Deep Recurrent Neural Networks for Optical Flow Learning in Particle-Image Velocimetry

Christian Lagemann[1*] & Wolfgang Schröder[1] (supervisor)

[1] Chair of Fluid Mechanics and Institute of Aerodynamics, RWTH Aachen University, Wüllnerstraße 5a, 52062 Aachen,
[*] email: c.lagemann@aia.rwth-aachen.de, phone: +49 163 8359284

High-level summary

Particle-Image Velocimetry (PIV) is a key technique in modern experimental fluid mechanics used to determine the flow velocity in a wide range of engineering and physics related problems. PIV is a non-intrusive, optical method that adds buoyant particles to the flow, which adopt the velocity of the surrounding fluid. As shown in Figure 1 (a), these tracer particles are illuminated by a thin high-power laser light sheet and a camera records double images separated by a short time interval to capture the particles' displacements. Given data from a PIV experiment, the crucial task is to determine the underlying displacement field, i.e., the vector field describing the fluid flow. The gold-standard in traditional PIV processing divides the original particle image into finite-size interrogation windows and cross-correlates the corresponding interrogation windows of two consecutive images. The position of the correlation peak indicates the averaged displacement within the respective image section (see Figure 1 (b)). This approach substantially lowers the spatial resolution of the resulting velocity field compared to the original image size and smooths velocity gradients existing in each sub-window. Furthermore, the high effort needed to extend the capabilities of these tools represents a major limitation for further development in PIV. Motivated by these restrictions of classical methods, we combine ideas from deep optical flow learning with PIV analysis and introduce a novel pipeline for arbitrary PIV experiments called Recurrent All-pairs Field Transforms for PIV (RAFT-PIV). In contrast to the traditional PIV evaluation, RAFT-PIV provides a velocity vector for each pixel. Thus, the spatial resolution is massively increased, which enables the estimation of small-scale flow features and localized velocity gradients that cannot be resolved by traditional routines without changing the experimental setup. As demonstrated in various applications spanning synthetic and experimental data, this characteristic may state a turning point for experimental fluid mechanics since it enables complex downstream analysis tasks from standard PIV measurements. That is, RAFT-PIV uniquely allows for a precise analysis of near-wall flow quantities such as the friction velocity and the wall-shear stress as well as relevant flow metrics, e.g. turbulent stresses and turbulence spectra, from simple yet easily feasible standard PIV measurements. To the best of our knowledge, such general spectrum of analysis possibilities including temporally and spatially resolved results cannot be provided by any other fluid flow measurement technique.

Scientific summary

PIV is a key approach in experimental fluid dynamics and of fundamental importance in diverse applications, including automotive, aerospace, and biomedical engineering. The current state-of-the-art in PIV data processing involves traditional handcrafted models which are subject to limitations including the high manual effort required and difficulties in generalizing across conditions. Although widely used, these existing tools have a number of well-known shortcomings, including limited spatial output resolution, gradient smoothing, and peak-locking biases. In contrast, the deep learning-based approach introduced in this work is general, largely automated, and provides a high spatial resolution. Extensive experiments, including benchmark examples where true gold-standards are available for comparison, demonstrate that the proposed approach achieves state-of-the-art accuracy and generalization to new data, relative to both classical approaches and other optical flow learning schemes. Moreover, it enables a precise analysis of near-wall flow quantities and relevant flow metrics from simple yet easily feasible standard PIV measurements. To be precise, these quantities can be analyzed in a spatially and temporally resolved fashion without applying any further modeling assumptions as typically required by other methods. To the best of our knowledge, RAFT-PIV is the only experimental workflow pipeline which can offer such a general flexibility. Before delving into the results, a short overview of the method is provided.

Architecture: RAFT-PIV is a neural optical flow estimator which is specifically designed for the use case of PIV images. It mainly consists of three stages: a feature extracting block, the computation of a pixelwise correlation volume, and iterative updates based on a Convolutional Gated Recurrent Unit (ConvGRU) as shown in Figure 1 (c). In a first step, the shared feature encoder derives latent embeddings for each input image individually using three convolutional neural network modules. In the second stage, the similarity of both image features is calculated using a full correlation volume between all pairs of both feature maps. Subsequently, a 4-layer correlation pyramid is formed by pooling the last two dimensions sequentially from level to level. Thus, RAFT-PIV maintains high resolution information of the first image while effectively addressing large object displacements. In the final stage, the estimated flow field is updated using a ConvGRU, which is a recurrent neural network that stores hidden state information of previous steps to modulate a limited content memory. Hence, it allows the network to balance the prediction of earlier optical flows and its current hidden state to compute a new optical flow update. To enhance the generalization ability of RAFT-PIV, a complementary global motion aggregation (GMA) module is implemented. It considers long-range feature connections extracted by a self-attention mechanism and estimates motion features which are propagated to occluded regions. Unlike convolutions, the self-aware attention formulation principally grants access to all parts of the entire input sequence such that all pixel embeddings are considered simultaneously and the network learns to pay attention to the most relevant image features.

Results: To demonstrate the superior spatial resolution of RAFT-PIV, results of a standard benchmark test case introduced in the third international PIV challenge are shown in Figure 1 (d). The underlying flow field com-

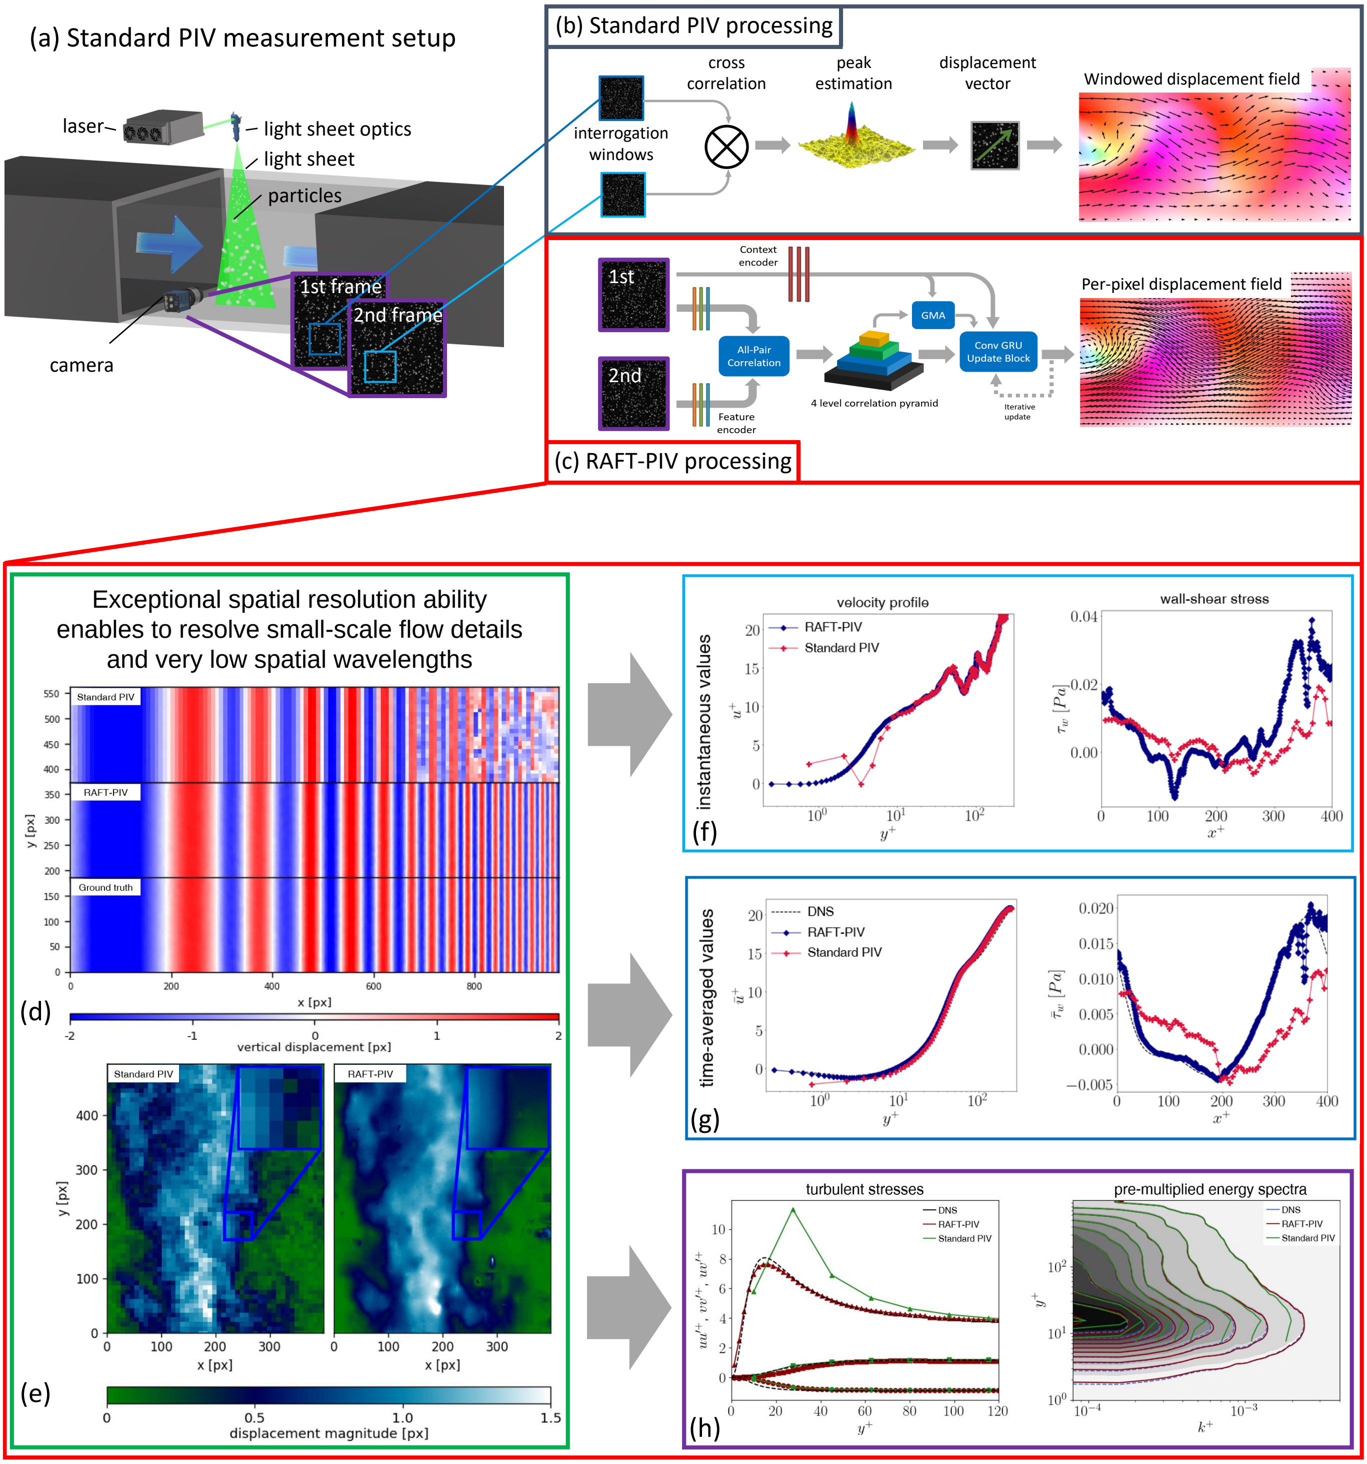

Figure 1: Sketch of a standard PIV setup (a), the standard vs. the novel PIV processing workflow (b & c), and examples of RAFT-PIV applications (d - h).

prises a wavelength-modulated one-dimensional shear displacement. The results illustrate that a standard cross-correlation based PIV approach is barely capable of resolving the low wavelengths. In contrast, our RAFT-PIV model closely matches the ground truth, i.e., it accurately estimates the local displacements even for the lowest wavelengths and very fine spatial scales. This is one of the most significant advantages of this neural method. However, since synthetic images cannot model all physical dependencies and measurement uncertainties properly, it is a must to evaluate the performance of RAFT-PIV on a variety of experimental data. In this thesis, an exceptional range of real-world measurements have been analyzed. Amongst others, these includes important generic flows of academic research such as various turbulent boundary layer, channel, and jet flows as well as engineering and human health related applications, e.g., laminar oscillating flows in flexible blood vessels, flows inside combustion engines, and high Re-number flows around airfoils. Exemplary results are shown in Figure 1 (e). For instance, instantaneous flow predictions of a turbulent jet flow demonstrate impressively the high-resolution advantage of RAFT-PIV. The obtained flow fields appear more realistic and enable a distinct identification of the bounding shear layer. Thus, small flow scales of the free shear layer and complementary mixing processes can be studied in greater detail which are lost when a standard PIV tool is used. To demonstrate this important advantage in more detail, Figures 1 (f-h) also depict results of extremely valuable but typically hardly measurable near-wall flow quantities. The blue tiles frame instantaneous and time-averaged velocity profiles and wall-shear stress distributions obtained from experimental turbulent wall-bounded flows. Note that complementary DNS data is also available for a comparison of the time-averaged flow quantities. Obviously, such a comparison is not possible for instantaneous data but synthetic test cases are studied in the thesis to demonstrate the impressive accuracy and reliability of RAFT-PIV in this context. For the experimental data presented in Figure 1(f & g), RAFT-PIV matches the DNS data exceptionally well. Contradictory, standard PIV tools underpredict the strong velocity gradients in the viscous sublayer and buffer layer and cannot provide valid velocity estimates in these regions. This issue is well reported in the literature and becomes obvious for the instantaneous velocity profile as well. As a result, wall-shear stress distributions derived from these spurious velocity estimates are barely physically meaningful and do not match the complementary DNS data. In contrast, since RAFT-PIV accurately predicts the friction velocity close to the wall, the wall-shear stress closely follows the DNS.

Similar findings can be reported when analyzing higher-order turbulence statistics, e.g. normal/shear stresses and energy spectra, of a turbulent channel flow (see Figure 1 (h)). RAFT-PIV accurately predicts the velocity fluctuations and the derived statistics and reliably estimates small-scale features even for very small wall distances. These can usually not be resolved with a standard PIV processing as shown in the provided plots.

To summarize, a deep learning based optical flow estimator called RAFT-PIV was developed in this thesis, which substantially outperforms current gold-standard PIV methods w.r.t. processing speed, spatial resolution ability, physical significance, robustness, and handling. RAFT-PIV can serve as a direct one-to-one replacement for existing PIV processing methods and proved its robustness and generalization ability for a large range of synthetic and experimental measurement data in an out-of-the-box fashion. Thorough analysis evidenced that RAFT-PIV can resolve much smaller flow structures compared to existing methods and allows to derive relevant turbulence statistics even in the low wavelength regime. Moreover, based on its tremendous spatial resolution, RAFT-PIV further allows to compute near-wall flow quantities accurately, which enables researchers to directly infer instantaneous wall-shear distributions over the entire spatial image domain from standard PIV images. This is not possible with any measurement technique in a similar way. Hence, RAFT-PIV will mark a turning point how we approach the measurement of traditionally challenging flow quantities and may constitute a game-changer for the experimental fluid dynamics community. It also opens the door for a completely new approach to validate reliably high-fidelity simulations favoring both, future and existing measurement data which can be reevaluated easily.

Publications during thesis

A large number of peer-reviewed publications and conference contributions have been published as part of this thesis highlighting the relevance and novelty of the RAFT-PIV approach. In total, 6 peer-reviewed publications (3 publications are currently under review) have been authored in scientific and renowned journals including Nature Machine Intelligence and Nature Computational Science. Additionally, 10 contributions to international conferences have been accepted including presentations at ICTAM, ETC, ISPIV, and LX-Laser conferences. Notably, two invited talks hosted by the ERCOFTAC PC Germany South and the Center of AI at the RWTH Aachen were given. A list of selected publications is provided below:

- Lagemann et al., "Deep recurrent optical flow learning for particle image velocimetry data", *Nat. Mach. Int.* 3:641-655 (2021).
- Lagemann et al., "Generalization of deep recurrent optical flow estimation for particle-image velocimetry data", *Meas. Sci. & Tech.* 33 (9) 094003 (2022).
- Lagemann et al., "Deep Learning of Causal Structures in High Dimensions", *Nat. Mach. Int.* (2022 - under review).
- Lagemann et al., "Wall-shear stress estimation with neural network enhanced particle-image velocimetry", *Nat. Comp. Sci.* (2023 - under review).
- Lagemann et al., "Challenges of deep unsupervised optical flow estimation for particle-image velocimetry data", *Exp. in Fl.* (2023 - under review).