Toward Machine-learned Discretizations and Differentiable CFD for Compressible Two-phase Flows

Application for the 20th ERCOFTAC da Vinci Competition 2025

Candidate:
Dr.-Ing. Deniz A. Bezgin
Technical University of Munich
Chair of Aerodynamics and Fluid Mechanics
deniz.bezgin@tum.de / +49 89 289 16742
Google Scholar / ResearchGate / ORCID

Supervisor:
Prof. Dr.-Ing. Nikolaus A. Adams
Technical University of Munich
Chair of Aerodynamics and Fluid Mechanics

Brief summary

Compressible single- and two-phase flows are central to many applications in aerospace, energy, process engineering, and biomedicine. However, numerically simulating these flows poses significant challenges due to complex flow phenomena such as turbulence, compressibility, combustion, and material interfaces. Since scale-resolving simulations of these processes will remain beyond the reach of available computational resources in the foreseeable future, computational fluid dynamics (CFD) continues to rely on suitable closure models (e.g., for turbulence) and nonlinear, solution-adaptive discretizations (e.g., for shock-capturing). While machine learning (ML) offers novel avenues to accelerate conventional numerical models and discretizations, questions arise on how to train ML models to ensure consistency, stability, and generalizability. This research addresses these points by examining constrained hybrid ML-CFD models and by developing a comprehensive differentiable CFD solver called JAX-Fluids. This solver facilitates the integrated training of these models, thereby preserving known physical laws and key numerical properties such as stability and convergence. JAX-Fluids combines high-order numerical schemes for compressible single- and two-phase flows with state-of-the-art ML methods within a unified solver framework, facilitating research at the intersection of ML and CFD. JAX-Fluids' automatic differentiation (AD) capabilities not only enable end-to-end training of machine-learned discretizations and closure models but also open up new solution strategies for challenging inverse problems, including active flow control or shape optimization in aerodynamics.

Scientific summary

Numerical simulations of shock-dominated flows require nonlinear, solution-adaptive discretizations, so-called shock-capturing schemes, which introduce sufficient – but not excessive – numerical dissipation to adequately represent discontinuities on the computational mesh. Among the most popular shock-capturing schemes are weighted essentially non-oscillatory (WENO) schemes, which use analytical smoothness measures to form convex combinations of candidate Harten-type polynomials. While analytical smoothness measures recover the underlying linear scheme in the asymptotic limit, they suffer from loss of accuracy in the pre-asymptotic regime. To address this, we substitute the analytical smoothness measures of a third-order WENO scheme by a shallow neural network (NN) surrogate [4], see Fig. A. The input stencil values are preprocessed via a so-called Delta-layer to ensure Galilean invariance. The network predicts the polynomial weights ω_k^{NN}, which are postprocessed to ensure the ENO-property (ENO-layer). We propose a novel loss function that dynamically balances two objectives – minimizing reconstruction error and minimizing the deviation of the polynomial weights ω_k^{NN} from the ideal linear weights – based on the smoothness of each training sample. The resulting WENO3-NN scheme is trained offline on analytical functions, including polynomials, trigonometric functions, and discontinuities. The overall scheme achieves maximum-order convergence in the asymptotic limit while outperforming conventional third-order WENO schemes at coarse resolutions. The combination of elements of conventional numerical methods with constrained ML surrogates enhances model interpretability and provides fallback mechanisms for unseen scenarios.

While the WENO3-NN model is trained offline and remains physics-agnostic, we resort to online training to develop a physics-informed, data-driven discretization for nonclassical undercompressive shocks [5]. By online training, we refer to a training paradigm in which the trainable model is integrated into the CFD solver, and AD is used to propagate gradients through the entire CFD framework. Using the example of a scalar hyperbolic partial differential equation (PDE) with a cubic flux, we demonstrate that end-to-end optimization allows shaping the truncation error of the learned data-driven discretization such that it is consistent with the underlying small-scale terms of the governing equations. Unlike offline training, the ML model in this case observers the entire dynamics of the discretized PDE. As a result, the model becomes physics-informed and generally extrapolates better to out-of-distribution samples while exhibiting enhanced long-term stability. In our experiments, the NN-based discretization was trained solely on Riemann problems. Notably, when tested on a smooth sinusoidal initial condition, the NN-based discretization was not only able to capture the dynamic formation of shock discontinuities – despite not having seen such scenarios during training – but also demonstrated significantly improved dissipation-dispersion characteristics compared to a well-controlled dissipation (WCD) scheme, see Fig. B.

Building upon aforementioned work on hybrid ML-CFD methods, the central contribution of this research is the development of JAX-Fluids (https://github.com/tumaer/JAXFLUIDS), a fully differentiable CFD solver for compressible single- and two-phase flows [2,3]. JAX-Fluids is a high-order Godunov-type finite-volume solver and supports two-phase flow simulations via two complementary approaches: a level-set-based sharp interface method and a five-equation diffuse interface method. Notably, the level-set method also serves as a conservative immersed boundary method, enabling flexible and accurate treatment of complex geometries.

Most importantly, JAX-Fluids enables the seamless integration of ML modules and leverages AD to compute gradients of the discretized PDE throughout the entire solver. The AD capabilities are essential for end-to-end optimization of ML-accelerated numerical models and discretizations. Fig. C provides a schematic of the integrated training process: given a ground truth trajectory, a forward trajectory is generated using an ML model embedded within the JAX-Fluids solver. A user-defined loss function $\mathcal{L}(\theta)$ measures the discrepancy between the ground truth and the model trajectory, where θ are trainable parameters of the ML model. Since the entire integration loop is automatically differentiable, the gradient of the loss function with respect to the parameters, $\nabla_\theta \mathcal{L}(\theta)$, can be computed directly, and model parameters can be optimized using gradient-based algorithms.

We leveraged JAX-Fluids' end-to-end optimization capabilities to implement and train machine-learned implicit large eddy simulation (ML-ILES) discretizations for compressible turbulence [1]. Using coarse-grained direct numerical simulation data as ground truth, the ML model was embedded into the cell-face reconstruction process, thereby shaping the schemes' truncation error. A key result of this work is that end-to-end trained discretizations can learn tailored numerical strategies for flow regions with strong compressibility effects and regions of high turbulent intensity. This demonstrates the potential of data-driven discretizations to handle complex flow phenomena.

While Fig. C shows 3D flow fields as training targets, the loss function can be computed from any (potentially low-dimensional) observable of the spatio-temporal trajectory. The training workflow has already been successfully adapted to solve classical inverse problems in CFD, including data assimilation, active flow control problems, and aerodynamic shape optimization. For example, we optimized the model parameters θ of an airfoil geometry by minimizing its drag coefficient. The differentiable CFD framework removes the need for manual derivation and implementation of adjoint operators and opens novel avenues for tackling longstanding inverse problems in fluid mechanics.

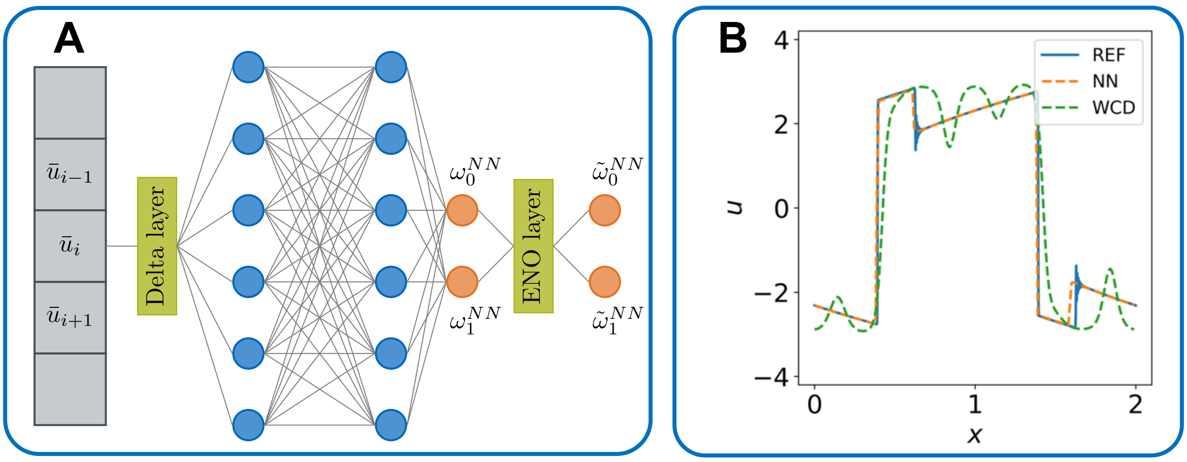

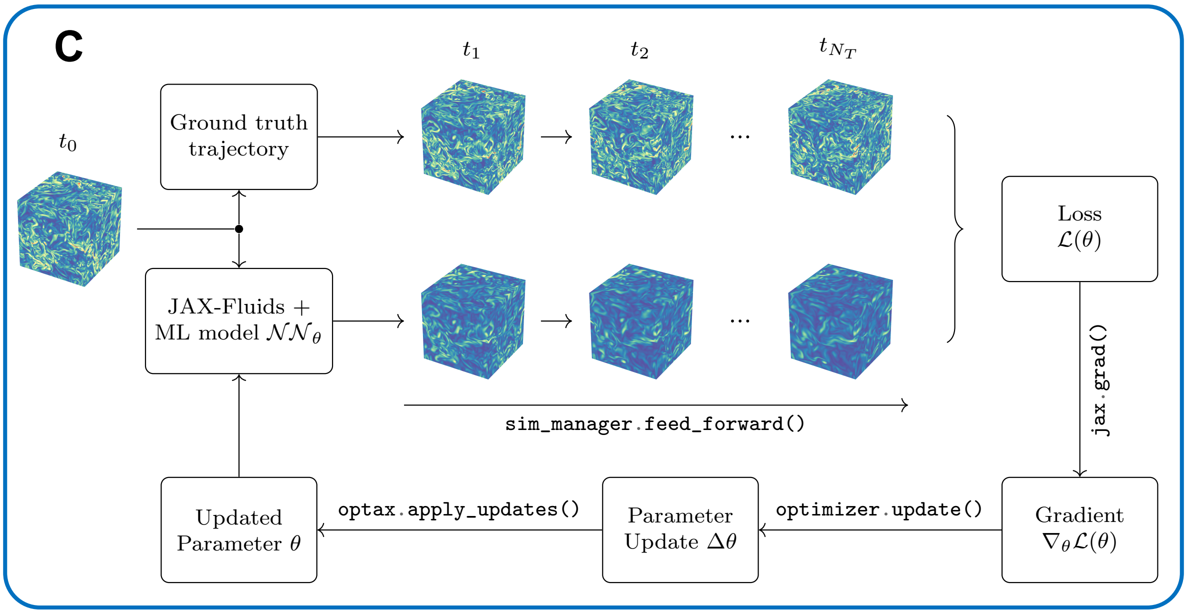

JAX-Fluids is parallelized and runs efficiently on modern high-performance computing hardware, including GPUs and TPUs, enabling scalable training and inference for large-scale flow simulations. By unifying high-fidelity numerical methods with differentiable programming, JAX-Fluids establishes a versatile open-source platform for advancing research at the intersection of ML and CFD. It not only enables the training of ML-accelerated models and discretizations in a physically consistent manner but also provides a powerful tool for solving inverse problems – bridging data-driven modeling with first-principles physics.

Selected publications & indication of the total research output

[1] **Bezgin, D. A.**, Buhendwa, A. B., Schmidt, S. J., & Adams, N. A. (2025). ML-ILES: End-to-end optimization of data-driven high-order Godunov-type finite-volume schemes for compressible homogeneous isotropic turbulence. *Journal of Computational Physics*, *522*, 113560.

[2] **Bezgin, D. A.**, Buhendwa, A. B., & Adams, N. A. (2025). JAX-Fluids 2.0: towards HPC for differentiable CFD of compressible two-phase flows. *Computer Physics Communications*, *308*, 109433.

[3] **Bezgin, D. A.**, Buhendwa, A. B., & Adams, N. A. (2023). JAX-Fluids: A fully-differentiable high-order computational fluid dynamics solver for compressible two-phase flows. *Computer Physics Communications*, *282*, 108527.

[4] **Bezgin, D. A.**, Schmidt, S. J., & Adams, N. A. (2022). WENO3-NN: A maximum-order three-point data-driven weighted essentially non-oscillatory scheme. *Journal of Computational Physics*, *452*, 110920.

[5] **Bezgin, D. A.**, Schmidt, S. J., & Adams, N. A. (2021). A data-driven physics-informed finite-volume scheme for nonclassical undercompressive shocks. *Journal of Computational Physics*, *437*, 110324.

In addition to the publications described above, the total research output includes a paper published in the *Journal of Computational Physics (JCP)*, an accepted paper at *Physical Review Fluids*, and a preprint currently under review at *JCP*. A summary of all publications is available on Google Scholar. This research has been continuously presented at national and international conferences, including ECCOMAS, APS DFD, and GAMM, and at several invited talks, including the *Harvard Rising Stars in CSE* event and the *Computer Physics Communications Seminar Series*. Under the supervision of the candidate, more than 10 students have employed the JAX-Fluids solver for their term papers and bachelor's and master's theses. Multiple collaborations with both national and international researchers have been established over the course of this doctorate, and the JAX-Fluids code base remains under active development.