

Simultaneous Dual-Exposure HDR Imaging for Particle-Streak Velocimetry

P. Owens^{1,*}, K. Wood¹, K. Rajamanickam¹, M. K. Quinn¹

1: Department of Mechanical and Aerospace Engineering, The University of Manchester, United Kingdom

*Corresponding author: padrig.owens@manchester.ac.uk

Keywords: Particle Streak Velocimetry, Quad High Dynamic Range Imaging, Multi-Exposure Flow Imaging, Blur Kernel Estimation, Compact Flow Diagnostics.

ABSTRACT

Particle Image Velocimetry (PIV) is widely employed for flow diagnostics; however, its application in confined and optically constrained environments is hindered by the need for bulky imaging systems and synchronised high-power pulsed illumination. In this paper, a framework is presented in which a synthetic replication of the Quad High Dynamic Range (Q-HDR) acquisition feature of the Sony IMX900 sensor is reformulated for Particle Streak Velocimetry under continuous illumination. The method encodes a short-exposure particle image and a corresponding long-exposure streak within a single effective acquisition, forming a deformed double-exposure analogue without pulse-straddling. A spatially resolved Fourier-domain deconvolution approach is used to recover local blur kernels from paired sharp and blurred images, providing a representation of spatially varying particle motion. Synthetic particle images with prescribed flow fields are used to assess the method and allow comparison with ground truth and conventional PIV analysis. The recovered kernel fields reproduce the dominant flow structures and show qualitative agreement with PIV-derived velocity fields. However, performance degrades in low-velocity regions due to sub-pixel blur extents and regularisation effects within the inversion. The results indicate that multi-exposure information can be encoded and interpreted within a single image, providing a basis for compact flow diagnostics in geometries not readily accessible to conventional PIV.

1. Introduction

Particle Image Velocimetry (PIV) is widely established as a robust technique for the measurement of external flow fields; however, its applicability is fundamentally constrained by the requirement for direct optical access. In conventional configurations, the measurement system is positioned outside the flow domain, typically observing a controlled facility such as a wind tunnel, where unobstructed imaging of tracer particles is achieved through optical access windows. As a result, the extension of PIV to confined or wall-bounded flows remains non-trivial.

One approach to overcoming these restrictions involves the integration of endoscopic illumination and imaging optics, enabling access through bores of small diameter. While this configuration allows measurements within otherwise inaccessible geometries, it introduces several significant limitations. Endoscopic systems are prone to severe optical distortions, necessitating extensive calibration and post-processing, while their low transmission efficiency imposes substantially increased laser power requirements. Furthermore, the restricted aperture of endoscopic probes limits both the achievable field of view and the measurable flow volume, constraining their utility in many practical applications].

Recent advances in compact, high-resolution imaging technologies have prompted renewed interest in alternative approaches to flow diagnostics. Dedicated high-speed cameras are well-suited for PIV using continuous illumination, as it can achieve shorter inter frame times , but remain prohibitively expensive for many applications. Consequently, there has been a growing shift towards the use of non-specialised machine vision cameras intended for use in robots and industrial automation. These sensors have been adopted in a range of recent research efforts, including approaches such as TRUflow (Özer & Quinn, 2022), SPECTRA-A (Lemarechal et al., 2023), and event-based flow diagnostics (Willert & Klinner, 2022; Rusch & Rösigen, 2023; Gallego et al., 2018), which collectively demonstrate the potential of such hardware to enable more compact and accessible velocity measurement systems.

The use of such cameras offers several practical advantages, including reduced system cost, lower power requirements, and enhanced portability. Their compact form factor facilitates deployment within confined or embedded measurement regions, enabling access to geometries that are not amenable to conventional optical arrangements. In addition, their modularity and ease of integration enable a rapid experimental campaign, which is highly beneficial for field measurements in industrial environments.

However, these benefits are accompanied by several important limitations that directly impact their suitability for velocity measurement. In particular, machine vision cameras typically exhibit reduced illumination capability, lower achievable frame rates, and longer inter-frame times relative to dedicated PIV systems. Timing precision is further degraded by exposure delays, signal propagation through I/O interfaces, and system-level latency, limiting the attainable timing precision and temporal resolution. Critically, such systems lack true double-frame acquisition capability: whereas conventional PIV cameras capture two precisely timed exposures within a single acquisition cycle, enabling accurate synchronisation with pulsed illumination, machine vision cameras operate using sequential frame acquisition. This prevents reliable “pulse-straddling” and restricts the achievable inter-frame separation, precluding sub-microsecond Δt typically required for high-speed PIV measurements. Additionally, the inherent sensor readout latency in these cameras influences the timing precision and thus Δt .

These constraints motivate the exploration of alternative measurement strategies that are better

aligned with the capabilities of machine vision hardware. In this context, Particle Streak Velocimetry (PStV) offers a compelling framework, as it relaxes the requirement for precisely timed double-frame acquisition by extracting velocity information from temporally integrated particle streaks within a single exposure. This enables the utilisation of lower-cost, compact imaging systems while retaining the ability to interrogate flows within confined geometries.

Accordingly, the present work investigates the application of machine vision sensors to PStV for the measurement of internal and wall-bounded flows, with a particular focus on enabling compact, accessible, and deployable flow diagnostic systems suitable for geometrically constrained environments.

2. IMX900 Based Imaging System

The IMX900 sensor employs a backside-illuminated (BSI) CMOS architecture, in which the wiring layer is positioned behind the photodiode, increasing quantum efficiency and enabling shorter exposure times under reduced illumination. This is particularly advantageous in compact measurement setups where high power laser illumination and optical access are limited.

A small pixel pitch within a 1/3.1" optical format enables high spatial resolution while maintaining a compact footprint (42.8 mm x 29 mm x 29 mm (Basler, 2025)), facilitating integration into confined or embedded flow environments.

Global shutter operation ensures all pixels are exposed simultaneously, eliminating temporal skew and preventing motion-induced distortion, thereby preserving the integrity of particle and streak imaging.

The sensor adopts a stacked architecture, separating the pixel and processing layers, enabling high-speed operation and advanced readout modes within a compact form factor.


Programmable, sensor-level exposure control allows precise definition of integration times and is realised through a 2x2 pixel grouping structure, where subgroups can be assigned independent exposures while sharing a common readout. This enables multiple effective exposures within a single frame, forming the basis of Quad High Dynamic Range (Q-HDR) imaging and extending measurable intensity and velocity ranges.

Taken together, these features provide the sensitivity and deterministic capture required for accurate flow imaging, while enabling capabilities beyond conventional machine vision systems that underpin the Q-HDR methodology.


2.1. Quad High Dynamic Range

The Sony IMX900 sensor (Sony Semiconductor Solutions Group, 2025) is organised into 2×2 pixel unit cells, indexed from 1 to 4, with each index corresponding to a distinct exposure stage (Fig.1a). This architecture enables multiple exposure intervals to be encoded within a single frame, removing the need for sequential multi-frame acquisition. Each pixel subgroup is assigned an independent integration time, allowing simultaneous short- and long-duration exposures and thus multiple effective sampling times within a single measurement.

A key constraint of this approach is that all exposures must terminate simultaneously due to the shared readout (as visualised in Fig.1b), imposing a fixed synchronisation condition. As a result, exposure intervals are nested rather than sequential, which directly influences the interpretation of and the modelling of streak formation and motion blur.



(a) Demosaicing of raw images



(b) Example sub-image exposure breakdown

Figure 1. Quad High Dynamic Range on the Sony Pregius S IMX900.

The recorded image consists of interleaved pixel groups associated with different exposure stages, and must therefore be separated into four sub-images during reconstruction, each corresponding

to a distinct integration period. This demultiplexing reduces the effective spatial resolution and alters the spatial sampling characteristics.

Although originally intended for high dynamic range imaging of static scenes, this capability can be reframed as multi-exposure imaging in the context of flow diagnostics. Here, each pixel sub-group corresponds to a different integration time, such that motion is encoded over multiple temporal intervals within a single frame. Unlike conventional high-speed imaging, where exposure sequences are separated temporally, the Quad HDR architecture encodes multiple exposure intervals within a single frame through spatial multiplexing, enabling simultaneous acquisition of multi-timescale flow information.

While offering significant benefits, this approach is subject to some practical constraints, including a reduced exposure range when using Q-HDR as compared to global exposure, fixed synchronisation requirements, additional challenges associated with firmware-level processing and reduced spatial resolution following reconstruction. The latter of which is a consequence of each exposure stage sub-image making up quarter of the total resolution after image reconstruction due to the nature of the pixel grouping as shown in Fig.1a. The most notable pitfall however is the onboard static defect pixel correction which operates on full-frame neighbourhood information, and does not consider the pixel grouping discussed in (Fig.1). Therefore, when applied to Q-HDR images, where each pixel group has a different exposure time, corrections are applied based on intensity values that are not representative of the exposure conditions at a sub-image level, thereby introducing artefacts. This correction is typically enabled by default at the camera level, and its use in the present context was found to increase spurious noise in the reconstructed sub-images. To address this, the manufacturer-provided defect pixel map was noted, and the static defect pixel correction was disabled before images were captured. Defective pixels are instead corrected after sub-image reconstruction using a neighbourhood based scaling approach applied within each sub-image, ensuring that corrections are consistent with the corresponding exposure level. At present, it remains unclear whether these artefacts arise from fundamental constraints of the sensor architecture, or from the specific implementation of the sensor by the camera manufacturer.

2.2. Underlying Concept for Application of Q-HDR to Flow Measurement

Because the Q-HDR acquisition scheme requires overlapping exposure intervals, it is not directly compatible with conventional frame-straddling Particle Image Velocimetry, in which one exposure is assigned to each frame. This constraint instead motivates the consideration of two alternative routes for velocity estimation. The first is multi-exposure autocorrelation PIV, as applied by Özer & Quinn (2022) in the TRUflow system, in which multiple laser pulses are recorded within each frame and the resulting displacement information is recovered through autocorrelation. While this approach relaxes the frame-rate requirement of the camera and partially mitigates directional ambiguity through the ordering of consecutive exposures, it remains dependent on tightly controlled

pulsed illumination and the associated timing circuitry, which can add significantly to cost and system size. In addition, the well-known limitations of autocorrelation-based approaches, including ghost particles and particle overlap, remain present.

An alternative and potentially more suitable approach is Particle Streak Velocimetry under continuous illumination. In this case, the use of continuous rather than pulsed light sources simplifies the illumination systems and eliminates the need for high-speed pulse-driver electronics, which is attractive for compact and affordable measurement hardware. Although streak-based methods remain susceptible to overlap in regions of high seeding density, the multiple nested exposure intervals provided by Q-HDR introduce an additional temporal structure into the recorded image. Rather than observing a single integrated streak alone, the evolution of the streak across the available exposure stages may be used to assist in interpreting particle motion and improve the resolution of ambiguous or overlapping trajectories (Fig.2b). Most importantly, the shortest exposure yields a localised particle image at the termination of the streak, thereby addressing one of the principal limitations of conventional Particle Streak Velocimetry, namely directional ambiguity.

This capability closely mirrors the principle proposed by Kato & Shimizu (2000), who employed a single-frame multi-exposure method in which a short pulse was used to identify one end of the particle trajectory and a longer pulse to determine the particle displacement (Fig.Kato & Shimizu (2000)). Q-HDR enables an analogous measurement arrangement within a single acquisition by simultaneously providing both a sharp particle image and a temporally integrated streak image. In this sense, the sensor facilitates a deformed double-exposure measurement without requiring two separately timed illumination events. The principal trade-off is a reduction in effective spatial resolution, since each exposure stage is reconstructed from only one quarter of the sensor samples. Nevertheless, this compromise may be acceptable if it enables longer effective measurement intervals and improved extraction of velocity information from compact imaging hardware. Although the short exposure occurs at the end rather than at the start of the streak, this does not alter the fundamental principle, since the corresponding analysis may be applied in reverse.

The relationship between the short- and long-exposure images also suggests a useful analogy with motion deblurring. The long-exposure image may be interpreted as a motion-blurred representation of the particle field, while the short-exposure image provides a sharper reference of the same flow state. If the blur kernel required to transform the short-exposure image into the corresponding long-exposure image can be estimated, this kernel would provide a direct description of the local particle motion within each interrogation region. On this basis, the initial objective is to establish a proof of concept using only one pair of short- and long-exposure images. Demonstrating velocity recovery from this reduced image pair would provide a foundation for a more advanced methodology that exploits all four Q-HDR exposure stages to recover three successive velocity estimates at distinct values of Δt .




Figure 2. Particle Images and Timing Diagrams

3. Use of Synthetic Images for Algorithm Development

Synthetic images are employed to provide a controlled and quantitative basis for algorithm development and validation. By prescribing the flow field, they yield exact particle displacements and velocity information, enabling direct assessment of algorithm accuracy against a known ground truth. In addition, key imaging and flow parameters, including particle seeding density, noise level, displacement magnitude, and illumination conditions, can be varied independently. This allows systematic evaluation of sensitivity and robustness, while ensuring repeatability through the regeneration of identical datasets. As a result, synthetic data serve as an efficient testbed for method development and parameter tuning prior to application on experimental measurements.

However, this approach is subject to limitations. Synthetic images may not fully replicate experimental conditions, leading to a sim-to-real gap that can affect the transferability of observed performance. In particular, accurately modelling particle appearance, noise characteristics, and

optical artefacts such as aberrations, saturation, and reflections remains challenging. Furthermore, there is a risk that algorithms become implicitly tuned to the assumptions or artefacts of the specific image generator, potentially reducing generalisability.




Figure 3. Synthetic image input data generated using the modified PIVLab tool.

The synthetic image generator used in this work was based on a modified version of the PIVLab synthetic particle image generation tool (Thielicke, 2014). It was adapted to produce sequences of frames with minimal inter-frame displacement, which were subsequently combined and normalised to emulate a long-exposure streak image (Fig.3b). The final frame in each sequence was retained separately to represent the corresponding short-exposure image (Fig.3a). This process yields image pairs with a known ground-truth flow field and associated error, providing a consistent framework for algorithm development and validation prior to application to experimental data.

4. Methods - Blur Kernel Estimation

To recover the local blur kernels from paired short- and long-exposure images, a spatially resolved deconvolution strategy was developed. A local, Fourier domain deconvolution framework was implemented using PIV-style interrogation windows to spatially resolve blur characteristics across the image field. First, the sharp and blurred images were deconstructed into interrogation windows of 64×64 pixels, with 50% overlap to enable localised processing. For each window, the sub-images were embedded into a zero-padded domain sized to avoid circular convolution artefacts. This padding ensures that the subsequent Fourier operations approximate linear convolution, which is critical for accurate kernel estimation.

Deconvolution was performed in the frequency domain using a noise-aware Wiener formulation,

in which the long-exposure image, B , is assumed to arise from the convolution of a short-exposure image, I , with an unknown kernel, K , such that $B = I \otimes K$. Taking the Fourier transform converts this relationship into a multiplicative form, allowing the kernel spectrum to be estimated as:

$$F_K(\mathbf{k}) = \frac{F_B(\mathbf{k}) F_I^*(\mathbf{k})}{|F_I(\mathbf{k})|^2 + S_n(\mathbf{k})} \quad (1)$$

Where F_B and F_I denote the Fourier transforms of the blurred and reference windows, respectively, F_I^* is the complex conjugate, $|F_I|^2$ represents the signal power spectrum, and S_n is an estimate of the noise power. In practice, the power spectrum of the reference window is computed and smoothed to provide a stable estimate of signal energy, while the noise level is inferred from the low-magnitude frequency components which are assumed to be noise dominated. The inclusion of S_n acts as a regularisation term, stabilising the inversion and preventing amplification of high-frequency noise in poorly conditioned regions. The spatial domain kernel is then recovered via inverse Fourier transform of each interrogation region and recorded within a blur kernel field as shown in Fig.5.




Figure 4. Algorithm workflow for blur kernel estimation.

To further refine the kernel estimate, a frequency domain filtering step is applied in which symmetric peak pairs, corresponding to structured artefacts or oscillatory components, are identified relative to the zero frequency component of the Fourier domain and removed. The filtering step operates by identifying and removing conjugate-symmetric peak pairs in the Fourier domain, which correspond to oscillatory artefacts introduced during deconvolution. Exploiting the inherent symmetry of the Fourier transform for real valued signals, dominant spectral peaks are detected and paired via inversion about the frequency origin. Removal of these paired components suppresses




Figure 5. Full Kernel Fields Produces using Noise-Aware Deconvolution for the Synthetic Images in Fig.3

unstable oscillatory modes, analogous to an adaptive notch filtering process, thereby reducing ringing artefacts in the reconstructed kernel. The filtered spectrum is subsequently inverted, and residual noise is attenuated using guided filtering to preserve coherent spatial structure while improving overall kernel smoothness. An example of the improved signal-to-noise ratio from this filtering technique is evident when comparing the unfiltered and filtered kernel maps in Fig.5. It can be observed that the background noise is attenuated while the streaks remain sharp, thus making it easier to discern the blur streak from the background.

The proposed deconvolution framework, while effective for spatially resolved kernel estimation, is subject to several limitations that are partially mitigated through its design. The use of fixed size interrogation windows assumes locally spatial invariance of the blur kernel, which may not hold in regions with strong gradients or rapidly varying imaging conditions, potentially introducing inconsistencies between neighbouring windows. This limitation is addressed through the use of overlapping interrogation windows, which promote smoother spatial transitions and improve robustness to moderate spatial variations. Similarly, the choice of window size introduces an inherent trade-off between spatial resolution and spectral robustness; however, the selected overlap and window configuration provide a practical balance that maintains both localisation and stability.

Rectangular windowing and padding can introduce spectral leakage and boundary artefacts that bias the Fourier representation. To mitigate this, each window is embedded within an enlarged padding domain, ensures that the deconvolution more closely approximates linear convolution and reducing edge induced distortions. Within the inversion itself, the noise-aware Wiener formulation relies on estimating noise power from a low-magnitude frequency components, implicitly assuming that these regions are noise-dominated and spectrally uniform. While deviations from these assumptions can lead to inaccurate regularisation, the approach is stabilised by smoothing

the signal power spectrum and incorporating the estimated noise term, which suppresses amplification of poorly conditioned frequency components and improves robustness in low-energy regions. Nonetheless, the method remains sensitive to the quality and alignment of the reference image, although the local, windowed formulation helps to limit the spatial propagation of such errors.

Further limitations arise in the frequency domain filtering stage, where symmetric peak pairs are removed using a tolerance based heuristic. Although this strategy risks misclassifying physically meaningful spectral content or failing to remove artefacts, it provides an efficient means of suppressing structured oscillatory components without requiring explicit modelling. This is complemented by a guided filtering step in the spatial domain, which mitigates residual noise while preserving dominant kernel structure, thereby balancing noise suppression with feature retention. Finally, while the computational cost of performing localised Fourier operations across a dense interrogation grid is significant, this is alleviated through parallelisation, enabling dense spatial sampling and making the approach practically tractable for large-scale image analysis.

4.1. Algorithm Performance

The method, at its current stage of development, produces only qualitative outputs, which limits how quantitatively it can be evaluated. The analysis therefore adopts a more qualitative and holistic perspective to examine the performance and behaviour of the approach. To provide a reference point, the results are compared against an established solution. PIV analysis is performed on image pairs derived from the same source data and temporal separation, Δt . The resulting velocity field is superimposed on the corresponding blur kernel field, demonstrating good agreement between the flow structures recovered via deconvolution and those obtained using PIV (Fig.6). Both approaches recover the characteristic counter-rotating vortex structures, with streak or vector length reflecting local velocity magnitude and orientation indicating flow direction, showing strong overall agreement between the two methods at areas of high velocity.

However, the limitations of this visualisation technique become apparent at low relative velocities, such as those observed near the corners of Fig.6. In these regions, differences in velocity are difficult to distinguish due to the finite spatial resolution of the streaks. At sufficiently small displacements, motion is confined to a sub-pixel scale; that is, blur extents in the sub-pixel scale remain indistinguishable within the same central pixel neighbourhood. It's possible that this is also a consequence of the regularisation term in the deconvolution formulation. When the blur magnitude is weak, the noise related spectral term may become dominant relative to the informative kernel component, causing the estimator to behave similarly to a heavily regularised Wiener filter. As a result, kernels corresponding to small velocities may have been driven towards a common, spatially smooth central structure, reducing sensitivity to subtle variations. This highlights the need to investigate a more adaptive or less aggressive noise modelling strategy within the




Figure 6. Vector field obtained from PIV analysis of an image pair generated using the same underlying velocity field as the synthetic streak images, overlaid on the corresponding blur kernel map derived from the deconvolution method.

deconvolution formulation, that is able to reduce the risk of over-smoothing.

From a computational perspective, the method demonstrates moderate performance. For a representative image of 1032×776 pixels, using a window size of 64×64 pixels with 50% overlap, a full field of 31×23 kernels is recovered in an average time of 7.7 seconds, based on 10 runs using synthetic data. This provides a baseline for performance assessment, although direct comparison with established PIV implementations are at date unsuitable due to the large differences in their data output, nevertheless, once the algorithm has been developed further, a comprehensive analysis will be undertaken.

In its current form, the technique is best suited to flows that produce sufficiently long particle streaks within each interrogation window. Its performance degrades for sub-pixel displacements, limiting its ability to resolve low velocity regions, and small variations in velocity between windows. This constraint can be partially mitigated through experimental design, for example by selecting exposure times that yield longer streaks in low-speed regions, or by leveraging the camera's Q-HDR capability to have longer exposure times within regions of low velocity and shorter

exposure times in areas of higher velocity, however, this comes with its own complications. Consequently, the present implementation should be regarded as a purely qualitative tool, capable of visualising key flow features and capturing relative velocity distributions and flow direction, but not yet providing quantitatively resolved velocity measurements at sub-pixel scales.

5. Future Work

In the context of the deconvolution tool, future work will focus on improving the resolution of the recovered blur kernels and enabling discrimination between kernels that are visually indistinguishable. Sub-pixel analysis is therefore required to resolve small variations in the extent of blur that are otherwise indistinguishable within the native pixel grid. In parallel, low-velocity regions can be referenced against a ground-truth zero-velocity kernel acquired prior to experimentation, providing a baseline against which subtle deviations associated with weak motion may be identified.

Further work is also required to establish a reliable method for extracting velocity magnitude directly from the recovered streaks. While several techniques have been explored, no consistent approach has yet been identified that robustly relates streak length to the underlying velocity field across varying flow conditions. Persistent limitations arise from the low signal-to-noise ratio at the extremities of particle streaks, which inhibits consistent identification of their termination points. Although the application of the filtering strategies outlined above enhances streak visibility, this challenge is not fully resolved. This behaviour also further motivates the need for sub-pixel analysis in low-velocity regions, where streak structures do not exhibit the same well-defined tails observed at higher velocities. Consequently, both streak extent and orientation become increasingly ambiguous, reducing confidence in the inferred flow direction within the interrogation window.

Furthermore, an experimental campaign is proposed to evaluate the performance of the method under practical conditions. This will involve comparing a Q-HDR imaging system with a conventional whole-frame exposure camera in canonical flow configurations, using low-cost continuous illumination. The objective is to assess the applicability of the technique beyond synthetic datasets, while simultaneously identifying the mechanisms responsible for failure under real-world conditions and testing potential mitigation strategies. A side-by-side camera mounting arrangement will be used to acquire closely matched image pairs from two camera systems, ensuring consistent spatial resolution and optical characteristics. One camera will utilise the Sony IMX900 sensor within a Q-HDR imaging configuration, while a second camera will operate as a conventional, exposure-controlled PIV system. By synchronising the acquisition such that equivalent temporal separations are achieved, the conventional camera data can be used to generate benchmark velocity fields using validated PIV software. The acquisition timing for the conventional system is illustrated in Fig.7. This enables direct comparison between the deconvolution-based approach,




Figure 7. Streak and timing diagrams for each approach in the proposed experimental campaign demonstrating how the timing will work.

established measurement techniques, and analytical solutions for canonical flows.

6. Conclusions

The limited applicability of conventional PIV to confined and optically restricted geometries motivates the exploration of measurement strategies that are compatible with compact, lower-cost imaging hardware. In this context, the Quad High Dynamic Range acquisition scheme of the IMX900 sensor provides a fundamentally different mechanism for capturing temporal information, in which multiple exposure intervals are encoded within a single frame rather than resolved through sequential acquisition. By synthetically replicating this acquisition structure, it has been shown that the resulting combination of localised particle images and temporally integrated streaks

can be interpreted as a deformed double-exposure measurement under continuous illumination, without reliance on pulse-straddling or tightly synchronised timing. When coupled with a spatially resolved Fourier-domain deconvolution framework, this representation enables recovery of blur kernels that reflect spatial variations in particle motion, thereby providing a route for extracting flow structure from single-frame, multi-exposure data. Within the scope of the synthetic datasets considered here, the recovered kernel fields preserve the dominant flow features and remain consistent with corresponding PIV-derived velocity fields, indicating that the encoded exposure structure contains usable kinematic information.

The method nevertheless remains at a proof-of-concept stage and is currently limited to qualitative interpretation of the recovered kernel field. Sensitivity degrades in low-velocity regions, where sub-pixel blur extents are not cleanly distinguishable and the regularised deconvolution increasingly suppresses weak kernel differences. Quantitative recovery of velocity magnitude has not yet been established, in part because streak extent remains difficult to define robustly at low signal-to-noise ratio and because the recovered kernel is affected by both reference-image quality and the assumptions embedded in the deconvolution model. In addition, the current assessment is based on synthetic imagery, and experimental validation is still required to determine how reliably the method transfers to measured data. Even with these restrictions, the results indicate that Q-HDR-style acquisition can be repurposed as a mechanism for embedding multi-timescale motion information within a single image, in a manner that is inherently compatible with compact machine vision systems. This establishes a viable pathway towards compact machine-vision-camera-based flow diagnostic systems that can broaden the applicability of velocity measurement to geometries that are not practically accessible using conventional PIV.

References

- Basler. (2025). *Basler ace 2 a2A2048-114umBAS | NIR Camera | Basler AG — baslerweb.com*. <https://www.baslerweb.com/en/shop/a2a2048-114umbas/>. ([Accessed 03-06-2026])
- Gallego, G., Rebecq, H., & Scaramuzza, D. (2018). A unifying contrast maximization framework for event cameras, with applications to motion, depth, and optical flow estimation. In *2018 IEEE/CVF conference on computer vision and pattern recognition (cvpr)* (p. 3867-3876). New York: IEEE.
- Kato, F., & Shimizu, I. (2000). Optical processing of particle tracking velocimetry under deformed double exposure. *Measurement science and technology*, 11(6), 646-654.
- Lemarechal, J., Dimond, B. D., Barth, H. P., Hilfer, M., & Klein, C. (2023). Miniaturization and model-integration of the optical measurement system for temperature-sensitive paint investigations. *Sensors*, 23(16).

- Özer, Ö., & Quinn, M. K. (2022). Novel particle image velocimetry methods for visualization of thrust reverser's flow interactions. In *Aiaa aviation 2022 forum*. Reston: Amer Inst Aeronautics and Astronautics.
- Rusch, A., & Rösgen, T. (2023). Trackaer: real-time event-based quantitative flow visualization. *Experiments in fluids*, 64(8).
- Sony Semiconductor Solutions Group. (2025). *Compact, high resolution, high nir sensitivity 1/3.1" 3.2mp global shutter imx900*. Retrieved 16-12-2025, from <https://www.sony-semicon.com/en/products/is/industry/gs/imx900.html>
- Thielicke, W. (2014). *The flapping flight of birds: Analysis and application* (Doctoral dissertation, University of Groningen). doi: 10.13140/RG.2.2.18656.94728
- Willert, C. E., & Klinner, J. (2022). Event-based imaging velocimetry: an assessment of event-based cameras for the measurement of fluid flows. *Experiments in fluids*, 63(6).