

A lean probabilistic CNN for PIV enabling joint flow and reconstruction uncertainty estimation

C. Mucignat^{1,*}, K. Zydbal¹, H. Gorji¹, I. Lunati¹

¹: Laboratory for Computational Engineering, Swiss Federal Laboratories for Materials Science and Technology, Switzerland

*Corresponding author: claudio.mucignat@empa.ch

Keywords: PIV processing, CNN, Uncertainty.

ABSTRACT

Particle Image Velocimetry (PIV) is affected by multiple sources of error arising from imaging conditions, tracer dynamics, and three-dimensional flow effects. While established approaches estimate uncertainty from correlation statistics or particle matching, convolutional neural network (CNN)-based PIV methods typically provide only deterministic displacement fields, limiting quantitative assessment of reconstruction reliability and experimental optimization. In this work, we introduce *P-LIMA*, a lean probabilistic CNN for PIV that jointly reconstructs the displacement field and a spatially resolved, data-dependent uncertainty estimate. *P-LIMA* extends the lightweight *LIMA* architecture by augmenting its output with per-component scale parameters and training the network using a heteroscedastic negative log-likelihood formulation based on Gaussian or Student-t distributions, enabling robust modeling of aleatoric uncertainty. Physically motivated regularization terms on displacement gradients are retained to ensure spatial consistency, while a delayed activation of the probabilistic loss stabilizes training and improves the calibration of the predicted uncertainty. The method is evaluated on synthetic datasets generated from direct numerical simulations, experimental PIV measurements, and the PIV Challenge 3 (Case B and C) under degraded imaging conditions. Coverage analysis on validation data shows close agreement between empirical and nominal confidence levels, indicating well-calibrated uncertainty estimates. Compared to classical correlation-based approaches, *P-LIMA* produces uncertainty fields that are more coherent and more strongly correlated with reconstruction errors, while preserving the accuracy of the deterministic baseline model. With approximately 500k parameters and real-time to near-real-time inference on modern GPUs, *P-LIMA* provides uncertainty estimates without sacrificing computational efficiency.

1. Introduction

Quantifying the uncertainty of Particle Image Velocimetry (PIV) measurements is essential for assessing data reliability, comparing experimental conditions, and guiding the design of measurement setups. Classical uncertainty estimation approaches rely on correlation statistics Wieneke (2017) or particle-matching metrics Sciacchitano et al. (2013), which provide valuable indicators

of reconstruction quality based on imaging properties such as signal-to-noise ratio and particle loss. However, a complete uncertainty assessment would also require accounting for additional error sources, including out-of-plane motion, three-dimensional velocity gradients across the laser sheet, and tracer inertia effects Raffael et al. (2018), for which the necessary information is not directly accessible from standard PIV recordings.

In recent years, convolutional neural networks (CNNs) have been successfully adopted for PIV image processing, offering improved robustness and computational efficiency compared to classical algorithms such as *WIDIM*. Architectures originally developed for optical flow estimation, such as FlowNet Dosovitskiy et al. (2015), PWC-Net Sun et al. (2018), and RAFT Teed & Deng (2020), have been adapted to PIV and shown to deliver accurate displacement fields. Despite these advances, most CNN-based PIV methods remain purely deterministic, forcing users to judge the reliability of the reconstructed flow fields based on experience rather than quantitative uncertainty measures. This limitation represents a major obstacle to the adoption of machine-learning-based PIV in uncertainty-aware experimental workflows.

In the context of machine learning, it is customary to distinguish between epistemic and aleatoric uncertainty Kendall & Gal (2017). Epistemic uncertainty reflects model uncertainty due to limited or biased training data and is typically estimated using Bayesian neural networks or ensemble approaches. While conceptually appealing, these methods are computationally expensive and poorly suited to high-throughput or real-time PIV applications. Aleatoric uncertainty, on the other hand, captures data-dependent noise intrinsic to the observations, such as particle loss, low illumination, poor focus, or background contamination. This form of uncertainty is particularly relevant for PIV, as imaging conditions can vary significantly across the field of view and over time.

Aleatoric uncertainty can be estimated efficiently by augmenting deterministic CNNs with additional output channels that parameterize an assumed error distribution and training the model using a heteroscedastic negative log-likelihood (NLL) loss. Although this approach is widely used in probabilistic regression, it has seen limited adoption in PIV applications. Existing attempts are either restricted to coarse spatial resolution Morrell et al. (2020) or rely on uncertainty heads attached to large optical-flow architectures such as RAFT, resulting in substantial computational overhead Wang et al. (2025); Mucignat et al. (2025).

In this work, we introduce *P-LIMA*, a probabilistic extension of the lightweight *LIMA* architecture specifically designed for PIV. The proposed model jointly reconstructs dense displacement fields and spatially resolved aleatoric uncertainty estimates at minimal additional computational cost. The main contributions of this work are:

- a lean probabilistic PIV architecture that scales uncertainty estimation to dense, high-resolution flow fields;
- a stabilized training strategy based on delayed activation of the NLL loss to prevent degen-




Figure 1. Schematic of P-LIMA , need to be changed!!!.

erate uncertainty inflation;

- a systematic evaluation of uncertainty calibration and spatial structure on synthetic, experimental, and benchmark PIV datasets.

By preserving the efficiency and accuracy of the deterministic LIMA model while providing calibrated uncertainty estimates, P-LIMA enables uncertainty-aware CNN-based PIV suitable for on-the-fly experimental diagnostics.

2. Methods

2.1. Baseline architecture

LIMA (Manickathan et al., 2023; Mucignat et al., 2026) is a lightweight CNN framework for estimating tracer displacement fields from PIV image pairs using an iterative coarse-to-fine strategy. The network consists of a convolutional pyramid encoder and a correlation-based decoder operating on progressively refined displacement estimates. The architecture and its optimization for PIV displacement scales are described in detail in Mucignat et al. (2026).

In the present work, we adopt a three-level encoding configuration with reconstruction truncated at the first encoding level, corresponding to the LIMA^{3,1} variant. This configuration provides a favorable compromise between spatial resolution, receptive field size, and computational efficiency

2.2. Probabilistic extension (P-LIMA)

To enable uncertainty quantification, the deterministic LIMA architecture is extended into a probabilistic formulation, referred to here as P-LIMA, by introducing an additional decoder head dedicated to the estimation of per-component uncertainty scales (σ_u, σ_v) . The displacement field

$\mathbf{d} = (u, v)$ and the associated uncertainty estimates are produced by two distinct decoder branches operating at the output pyramid level.

The displacement decoder follows the standard LIMA formulation and predicts incremental flow updates based on correlation features and the current displacement estimate. In contrast, the uncertainty decoder is evaluated only at the final reconstruction level and is specifically designed to estimate heteroscedastic, spatially varying uncertainty fields. Its input consists of the same correlation features used by the displacement decoder, augmented with the final displacement estimate and the latest displacement increment. In this way, the uncertainty prediction is explicitly conditioned on both local matching ambiguity (as encoded by the correlation volume) and the recovered flow structure.

The uncertainty decoder outputs logarithmic uncertainty scales $(\log \sigma_u, \log \sigma_v)$, which are subsequently clamped to physically plausible bounds and interpreted as per-component standard deviations expressed in pixel units. No uncertainty prediction is performed at coarser pyramid levels, where the flow estimate remains provisional.

Importantly, the probabilistic extension is formulated such that the mean displacement estimator and the uncertainty estimator are decoupled. In the experiments presented here, the displacement network is first trained deterministically. Subsequently, the displacement estimator is frozen and only the uncertainty decoder is optimized by minimizing a negative log-likelihood objective. This separation avoids degradation of the deterministic flow accuracy and ensures that the predicted uncertainty reflects the statistics of the residual displacement error rather than compensating for changes in the mean estimate.

The introduction of the second decoder head increases the number of trainable parameters from approximately 500 k to about 1 M, while remaining substantially more lightweight than optical-flow-based architectures such as RAFT.

2.3. Loss Functions and Uncertainty Modeling

When uncertainty estimation is not needed, the network is trained using a deterministic error computed on the reconstructed displacement field $\mathbf{d}$. Standard choices include the $\mathcal{L}_{L2}$ and $\mathcal{L}_{L1}$ losses,

$$\mathcal{L}_{L2} = \|\mathbf{d} - \mathbf{d}_{GT}\|_2^2, \quad \mathcal{L}_{L1} = \|\mathbf{d} - \mathbf{d}_{GT}\|_1, \quad (1)$$

where $\mathbf{d}_{GT}$ denotes the reference displacement. In practice, a Charbonnier loss is preferred for improved numerical stability,

$$\mathcal{L}_{charb} = \sqrt{\|\mathbf{d} - \mathbf{d}_{GT}\|_2^2 + \epsilon^2}, \quad (2)$$

with $\epsilon = 10^{-3}$.

When confidence estimation is required, a heteroscedastic negative log-likelihood (NLL) formulation is adopted Kendall & Gal (2017). Assuming Gaussian-distributed errors, the per-pixel loss reads

$$\mathcal{L}_{Gauss} = \frac{1}{2} \left(\frac{\|\mathbf{d} - \mathbf{d}_{GT}\|_2^2}{\sigma^2} + 2 \log \sigma \right), \quad (3)$$

where $\mathbf{d}$ and $\mathbf{d}_{GT}$ denote the predicted and ground-truth displacement vectors, respectively, and σ represents a per-component uncertainty scale associated with the likelihood. To improve robustness to outliers, a Student- t likelihood can be used instead,

$$\mathcal{L}_t = \frac{\nu + 1}{2} \log \left(1 + \frac{\|\mathbf{d} - \mathbf{d}_{GT}\|_2^2}{\nu \sigma^2} \right) + \log \sigma, \quad (4)$$

where ν is the number of degrees of freedom. While the t-Student formulation allows to change the behavior to outliers, the gaussian one is more straightforward to interpret and thus was chosen for this work. Note that the uncertainty scale σ is constrained to lie within the interval $[10^{-4}, 10]$ for numerical stability. As σ corresponds to a per-component uncertainty, we define a vector uncertainty as

$$\sigma_{vec} = \sqrt{\sigma_u^2 + \sigma_v^2}, \quad (5)$$

where σ_u and σ_v denote the predicted per-component scales. In addition to these data fidelity terms, regularization is introduced to enforce spatial consistency of the displacement gradients. Following Manickathan et al. (2023); Mucignat et al. (2023), two penalties based on the Jacobian $\mathbf{J}$ of $\mathbf{d}$ are considered. The first enforces consistency of local gradient changes,

$$\mathcal{L}_{Jac} = |\det(\mathbf{J}) - \det(\mathbf{J}_{GT})|, \quad (6)$$

while the second constrains the overall deformation magnitude,

$$\mathcal{L}_{Fro} = \|\mathbf{J}\|_F - \|\mathbf{J}_{GT}\|_F. \quad (7)$$

Note that these terms are weighted with coefficients of 0.1 and 0.01, and the total loss at pyramid level i is defined as

$$\mathcal{L}^{(i)} = \lambda_1 \mathcal{L}_d^{(i)} + \lambda_2 \mathcal{L}_{Jac}^{(i)} + \lambda_3 \mathcal{L}_{Fro}^{(i)}, \quad (8)$$

where $\mathcal{L}_d$ denotes the probabilistic (NLL) data term. The overall objective is obtained by summing over all pyramid levels, while invalid regions (e.g., solid boundaries) are excluded using a binary mask.

We train the network gradually transitioning from a deterministic Charbonnier loss and transition to the NLL formulation in 5 epochs. This strategy stabilizes optimization and prevents the model

from exploiting the probabilistic formulation by inflating uncertainty values to artificially reduce the loss before accurate flow estimates can be obtained. This delayed activation encourages the network to first learn meaningful motion representations and subsequently calibrate uncertainty in a physically consistent manner. Similar training strategies are commonly employed in probabilistic regression to avoid degenerate solutions and improve uncertainty calibration.

We train P-LIMA by means of supervised learning on a dataset of 60k samples with a spatial resolution of 256×256 pixels, generated using the kinematic training strategy proposed in Manickathan et al. (2022). Training is stopped after 250 epochs. We augment the training set by implementing random mirroring, intensity rescaling, Poisson noise, random occlusions, and random backgrounds. We use the Adam optimizer and initialize the network weights sampling from a Kaiming distribution.

3. Results

The performance of P-LIMA is assessed by analyzing three complementary aspects: (i) the accuracy of the reconstructed displacement fields relative to the deterministic baseline, (ii) the spatial structure of the predicted uncertainty fields, and (iii) the statistical calibration of the uncertainty estimates. Results are reported for DNS-based synthetic data, experimental PIV measurements, and degraded benchmark cases from the PIV Challenge.

3.1. DNS test

We benchmark the models using publicly available synthetic image dataset Carlier (2005) generated from DNS snapshot of isotropic turbulence. We select one sample of the dataset and report the reconstructed displacement maps and the related error distributions for all models in Figure 2.

Note that we define the error ϵ as

$$\epsilon = \sqrt{(u - u_{GT})^2 + (v - v_{GT})^2} \quad (9)$$

and compare it with the reconstructed uncertainty σ . It is noteworthy that LIMA and P-LIMA yield ground-truth error distributions that are comparable to those obtained with WIDIM. A further reduction of the error would be possible by introducing an additional flow reconstruction step at the image level, albeit at increased computational cost. Moreover, the uncertainty estimate σ provided by P-LIMA shows a mild correspondence with the ground-truth error, whereas the one obtained from WIDIM appears more spatially unstructured and, in some regions, much lower than the ground truth error. This indicates that P-LIMA is better able to relate uncertainty levels to local image conditions and displacement field features than WIDIM.




Figure 2. Top panels: contour maps of displacement magnitude obtained with , LIMA^{3,1} , P-LIMA^{3,1} , and WIDIM. Mid panels: contour maps of the endpoint error ϵ . Bottom panels: contour maps of the uncertainty σ . Data obtained for sample N°82 of DNS database (Carrier, 2005).

To assess the statistical consistency of σ we evaluate against ϵ on 100 samples of the DNS dataset, assuming independent uncertainties on the two velocity components. The resulting curves are reported in Figure 3 for P-LIMA and WIDIM. In Figure 3 lines lying on the diagonal represent perfect coverage, while if they are above diagonal they indicate under-confidence, or σ values too high compared to ϵ . Conversely, points below the diagonal indicate over-confidence, i.e. σ is too small compared to ϵ . This test allows us to prove that P-LIMA provides optimal coverage compared to WIDIM, offering a more robust estimation of the uncertainty for the considered data.




Figure 3. $\sigma - \epsilon$ curves for P-LIMA and WIDIM. evaluated on 100 samples

3.2. Experimental images

Figure 4 reports the reconstructed displacement fields and associated uncertainty values for the flow past a cylinder data, also reported in Manickathan et al. (2023). It is noteworthy that, also in this case, P-LIMA shows displacement fields which are consisted with LIMA, which does not provide any uncertainty estimate. Moreover, P-LIMA labels with higher uncertainty the flow regions where 3D motion and shear are more pronounced, allowing the experimentalist eventually optimize experimental parameters to increase confidence.

3.3. PIV Challenge 3 – Case B

We use the dataset from Case B of the third PIV Challenge to evaluate the behaviour of P-LIMA under degraded imaging conditions. This case is designed to represent challenging PIV measurements, leading to reduced correlation quality and increased uncertainty in the displacement estimation.

Figure 5–Figure 8 show the velocity fields and corresponding uncertainty maps for samples 10, 85, and 115, obtained with P-LIMA and WIDIM. Note that in this dataset, the image quality degrades from the initial samples to pathological effects at the end (sample 119). Both methods exhibit increased uncertainty levels, particularly in regions where the image quality is reduced and the flow is more complex. Overall, the uncertainty predicted by P-LIMA follows the spatial variations observed in the data, indicating that the model remains sensitive to local degradation effects.

3.4. PIV Challenge 3 – Case C

We use the dataset from Case C of the third PIV Challenge to evaluate the behavior of P-LIMA where significant differences in the seeding density appear across the image. Also in this bench-




Figure 4. Vortex shedding past a cylinder. Top panels: contour maps of displacement magnitude obtained with LIMA^{3,1}, P-LIMA^{3,1}, and WIDIM. Bottom panels: contour maps of reconstructed uncertainty ϵ_d .

mark, P-LIMA shows its capability to label with very high uncertainty values the regions where the seeding density is insufficient to reconstruct meaningful velocity vectors.

4. Conclusions

We presented P-LIMA, a lean probabilistic convolutional neural network for Particle Image Velocimetry that jointly estimates dense displacement fields and spatially resolved, data-dependent uncertainty. By extending the lightweight LIMA architecture with a heteroscedastic negative log-likelihood formulation, P-LIMA provides calibrated aleatoric uncertainty at a modest computational overhead.

Coverage and reliability analyses on synthetic validation data demonstrate that the predicted uncertainty is well calibrated and more spatially structured than classical correlation-based uncertainty estimates. Benchmark results on experimental and degraded PIV Challenge datasets show that the probabilistic extension preserves the accuracy of the deterministic model while providing additional information on reconstruction reliability, particularly in regions affected by particle loss, low signal quality, or three-dimensional flow effects.




Figure 5. PIV challenge 3 case B, snapshot $N^\circ 10$. Top row: displacement maps obtained with LIMA, P-LIMA, and WIDIM. Bottom panels: uncertainty maps.




Figure 6. PIV challenge 3 case B, snapshot $N^\circ 85$. Top row: displacement maps obtained with LIMA, P-LIMA, and WIDIM. Bottom panels: uncertainty maps. Degradation of the signal-to-noise leads to an increase of the estimated uncertainty.




Figure 7. PIV challenge 3 case B, snapshot $N^\circ 115$. Top row: displacement maps obtained with LIMA, P-LIMA, and WIDIM. Bottom panels: uncertainty maps. Corrupted signal-to-noise leads to high uncertainty and increased number of non-valid-vectors in WIDIM.




Figure 8. PIV challenge 3 case C, snapshot $N^\circ 1$. Top row: displacement maps obtained with P-LIMA and WIDIM. Bottom panels: uncertainty maps. Corrupted signal-to-noise leads to high uncertainty and increased number of non-valid-vectors in WIDIM.

With real-time to near-real-time performance and a parameter count of approximately 1 M, P-LIMA enables uncertainty-aware CNN-based PIV suitable for practical experimental workflows. Future work will extend the present framework to multiple resolution variants, allowing controlled trade-offs between spatial resolution, computational throughput, and uncertainty fidelity.

Acknowledgements

We would like to thank Stephan Kunz in assisting with the code.

References

- Carlier, J. (2005). *Images with synthetic particles and scalar based on dns of self sustained 2d turbulence*. Retrieved from <http://fluid.irisa.fr/data-eng.htm>
- Dosovitskiy, A., Fischer, P., Ilg, E., Hausser, P., Hazirbas, C., Golkov, V., . . . Brox, T. (2015). FlowNet: Learning optical flow with convolutional networks. In *Proceedings of the ieee international conference on computer vision* (pp. 2758–2766).
- Kendall, A., & Gal, Y. (2017, 10). What uncertainties do we need in bayesian deep learning for computer vision? Retrieved from <http://arxiv.org/abs/1703.04977>
- Manickathan, L., Mucignat, C., & Lunati, I. (2022). Kinematic training of convolutional neural networks for particle image velocimetry. *Measurement Science and Technology*, 33(12), 124006.
- Manickathan, L., Mucignat, C., & Lunati, I. (2023). A lightweight neural network designed for fluid velocimetry. *Experiments in Fluids*, 64(10), 161.
- Morrell, M., Hickmann, K., & Wilson, B. (2020, 12). Particle image velocimetry analysis with simultaneous uncertainty quantification using bayesian neural networks. Retrieved from <http://arxiv.org/abs/2012.00642><http://dx.doi.org/10.1088/1361-6501/abf78f> doi: 10.1088/1361-6501/abf78f
- Mucignat, C., Manickathan, L., Shah, J., Rösger, T., & Lunati, I. (2023). A lightweight convolutional neural network to reconstruct deformation in bos recordings. *Experiments in Fluids*, 64(4), 72.
- Mucignat, C., Zdybał, K., & Lunati, I. (2025, 10). Improving the performance of a lightweight convolutional neural network for particle image velocimetry through hyper-parameter and padding optimization. *Physics of Fluids*, 37. doi: 10.1063/5.0283779
- Mucignat, C., Zdybał, K., & Lunati, I. (2026). Lima-3: a minimal architecture for high-accuracy particle image velocimetry. *in preparation*, 37.

- Raffael, M., Willert, C., Scarano, F., Kähler, C. J., Wereley, S. T., & Kompenhans, J. (2018). *Particle Image Velocimetry (the Third Edition)* (3rd ed.). Berlin: Springer Berlin Heidelberg. doi: 10.1007/978-3-540-72308-0
- Sciacchitano, A., Wieneke, B., Scarano, F., & Wieneke, B. (2013). Piv uncertainty quantification by image matching. *Measurement Science and Technology*, 24, 1-20. doi: 10.1088/0957-0233/24/4/045302
- Sun, D., Yang, X., Liu, M.-Y., & Kautz, J. (2018). PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 8934–8943).
- Teed, Z., & Deng, J. (2020). Raft: Recurrent all-pairs field transforms for optical flow. In *Computer vision—eccv 2020: 16th european conference, glasgow, uk, august 23–28, 2020, proceedings, part ii* 16 (pp. 402–419).
- Wang, W., Hu, J., Ai, J., & Lee, Y. (2025, 7). On uncertainty prediction for deep-learning-based particle image velocimetry. Retrieved from <http://arxiv.org/abs/2507.20102>
- Wieneke, B. (2017). *Piv uncertainty quantification and beyond* (Doctoral dissertation). doi: <https://doi.org/10.4233/uuid:4ca8c0b8-0835-47c3-8523-12fc356768f3>