CFDML: First International Workshop on the Application of Machine Learning Techniques to Computational Fluid Dynamics Simulations and Analysis

Workshop scope

We propose to organize a workshop focused on the emerging field of application of machine learning (ML) techniques to computational fluid dynamics (CFD) based simulations. The rapid development of computing technology in the last few years resulted in an explosive adoption of ML that has also fueled a number of extraordinary scientific, algorithmic, software, and hardware developments such as image and voice recognition, machine vision for self-driving cars, robotics, natural language processing, and much more. Many engineering problems are solved in industry and academia by applying CFD-based simulations. The combination of CFD and ML is a newly emerging research area with the potential to become the key to solve so far unsolved problems in many application domains such as turbulence, combustion, hidden feature detection, provision of expert systems for simulative setups, to name a few. CFD+ML is an emerging field with numerous new efforts and promising results, yet the early work still has a long way to go in demonstrating its full potential. The proposed workshop will stimulate this research by providing a venue to exchange new ideas and discuss challenges and opportunities as well as expose this newly emerging field to a broader research community.

ML is increasingly used to replace portions (sub-models) of traditional simulations in many fields that are based on partial differential equations (PDEs). Apart from fluid dynamic simulations, other research areas, such as constitutive modeling of heterogeneous materials, computational chemistry, computational astrophysics, dynamics of the atmospheric/ocean/climate system, and combustion/chemical reactions are working on similar techniques. Therefore, the proposed workshop will have a broad appeal and will attract a wide community of researchers working in these domains.

Specifically in CFD application, the CFD+ML field primarily consists of several distinct directions: 1) Physics-based modeling with the main focus on fluid physics, such as reduced modeling for dimensionality reduction and the Reynolds-averaged Navier-Stokes (RANS) turbulence modeling; 2) Shape and topology optimization; 3) Uncertainty quantification and reliability analysis; and 4) more recently, reinforcement learning for the design of active/passive flow control.

In the field of turbulence modeling for engineering systems, finding the optimal mathematical/physical model that allows for realistic simulations is a challenging task. It has been identified as a critical area in, e.g., NASA's vision 2030 Computational Fluid Dynamics report. Higher fidelity methods like Direct Numerical Simulation (DNS) or Large Eddy Simulation (LES) are restricted by the need to resolve the domain by meshes that are small enough to capture all the relevant scales of turbulence and hence highly increase the computational costs. Turbulence simulation is typically done using RANS approach since it is more computationally affordable. Such models have been developed over the years based on theoretical, experimental, and empirical approaches. Their calibration is based on a certain set of applications and hence they cannot perform as accurately when applied across a wide range of fluid flows, i.e., individual model constants need to be tuned per flow type to achieve realistic results. This approach is also shared in many other application domains, such as cosmology, where the governing equation system cannot be analytically closed, and theoretical ansatz or empirical data are used to supply this information; the tools discussed in this workshop are of high interest to practitioners in those communities too. ML-based methods for improving turbulence modeling in RANS simulations have recently gained significant interest. The main concepts behind ML-based turbulence modeling is that instead of looking for the complete closure form, a relationship between the unclosed term with other known features and using neural network to fit unknown coefficients could presumably be postulated. The integration of ML with turbulence models is generally accomplished via three distinct strategies: (1) modeling the Reynolds stress anisotropic tensor directly from either input features guided by physics intuition, or Galilean invariant input features (e.g. mean strain rate and rotation rate tensor); (2) modeling the deficiencies in the functional form of the production, destruction, and diffusion term in the turbulence model transport equation; (3) modeling the discrepancy of the Reynolds stress tensor between RANS and high-fidelity data. All three approaches discussed above are essentially targeting the same objective: the improvement of RANS modeling by using standard turbulence models in conjunction with existing high-fidelity data.

Apart from turbulence modeling, ML has been heavily explored recently to tackle the notoriously difficult to solve inverse problem in CFD application: shape optimization and flow control. Solving inverse problem using conventional numerical methods has been prohibitively expensive as it requires multiple runs of the computational solvers with grid regeneration, making most conventional approach almost impractical. In shape optimization, like airfoil or other fluid-structure interactions, researchers are able to build approximate models to directly map the nonlinear relationship between structure geometries with flow fields. Order of magnitude computational times can be saved by replacing the repetitive CFD simulations with purely data-driven models like convolutional neural networks in the initial structure design stage. In flow control, to suppress the onset of transition to turbulence for complex geometry, radiated noise, maneuverability, and vibration reduction, only a fairly limited range of flow configurations are usable due to that turbulent flow exhibits both a broad spectrum of spatial scales and a very rich temporal dynamics. These high-frequency phenomena therefore require sufficiently fast control that are able to adapt to changes in real-time. Purely data-driven approach like Deep Reinforcement Learning (DRL), rather than physical models coupled with traditional closed-loop control, are more suited to these fast response problem. DRL is an optimization framework for problem solving that implies goal-directed interactions of an agent with its environment. In flow control, the agent is the actuator (e.g. plasma or mechanical actuators) while the environment is the flow field from either experimental probes or CFD simulation.

Besides of its computational merits, ML can find the unrevealed physics involved in the complex fluid dynamics with a deeper insight with the abundance of high-fidelity data. Considering furthermore that LES and DNS computations of turbulent flow are becoming more and more complex due to the massive increase of computational power, the data associated with such simulations is even more complex to analyze. In some applications it is already almost impossible to analyze the simulated data in a post-simulative manner, i.e., in-situ analyses become an essential part of the analysis workflow. The massive amount of data produced, however, prevents a full analysis of the flow field and the analysis is often restricted to certain flow features or regions. ML techniques have the potential to support the identification and extraction of not-easy-to-find hidden features in large-scale flow computations, hence allowing to shift the focus from time-consuming feature detection to in-depth examinations of such features. Furthermore, ML techniques have the ability to find undetected correlations between phenomena in the flow, which will lead to deeper insight of the physics involved in complex natural processes.

The development of exascale-ready multi-physics simulation software is key to the success of applied research activities. The trends in hard- and software evolution evidence the unavoidable convergence of both development strains towards an integrated and inter-operable simulation unit with flexible inter-connected hierarchies, data-exchange capabilities, and performance dependencies. That is, future hard- and software cannot be seen as individual components of a simulation anymore, but rather need to interact hand in hand to converge to highly efficient simulation frameworks and a desired predictivity under limited computational resources. A key to this will be the optimal utilization of the available hardware and tailored software that decouples the specification of data on all model, data, and method hierarchies from the mapping of tasks and simulation data to processing units, let it be CPUs, accelerators, or combined setups. This constitutes an automatic method that decides a priori the scheduling of the different simulation components on the available hardware hierarchy and updates at runtime the choice of models, methods, and the parallelization strategy under the constraints of minimal data movement, communication effort, and I/O and predictivity requirements. It is furthermore necessary to automatically react to computational imbalances caused, e.g., by dynamic refinement methods, model, or method changes. This task can be tackled by setting up large-scale simulation databases that include simulative data produced with various numerical methods, error estimates from juxtapositions to experiments, and information on the corresponding computational costs on different architectures. Applying ML methods on such a database to derive the optimal choice for a simulation with respect to the employed numerical methods, error reduction, and computational efficiency under given hardware limitations, is a promising approach and hence of high interest to the CFD community.

To summarize, this workshop will bring together researchers working on all aspects of CFD-based simulations concerned with applying ML as part of the computation or data post-processing with the objective of presenting early results, discussing current state of the art, and establishing new collaborations among the workshop participants. The proposed workshop will give a venue for the community of early adopters to share early successes and lessons learned. The field of application is not limited to the above-mentioned applications, i.e., the derivation of novel coupled simulation techniques that employ RANS models and high-fidelity methods, novel post-processing tools for the analysis of complex flow fields, and the development of expert systems that propose optimal simulation setups.

Relevance and impact of the workshop for ISC

Though the high-fidelity simulations hold promise for better predictive capabilities, the feasibility of running such simulations for design iterations as the norm in an engineering setting is unlikely. With the breakthrough in computational power, the evolution of data science techniques, and the ability to generate terabytes of data from high-fidelity simulations, we aim in this workshop to demonstrate the use of high-fidelity simulations to generate data and utilize it to train ML to better predict the underlying physics in fluid dynamics.

We believe the workshop will be of interest to a broad range of ISC participants. ISC is a premier world conference on HPC technology and applications, with ML becoming a major focus area in recent years. CFD is of interest to both HPC user community and HPC application developers. CFD simulations are responsible for a significant use of HPC resources world-wide. There are numerous efforts in the community to parallelize and improve performance of such codes to run on HPC systems at scale, with some of most recent efforts focusing on harvesting the predictive capabilities of machine learning techniques to replace time-consuming numerical calculations. The applications that rely on CFD-based simulations are very broad, ranging from applied engineering and product design in automotive and aerospace areas to purely research-focused simulations, e.g., in astrophysics or bio-fluid mechanics. The proposed workshop will also be of a high interest for industry interested in vastly reducing its computational efforts by using data from high-fidelity CFD calculations to train the ML model and therefore using it for any subsequent design change with different input parameters.

Convergence of HPC and ML was already the subject of another ISC event, *Workshop on the Convergence of Large Scale Simulation and Artificial Intelligence*. That workshop, however, covered many science domains, from particle physics to earth science, as well as ML and HPC methodology, which makes it very hard for anybody to fully comprehend all the contributions. Compared to it, this proposed workshop is focused on a specific science domain and is very likely to attract a well-defined community of people working in this specific area, while remaining open to closely related fields of science and engineering. We also expect it to be very attractive to the industrial attendees of the conference.

List of tentative program committee members

Jiahuan Cui, Zhejiang University, China

Eloisa Bentivegna, IBM Research, UK

Ashley Scillitoe, The Alan Turing Institute, UK

Charalambos Chrysostomou, The Cyprus Institute, Cyprus

Morris Riedel, University of Iceland, Iceland

Andreas Lintermann, Jülich Supercomputing Centre, Forschungszentrum Jülich, Germany

Jenia Jitsev, Jülich Supercomputing Centre, Forschungszentrum Jülich, Germany

Seid Koric, National Center for Supercomputing Applications, USA

Shirui Luo, National Center for Supercomputing Applications, USA

Madhu Vellakal, National Center for Supercomputing Applications, USA

Volodymyr Kindratenko, National Center for Supercomputing Applications, USA

Jeyan Thiyagalingam, Science and Technology Facilities Council, UK

Format for the workshop (e.g., talks, panel sessions, keynote)

The workshop will consist of 25-minute talks and will conclude with a panel session in which experts working in the field will discuss most pressing challenges and attempt to identify most promising directions for this field to continue developing in the near future. Below is a tentative schedule.

09:00 – 09:15 am Welcome and Introduction

09:15 – 09:45 am Invited Presentation or Keynote

09:45 – 10:10 am Presentation 1

10:10 – 10:35 am Presentation 2

10:35 – 11:00 pm Presentation 3

11:00 – 11:30 am Break and Discussions

11:30 – 11:55 pm Presentation 4

11:55 – 12:20 pm Presentation 5

12:20 – 12:50 pm Panel and Q&A

12:50 – 1:00 pm Wrap-up

This schedule will accommodate 5 technical presentations and it can be modified to add one more, depending on the number of submissions.

Expected outcome from the workshop

We expect to publish papers presented at the workshop. Tentative schedule for soliciting contributions is provided below.

Paper due: March 30th, 2020

Acceptance notification: April 27, 2020

Camera ready: June 1st, 2020

Workshop day: June 25, 2020

We also expect to form a community of researchers working on CFD/ML convergence that will continue to explore these converging topics. The workshop can also facilitate the creation of a database or the identification of canonical problems in fluid dynamics which can be used to fine tune the machine learning algorithms towards better predictability.

Strategies for advertise and attract attendees

The workshop organizing committee will take the responsibility to promote the event among their peers. We will leverage each workshop organizer's professional network and his/her organization's outreach capabilities to reach out to the relevant communities. We will also advertise the event through directed email to relevant mailing lists and social networks. An excellent list of relevant CFD communities to be targeted by this advertisement is provided at https://www.cfd-online.com/Links/cmc.html. We will also advertise the event to HPC-focused communities through NCSA connections as well as relevant industry participants through NCSA Industry program, which is one of the largest academic research industrial HPC programs in the world.