Proceedings of the 1st international Symposium on AI and Fluid Mechanics
Abstract No S18 A1

PICT: Adaptive GPU Accelerated and Differentiable Fluid Simulation for Machine Learning

Erik Franz[1] ; Hao Wei[1] ; Luca Guastoni[1] ; Nils Thuerey[1]

[1] *School of Computation, Information and Technology; Technical University of Munich (TUM)*

Abstract

The paper introduces the open-source differentiable fluid simulation framework *PICT* tailored for machine learning applications. Developed within the PyTorch ecosystem, it leverages GPU acceleration through custom CUDA operations and provides users with a powerful tool for fluid dynamics research and optimization tasks.

At its core, PICT implements the *PISO* algorithm (Pressure-Implicit with Splitting of Operators) for simulating incompressible flows. This algorithm is integrated into PyTorch with a hybrid design: CUDA handles the computationally intensive components, while Python ensures ease of use and customizability. The simulator features spatially adaptive *multi-block grids* using a generalized coordinate system, enabling flexible grid refinement and alignment with complex boundaries. Despite this adaptability, it retains a Cartesian grid structure, simplifying memory management and maintaining compatibility with convolutional neural networks (CNNs) for machine learning applications.

A primary focus of PICT is its differentiability. Efficient analytical gradients for key components, such as the implicit Euler step, have been implemented to support PyTorch's automatic differentiation. This makes it suitable for end-to-end training and gradient-based optimization tasks. The framework enables stable backpropagation through long simulations, allowing users to optimize physical parameters (e.g., viscosity or boundary velocity) and learn control strategies using neural networks.

To validate its performance, the simulator's forward simulations were compared against analytical and numerical benchmarks, demonstrating accuracy and stability even during extended turbulent rollouts. Numerical gradient checks were performed for individual operations, and optimization experiments confirm the reliability of its gradients in challenging scenarios. We will present these validation experiments and gradient calculations together with a first set of optimization tasks and deep learning examples. Together, these experiments demonstrate the capabilities for flexible learning and inverse problem solving tasks in the field of computational fluid dynamics.

PICT addresses a critical gap in the intersection of machine learning and computational fluid dynamics (CFD). By providing a differentiable, GPU-accelerated simulation framework, it opens new possibilities for applications such as hybrid solvers, learned control tasks, and design optimization. Its open-source availability and focus on performance ensures that it can serve as a valuable tool for researchers seeking to integrate machine learning techniques into their simulation workflows.

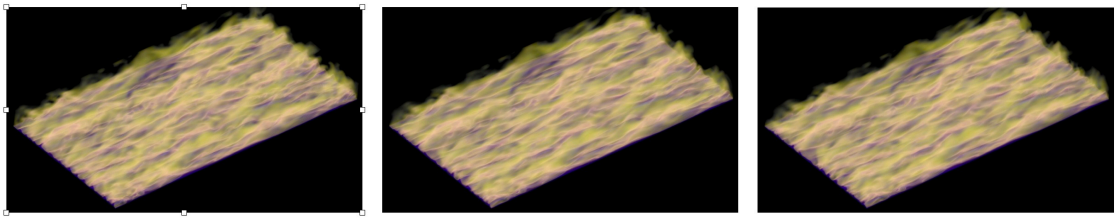

Figure 1: A three-dimensional PICT simulation of a turbulent channel flow. The boundary layer, adaptively resolved near the bottom wall of the channel, is shown in terms of streamwise velocity magnitude.

Keywords: deep learning, differentiable programming, computational fluid dynamics