Proceedings of the Canadian Society for Mechanical Engineering International Congress
32nd Annual Conference of the Computational Fluid Dynamics Society of Canada
Canadian Society of Rheology Symposium
CSME/CFDSC/CSR 2025
May 25–28, 2025, Montréal, Québec, Canada

ID 234

Enhancing noisy PIV measurements through Signal Processing and Machine Learning Techniques

Blaise Madiega[1*], Mathieu Olivier[1], Konrad Tagnon Amen Alahassa[2], Christ-Marie Anzian Brou[2]
[1]Mechanical Engineering, Université Laval, Quebec City, Canada
[2]Computer Engineering, Université Laval, Quebec City, Canada
[*]blaise.madiega.1@ulaval.ca

***Abstract*—Particle Image Velocimetry (PIV) is a cornerstone technique in experimental fluid dynamics, yet its accuracy and reliability can be severely affected by noise and insufficient resolution in the measured velocity fields. In this work, we propose a deep learning framework that integrates signal processing methods with a hybrid U-Net-LSTM convolutional neural network to enhance data quality and resolution. Our approach applies filtering techniques for noise reduction and leverages learned feature extraction for improved particle segmentation, ultimately boosting measurement fidelity. We evaluate this pipeline using a specialized PIV-UQ dataset containing raw images (PIV-MS, low resolution measurement system) and high-fidelity reference measurements from a stereoscopic system (PIV-HDR). Results indicate that effective denoising in conjunction with the U-Net-LSTM architecture significantly refines pixel-level velocity estimations. (Code available at Noisy PIV Enhanced by MLSP-GitHub Repo)

***Keywords-component*—Deep Optical Flow Learning, Image Denoising, PIV, Experimental Fluid Mechanics**

I. Introduction

Particle Image Velocimetry (PIV) is almost omnipresent in experimental fluid mechanics, providing spatiotemporal measurements that enable detailed insights into flow behavior. Despite its utility, PIV remains highly susceptible to noise and resolution limitations, which can degrade the accuracy and reliability of velocity estimations.

Recent advancements have predominantly focused on leveraging deep learning to enhance PIV measurements. For example, the Deep-learning Optical Flow (DLOF) algorithm demonstrated superior performance over traditional PIV by providing higher-resolution and smoother velocity fields, especially in densely labeled active nematics systems [1]. Similarly, RAFT-StereoPIV introduced a deep learning model that integrates convolutional gated recurrent units (Conv-GRU) with stereoscopic PIV data [2], achieving significant error reductions and robust performance in turbulent aerodynamic measurements.

Additionally, CNN-based approaches like PIV-NetS [3] have improved dense motion estimation in PIV by improving accuracy and computational efficiency, particularly in resolving small-scale flow structures within synthetic datasets. LiteFlowNet [4] further contributes to the field by offering a lightweight CNN architecture for optical flow estimation, achieving high accuracy with fewer parameters and faster processing times.

Among these advanced methods, UnLiteFlowNet-PIV [5] has emerged as a state-of-the-art solution tailored specifically for PIV measurements. UnLiteFlowNet-PIV utilizes a deep learning architecture that excels in estimating optical flow for PIV data. It builds upon the strengths of prior optical flow methods and enhances them by providing accurate flow estimations with a lower computational cost. In fact, UnLiteFlowNet-PIV has demonstrated competitive performance against classical PIV techniques, as well as supervised learning-based approaches, often outperforming them in challenging flow scenarios.

Despite their success, most of these methods rely heavily on synthetic datasets and do not robustly address real-world noise, leaving a noticeable gap in integrating signal processing techniques with machine learning for enhanced data fidelity.

This study bridges this gap by combining classical denoising methods with a hybrid U-Net-LSTM architecture to improve the quality and resolution of noisy PIV measurements, offering a noise-tolerant solution for experimental fluid mechanics.

II. Methodology

A. Dataset Description

The dataset (PIV-UQ) [6] was obtained from experiments conducted at the Experimental Fluid Dynamics Laboratory (EFDL), Utah State University, using a rectangular jet facility (aspect ratio 7.2) with a jet exit velocity of $u_0 = 5\,\text{m/s}$, yielding a Reynolds number of 3,000 (based on the jet height, $h = 10.2\,\text{mm}$).

Two PIV systems were used to obtain the database PIV-UQ:

(1) **PIV-MS**, a single-camera system providing 2D-2C velocity data with a digital resolution of $7.2\,\text{px/mm}$ at $3000\,\text{fps}$;

(2) **PIV-HDR**, a stereoscopic system capturing 2D-2C or 3D-3C data with $26.7\,\text{px/mm}$ at $5400\,\text{fps}$. The PIV-HDR system, characterized by its superior resolution and lower measurement error (reduced by a factor of 3–4), serves as the high-resolution ground truth for this study. In contrast, the PIV-MS system, with its lower resolution, provides noisy measurement data, emulating input conditions for image-processing and flow-estimation tasks.

B. Denoising Techniques

To preprocess the noisy PIV-MS data, we applied several denoising techniques tailored to enhance the signal-to-noise ratio while preserving fine-scale flow structures.

- **Gaussian Filtering:** Convolution with a Gaussian kernel of standard deviation σ ($\sigma = 1 - 1.5$).

- **Median Filtering:** Non-linear filter using a square kernel of size $s \times s$ ($s = 3 - 5$).

- **Non-Local Means (NLM):** Leveraging patch-based similarity with parameters h (smoothing factor), patch size, and search distance.

We also applied advanced denoising techniques capable of handling complex noise distributions that may not be strictly additive. These methods aim to **enhance the signal-to-noise ratio** while preserving fine-scale flow features and include:

- **Total Variation (TV) Minimization [7]:** Reduces noise by minimizing the total variation norm of the image, effectively preserving sharp edges and discontinuities in flow features.

- **Anisotropic Diffusion (Perona-Malik Filter):** An iterative partial differential equation (PDE)-based method that smooths noise selectively while maintaining strong gradients across edges.

For each method, the noisy input data was passed through an iterative denoising pipeline, and the hyperparameters were optimized to minimize residual noise while preserving the flow field's velocity gradients.

C. Models' Architecture

Baseline model

For our baseline experiments, we maintain the original architecture of the UnLiteFlowNet-PIV model as described in the reference paper. The model is specifically designed for optical flow estimation, with a lightweight and efficient design that ensures competitive performance while minimizing computational overhead. The training process employs an unsupervised loss function, which is a combination of three components: Photometric loss, Flow smoothness loss and the Consistency loss. Additionally, as outlined in the original paper, we utilize a multi-scale resolution loss. This loss is computed as a weighted sum of the estimation losses from each of the intermediate layers in the model. For both the unsupervised loss and the multi-scale resolution loss, we retain the original weights assigned to each component as described in the baseline paper. This ensures consistency with the baseline methodology and provides a solid foundation for comparative analysis.

U-Net-LSTM

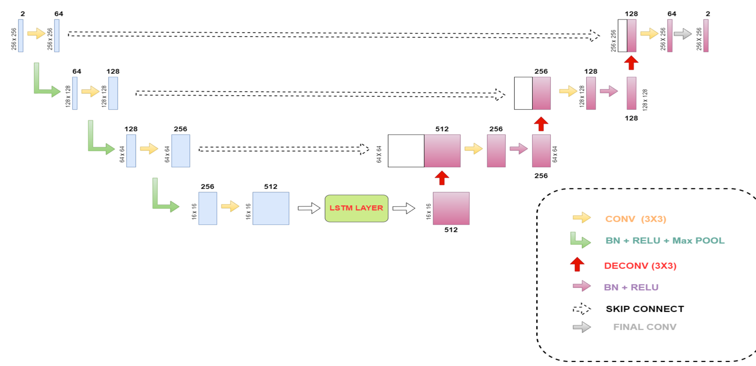

Figure. 1. U-Net-LSTM Architecture

The proposed architecture combines a UNet-based convolutional structure with LSTM layers to process sequential PIV images. The UNet extracts multi-scale spatial features, while the LSTM captures temporal dependencies across image sequences. This hybrid design ensures accurate reconstruction of dynamic flow fields by integrating spatial and temporal information effectively.

Methodology Illustration

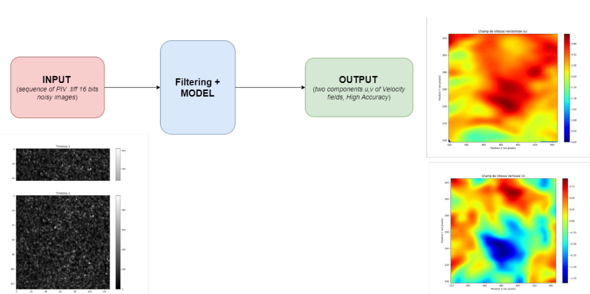

Figure. 2. Methodology Illustration

Figure 2 presents the pipeline for particle velocity prediction: images captured at two consecutive time points undergo filtering techniques and are subsequently processed by the U-Net-LSTM model to estimate particle velocities.

III. Experimentation

A. Baseline

To establish a reliable baseline for our proposed pipeline, we utilize the UnLiteFlowNet-PIV model, a deep learning framework specifically designed for optical flow estimation in PIV data. For this study, the baseline model is trained from scratch using the SQG and PIV-UQ datasets. These datasets provide diverse flow conditions and noise levels, allowing us to assess the model's performance comprehensively.

The training process adheres closely to the methodology outlined in UnLiteFlowNet-PIV, with a few modifications aimed at improving the learning dynamics. Specifically, we introduce a StepLR learning rate scheduler, which reduces the learning rate by a fixed factor every 20 epochs, starting from an initial value of 1×10^{-4}. This gradual reduction helps maintain stability during training and prevents overfitting.

To ensure compatibility with the baseline model's architecture, we preprocess the PIV-UQ dataset by resizing the images from their original resolution of 120×120 pixels to 256×256 pixels using bilinear interpolation. This resizing step standardizes the input size across datasets and aligns with the requirements of the UnLiteFlowNet-PIV model.

The baseline model is trained for 100 epochs on both the SQG and PIV-UQ datasets under identical training conditions. Additionally, to evaluate the impact of noise reduction, we train separate baseline models incorporating various denoising techniques. These methods, applied as preprocessing steps allows us to compare the model's performance under different noise conditions and highlights the benefits of integrating signal processing into the training pipeline.

B. U-Net-LSTM

The hyperparameters of the proposed U-Net-LSTM network were optimized to achieve better performance for denoising and reconstructing PIV velocity fields. Key parameters include the batch size (32), learning rate (10^{-2} with a StepLR scheduler), and sequence length (1).

The network's architecture starts with 16 base filters in the convolutional layers and employs a single LSTM layer to handle temporal dependencies. The number of epochs was set to 600 to ensure convergence. Advanced hyperparameter (including denoising methods' parameters) optimization was conducted using the **Optuna** library, leveraging its efficient search algorithms to systematically explore the hyperparameter space. This process allowed the fine-tuning of architectural elements and training parameters to maximize the model's accuracy while maintaining computational efficiency.

IV. Results

A. Baseline

Table I presents the Average Endpoint Error (AEE) (averaged over all impage pairs) of the baseline model under different training and testing setups. The Average Endpoint Error (AEE) can be described as the L_2-norm of the difference between the flow estimation $\mathbf{F}_e$ and the flow ground truth $\mathbf{F}_g$:

$$\text{AEE} = \|\mathbf{F}_e - \mathbf{F}_g\|_2$$

TABLE. I
AEE of the baseline model on different setups.

Dataset	Train AEE	Test AEE
SQG	0.20606	0.20631
PIV-UQ	3.79355	2.03283
SQG + Gaussian smoothing	0.22859	0.23015
SQG + Median filter	0.26701	0.26593

As shown in Table I, the model achieves its lowest AEE when trained on the SQG dataset. Although the AEE appears lower when the model is trained on the PIV-UQ dataset, further investigation reveals that the model trained on SQG generalizes better to unseen data.

To illustrate this, we analyze an example of particle velocity prediction at a specific coordinate on the test set over time, shown in Figures 3 and 4. At the coordinate (60,60), predictions from the model trained on SQG align more closely with the ground truth compared to those from the model trained on PIV-UQ. This indicates that the model's ability to generalize is better when trained on the synthetic SQG dataset.

Furthermore, the application of denoising techniques—such as Gaussian smoothing and median filtering—prior to inputting image pairs into the model does not improve the model's AEE. This is particularly evident in the SQG dataset, where denoising slightly increases the AEE compared to the unprocessed data. Due to the poor performance of the baseline on the PIV-UQ dataset, denoising techniques were not applied when training the model on this dataset.

These results suggest that the baseline model struggles to perform well on real-world data like PIV-UQ, and that applying denoising techniques to synthetic datasets such as SQG does not enhance performance. The discrepancy can likely be attributed to the differences in the nature of the datasets: SQG consists of synthetic data, whereas PIV-UQ is derived from real experimental setups. The noise characteristics and complexity of synthetic data may limit the effectiveness of denoising techniques we use here in this context.

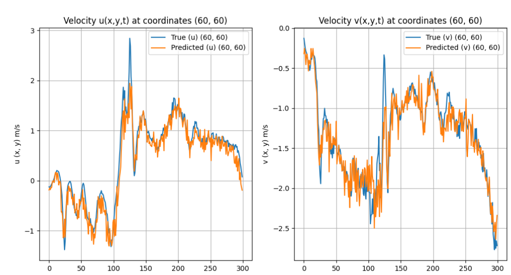

Figure. 3. Prediction of particle velocity at (60,60) when the baseline model (UnLiteFlowNet-PIV) model is trained on synthetic SQG dataset.

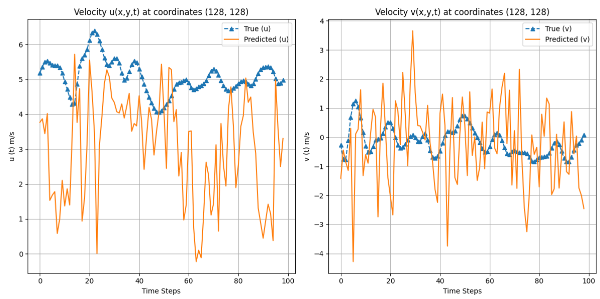

Figure. 4. Prediction of particle velocity at (128,128) when the baseline model (UnLiteFlowNet-PIV) is trained on PIV-UQ dataset.

B. U-Net-LSTM

The velocity profile comparison, illustrated in Figure 5, highlights the effectiveness of integrating denoising methods with the U-Net-LSTM architecture. The profiles obtained after applying anisotropic diffusion (Perona-Malik Filtering) demonstrate the best alignment with the ground truth data across both the $u(t)$ and $v(t)$ directions.

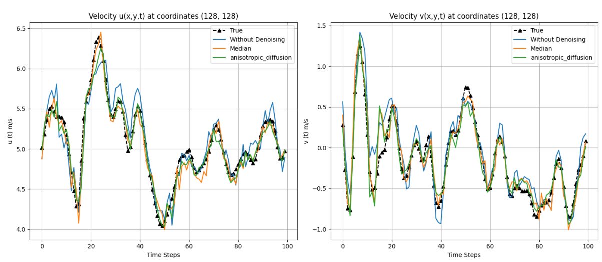

Figure. 5. Comparison of particle velocity's estimation at (128,128) on x and y-direction, u(t) & v(t)

This contrasts with the profiles derived without denoising or with other filtering techniques, which tend to deviate more prominently, particularly in regions of high velocity gradients. The smoothing effect of anisotropic diffusion appears to enhance the precision of the particle velocity estimations. Median filtering, while slightly less effective, also contributes significantly to reducing discrepancies, as evidenced by its improved velocity profiles over baseline predictions.

TABLE. II
AEE OF U-NET-LSTM USING DIFFERENT DENOISING METHODS ON PIV-UQ TEST DATA

Method	AEE	Ratio
Without Denoising	0.2178	—
Gaussian Filtering	0.1731	20.52 % ↓
Median Filtering	0.1723	20.89 % ↓
Non-Local Means	0.2325	6.75 % ↑
Total Variation (TV) Minimization	0.2061	5.37 % ↓
Anisotropic Diffusion	**0.1456**	**33.15** % ↓

Table II quantifies the Average Endpoint Error (AEE) for different denoising methods applied to the PIV-UQ dataset. Among the methods tested, anisotropic diffusion achieves the lowest AEE, reducing the error by **33.15** % compared to predictions without denoising. This substantial improvement underscores the method's capability to retain essential flow features while eliminating noise. Median filtering similarly yields a notable reduction in AEE (20.89%), outperforming Gaussian filtering and other traditional techniques. Interestingly, non-local means and total variation minimization, though effective in other contexts, do not perform as well in this scenario. These findings affirm the suitability of anisotropic diffusion and median filtering for enhancing the performance of U-Net-LSTM on noisy experimental datasets like PIV-UQ.

V. CONTRIBUTIONS

This work demonstrates that integrating denoising techniques with a hybrid U-Net-LSTM deep learning model significantly improves the processing of Particle Image Velocimetry (PIV) data. By combining classical denoising methods, such as Gaussian filtering and total variation minimization, with a neural network architecture, our approach enhances pixel-level velocity estimation and robustness, particularly on the challenging PIV-UQ dataset. Compared to prior methods like UnLiteFlowNet-PIV, our selective preprocessing reduces noise-induced errors and improves generalization to real-world data, showcasing the importance of combining denoising and deep learning for superior PIV measurement performance.

VI. CONCLUSION

In this study, we demonstrated that while the baseline model struggles to make accurate predictions on the noisy, real-world PIV-UQ dataset, it generalizes well when trained on synthetic datasets like SQG. In contrast, the U-Net-LSTM model shows robust performance even without denoising, achieving decent predictions on noisy experimental data. Furthermore, applying denoising techniques such as anisotropic diffusion and median filtering significantly enhances the model's accuracy, reducing noise-induced errors while preserving critical flow features. These findings highlight the potential of integrating deep learning architectures with signal processing methods to address challenges in processing noisy experimental PIV data, paving the way for more reliable and accurate velocity field estimations in complex flow scenarios. Expanding the dataset variety and benchmarking against SOTA architectures will further validate the robustness of our method and identify opportunities to refine it for broader applicability in real-world PIV scenarios.

VII. RELATION TO PRIOR WORK

Our project builds upon the advancements in deep learning methods for Particle Image Velocimetry (PIV) while addressing their limitations in handling noisy experimental data. Prior works, such as UnLiteFlowNet-PIV [5], have demonstrated the effectiveness of optical flow estimation using lightweight CNN architectures tailored for PIV, achieving competitive performance on synthetic datasets. However, these approaches often lack robustness against real-world noise and fail to integrate signal processing techniques. By combining classical denoising methods, such as Gaussian filtering and anisotropic diffusion, with a hybrid U-Net-LSTM architecture, our work advances the field by offering a noise-tolerant framework. This integration bridges the gap between signal processing and machine learning, significantly enhancing the fidelity of velocity estimations in noisy experimental environments, as demonstrated on the challenging PIV-UQ dataset.

REFERENCES

[1] Phu N. Tran, Sattvic Ray, Linnea Lemma, Yunrui Li, Reef Sweeney, Aparna Baskaran, Zvonimir Dogic, Pengyu Hong, and Michael F. Hagan, "Deep-learning optical flow outperforms piv in obtaining velocity fields from active nematics," 2024.

[2] Mohamed Elrefaie, Steffen Huttig, Mariia Gladkova, Timo Gericke, Daniel Cremers, and Christian Breitsamter, "Real-time and on-site aerodynamics using stereoscopic piv and deep optical flow learning," 2024.

[3] Shengze Cai, Shichao Zhou, Chao Xu, and Qi Gao, "Dense motion estimation of particle images via a convolutional neural network," *Experiments in Fluids*, vol. 60, 2019.

[4] Tak-Wai Hui, Xiaoou Tang, and Chen Change Loy, "Liteflownet: A lightweight convolutional neural network for optical flow estimation," 2018.

[5] Mingrui Zhang and Matthew D. Piggott, "Unsupervised learning of particle image velocimetry," *arXiv preprint arXiv:2007.14487*, 2020.

[6] Douglas Neal et al., "Collaborative framework for piv uncertainty quantification: the experimental database," 2015.

[7] Leonid I. Rudin, Stanley Osher, and Emad Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: Nonlinear Phenomena*, vol. 60, no. 1–4, pp. 259–268, 1992.