrivAerML: high-fidelity computational fluid dynamics dataset for road-car external aerodynamics. arXiv:2408.11969, 2024.

[2] Ashton N, Clark A, Heidt L, Ivey C, Bose S, Agrawal R, Goc K, Ranade R, Adams C, Sharpe P, Nidhan S, Akkurt S, Leibovici D and Kossaifi J. HiLiftAeroML: high-fidelity computational fluid dynamics dataset for high-lift aircraft aerodynamics. arXiv:2605.19565, 2026.

[3] Lu L, Jin P, Pang G, Zhang Z and Karniadakis G E. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. Nature Machine Intelligence, Vol. 3, pp 218–229, 2021.

[4] Li Z, Kovachki N, Azizzadenesheli K, Liu B, Bhattacharya K, Stuart A and Anandkumar A. Fourier neural operator for parametric partial differential equations. International Conference on Learning Representations, 2021.

[5] Ranade R, Nabian M A, Tangsali K, Kamenev A, Hennigh O, Cherukuri R and Choudhry S. DoMINO: a decomposable multi-scale iterative neural operator for modeling large scale engineering simulations. arXiv:2501.13350, 2025.

[6] Alkin B et al. AB-UPT: scaling neural CFD surrogates for high-fidelity automotive aerodynamics simulations via anchored-branched universal physics transformers. Transactions on Machine Learning Research, 2025 (arXiv:2502.09692).

[7] Ashton N, Brandstetter J and Mishra S. Fluid intelligence: a forward look on AI foundation models in CFD. arXiv:2511.20455, 2025.

[8] Slotnick J, Khodadoust A, Alonso J, Darmofal D, Gropp W, Lurie E and Mavriplis D. CFD Vision 2030 study: a path to revolutionary computational aerosciences. NASA/CR-2014-218178, 2014.

[9] Cambier L, Heib S and Plot S. The Onera elsA CFD software: input from research and feedback from industry. Mechanics & Industry, Vol. 14, No. 3, pp 159–174, 2013.

[10] Besson J, Bovet C, Chiaruttini V, Cortial J, Garaud J-D, Geoffre A, Kehr-Candille V, Marchand B, Medghoul G, Messahel R, Rannou J, Rogalski N, Sanchez R and Scherer J-M. A-set: outils et méthodes de simulation de nouvelle génération pour les matériaux et structures. 17e Colloque National en Calcul des Structures (CSMA), 2026.

[11] Gueyffier D, Plot S and Soismier M. SoNICS: a new generation CFD software for satisfying industrial users needs. NATO STO Workshop AVT-366, 2022.

[12] Lepage N, Beneddine S, Fiorini C, Mortazavi I, Sipp D and Thome N. Hybrid autoencoder/Galerkin approach for nonlinear reduced order modelling. Computers & Fluids, Vol. 302, 106811, 2025.

[13] Romanelli M, Beneddine S, Mary I, Beaugendre H, Bergmann M and Sipp D. Efficient and accurate data-driven wall modelling strategy for Reynolds-averaged Navier–Stokes simulations. Journal of Computational Physics, Vol. 538, 114128, 2025.

[14] Paris R, Beneddine S and Dandois J. Robust flow control and optimal sensor placement using deep reinforcement learning. Journal of Fluid Mechanics, Vol. 913, A25, 2021.

[15] Duraisamy K, Iaccarino G and Xiao H. Turbulence modeling in the age of data. Annual Review of Fluid Mechanics, Vol. 51, pp 357–377, 2019.

[16] Volpiani P S, Meyer M, Franceschini L, Dandois J, Renac F, Martin E, Marquet O and Sipp D. Machine learning-augmented turbulence modeling for RANS simulations of massively separated flows. Physical Review Fluids, Vol. 6, 064607, 2021.

[17] NEWMAC – New generation of machine-learning based models for aerodynamics computations. ANR-22-FAI2-0002, 2023–2027, coordinator P. S. Volpiani (ONERA).

[18] Zauner M, Mons V, Marquet O and Leclaire B. Nudging-based data assimilation of the turbulent flow around a square cylinder. Journal of Fluid Mechanics, Vol. 937, A38, 2022.

[19] Villiers R, Mons V, Sipp D, Lamballais E and Meldi M. Enhancing unsteady Reynolds-averaged Navier–Stokes modelling from sparse data through sequential data assimilation and machine learning. Flow, Turbulence and Combustion, Vol. 115, pp 991–1029, 2025.

[20] Raissi M, Perdikaris P and Karniadakis G E. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations.

Journal of Computational Physics, Vol. 378, pp 686–707, 2019.

[21] Patel Y, Mons V, Marquet O and Rigas G. Turbulence model augmented physics-informed neural networks for mean-flow reconstruction. Physical Review Fluids, Vol. 9, 034605, 2024.

[22] Peter J, Bennehard Q, Heib S, Hantrais-Gervois J-L and Moëns F. ONERA's CRM WBPN database for machine learning activities, related regression challenge and first results. Computers & Fluids, Vol. 302, 106838, 2025.

[23] Bar-Sinai Y, Hoyer S, Hickey J and Brenner M P. Learning data-driven discretizations for partial differential equations. Proceedings of the National Academy of Sciences, Vol. 116, No. 31, pp 15344–15349, 2019.

[24] Kochkov D, Smith J A, Alieva A, Wang Q, Brenner M P and Hoyer S. Machine learning–accelerated computational fluid dynamics. Proceedings of the National Academy of Sciences, Vol. 118, No. 21, e2101784118, 2021.

[25] Um K, Brand R, Fei Y, Holl P and Thuerey N. Solver-in-the-loop: learning from differentiable physics to interact with iterative PDE-solvers. Advances in Neural Information Processing Systems, Vol. 33, pp 6111–6122, 2020.

[26] de Romémont G, Renac F, Nunez J, Gueyffier D and Chinesta F. A data-driven learned discretization approach in finite volume schemes for hyperbolic conservation laws and varying boundary conditions. Computers & Fluids, Vol. 307, 106978, 2026.

[27] de Romémont G, Renac F, Chinesta F, Nunez J and Gueyffier D. Data-driven adaptive gradient recovery for unstructured finite volume computations. arXiv:2507.16571, 2025.

[28] Gueyffier D. Guarantees by construction for learned finite volume schemes on steady supersonic flow. arXiv:2607.20171, 2026.

[29] Gueyffier D. Solver exactness, learned flexibility: equivariant boundary-correction operators for Stokes flow. arXiv:2606.25075, 2026.

[30] Gueyffier D. Geometry encoding for out-of-distribution generalization in incompressible flow. In preparation.

[31] Saverio L et al. An end-to-end PyTorch interface for differentiable PDE solvers: a RANS model-correction study. Submitted to Data-Centric Engineering, 2026.

[32] Alkin B et al. AB-UPT for automotive and aerospace applications. arXiv:2510.15808, 2025.

[33] Lienhardt M and Michel B. SoNICS: design and workflow of a new CFD software. Comptes Rendus Mécanique, Vol. 353, pp 1261–1287, 2025.

[34] Delseny H et al. White paper: machine learning in certified systems. IRT Saint Exupéry / DEEL, hal-03176080, arXiv:2103.10529, 2021.