Velocity-field estimation techniques for on-disk solar coronal bright fronts (CBFs) and shock-wave propagation

Oleg Stepanyuk[1], Kamen Kozarev[1]
Institute of Astronomy and National Astronomical Observatory, Bulgarian Academy of Sciences, Sofia 1784, Bulgaria
stepanyuk.oleg@aol.com

(Submitted on 16.07.2025; Accepted on 17.08.2026)

Abstract. Coronal mass ejections (CMEs) are large expulsions of plasma and magnetic fields from the solar corona into the heliosphere. Coronal shock waves are driven by CMEs, and accurate Solar Energetic Particles (SEP) forecasting requires knowing when and where shocks form, how fast they propagate, and how they interact with the coronal magnetic field. By measuring the speed, geometry and temporal evolution of coronal bright fronts (CBFs), one can make more realistic predictions of SEP onset and intensity at Earth, thus velocity-field estimation techniques provide these critical parameters. High-resolution velocity field maps can provide even deeper insight. Instead of yielding only an overall speed of the front, techniques like FLCT and optical flow can show the spatial distribution of velocities across the entire wave dome or front. These analyses reveal which parts of the front are moving fastest and how the front might be deforming as it expands. Here we discuss a set of practical parameters for Lucas–Kanade, Horn–Schunck, FLCT, and Farnebäck algorithms, arrived at by combining analytic reasoning with empirical measurements on a single event—and demonstrate the potential of these techniques for on-disk shock-wave propagation during the X1.5-class event that took place 10 May 2022, processed by the wavelet-based Wavetrack technique. A better understanding of how to fine-tune these parameters could improve the precision of algorithms running in pipelines, and is especially valuable where per-event hand-tuning is impractical.

Key words: Coronal mass ejections, shock waves, velocity fields, solar energetic particles

Introduction

The motivation for studying coronal shock waves is tightly linked to space weather: shocks driven by CMEs are a primary source of solar energetic particles (SEPs), high-energy protons and ions that can pose radiation hazards to spacecraft and astronauts. Reliable SEP forecasting requires understanding when and where particles are accelerated in a solar eruption and how they travel to Earth. Velocity field estimation techniques contribute to this in several ways. First, identifying a CBF as a shock and measuring its speed, geometry, and evolution provides the initial conditions for particle acceleration models. The efficiency of particle acceleration at a shock depends on the shock's speed and angle relative to the magnetic field (quasi-parallel vs quasi-perpendicular), among other factors [Kozarev et al., 2017]. Second, understanding the propagation of SEPs through the corona and interplanetary space benefits from knowing the shock propagation history. SEPs released by a moving shock will travel along the interplanetary magnetic field to reach Earth or other spacecraft. The timing of SEP onset and their intensity can be better predicted if we know when the shock formed and how fast it traveled through the corona. Velocity maps, combined with coronagraph data in the higher corona (e.g., from SOHO/LASCO or STEREO), allow the construction of a continuous height–time profile of the front. With higher cadence data from STEREO and SDO, two-phase kinematics can be observed: an initial fast phase (several hundred km/s) that often decelerates over time [Long

et al., 2011]. High-resolution velocity field maps have provided even deeper insight. Instead of yielding only an overall speed of the front, techniques like FLCT and optical flow can show the spatial distribution of velocities across the entire wave dome or front. These analyses reveal, for instance, which parts of the front are moving fastest and how the front might be deforming as it expands. Often, the flanks of the coronal wave (the edges sweeping to the north/south of the eruption center) are observed to move faster than the wave segment directly above the eruption or on the disk.

Image-by-image segmentation, traditionally performed using algorithmic methods and more recently explored with data-driven approaches, often considers only the spatial aspects of image features, entirely overlooking their temporal evolution. Beyond their value in enhancing the qualitative understanding of physical processes, velocity field estimation techniques can further help differentiate genuine details from spurious features as these methods operate within a space-time context rather than being limited to purely spatial processing. Simply put, capturing feature movement helps to distinguish moving details from static ones, even if they have similar intensities or shapes. [Stepanyuk and Kozarev, 2025] The velocity field of a global EUV wave observed by SDO/AIA on 2021-05-07, derived using local correlation tracking, lets us measure instantaneous velocities of different parts of the CBF and how they change over time. In the 2021-05-07 event shown, the EUV wave's flanks reached $\sim$500 km/s early on, whereas later frames showed those flanks slowing down and the northern part of the front dropping below 200 km/s [Stepanyuk et al., 2024]. This kind of information is critical for understanding the wave's evolution and energy dissipation. Velocity fields can be used to derive physical parameters of the corona: combining an observed wave speed with assumptions of fast-mode wave behavior yields an estimate of the local magnetosonic speed, and hence the coronal magnetic field strength. In the 2010 event [Long et al., 2011], the range of observed CBF speeds (380–460 km/s) implied a quiet coronal magnetic field of order 1–2 Gauss in the low corona. Thus, velocity estimation techniques are not just about kinematics – they also provide a diagnostic tool for the coronal environment (plasma density and magnetism) through which the wave propagates.

In solar physics, the FLCT method is often considered a first-choice technique for determining horizontal flows, particularly in active regions. The method calculates local shifts between two images in the Fourier domain, which is particularly useful for tracking solar granulation and magnetic fields. By computing cross-correlations in Fourier space, it can be efficient and able to identify local displacements effectively [Fisher and Welsch, 2008].

[Byrne et al., 2013] addresses the FLCT sampling-cadence problem within the framework of numerical differentiation, emphasizing that naive implementations of basic difference formulas are inadequate in scenarios with infrequent or irregular data sampling. The fundamental difficulty arises from errors introduced in differentiation processes due to non-uniform or sparse sampling intervals. Consequently, more sophisticated methods, such as higher-order polynomial fits or adaptive numerical schemes, become necessary in these situations. This challenge prompted our exploration of alternative velocity field estimation techniques [Stepanyuk and Kozarev, 2024], where we mainly used Horn–Schunck [Horn and Schunck, 1981] and Lucas–Kanade [Lucas and Kanade,

1981] velocity estimation methods as an alternative for FLCT. Here we extend this research by performing a comparative analysis of velocity field estimation techniques for on-disk wave propagation.

1 Method

In the context of estimating velocity fields from temporally separated frames, Fourier Local Correlation Tracking (FLCT) and classical optical-flow algorithms such as Lucas–Kanade, Horn–Schunck an Farnebäckprovide distinct methodological approaches, each with tunable parameters that significantly influence their sensitivity to different motion scales and structures. This section first introduces the formal description of each algorithm (Sec. 1). The remainder of the section then works through a set of parameters for detecting moderately sized circular features, combining *analytic* reasoning with *empirical* statistics measured directly on the AIA 193 Å running-difference frames of the X1.5-class event of 10 May 2022 (Sec. 2). This is presented as a worked example of how such reasoning can give an algorithm a sensible starting point rather than being tuned blind; Sec. 3 then demonstrates the resulting algorithm performance on the same event. The 10 May 2022 event therefore serves as the running case study throughout: it provides the calibration data for the parameters and the demonstration case for the resulting algorithms.

Algorithms

Fourier Local Correlation Tracking (FLCT). FLCT estimates the velocity field by calculating the local cross-correlation of image patches in the Fourier domain. It is designed to find the displacement that best aligns local regions between two images. The core idea is to find the shift $(\delta x, \delta y)$ that maximizes the local cross-correlation function $C(\delta x, \delta y)$ between image patches. Let $f_1(x, y)$ and $f_2(x, y)$ be two consecutive image frames. The local correlation function is computed as

$$C(\delta x, \delta y) = \mathcal{F}^{-1}\left\{\mathcal{F}[G(x,y)f_1(x,y)] \cdot \overline{\mathcal{F}[G(x,y)f_2(x,y)]}\right\}(\delta x, \delta y), \quad (1)$$

where $G(x, y)$ is a Gaussian window, $\mathcal{F}$ denotes the Fourier transform, and the overline denotes the complex conjugate. The velocity vector is then estimated from the peak of the correlation function in each windowed region. The technique therefore exposes three control parameters: the correlation patch size W_{FLCT}, the Gaussian weighting width σ_G, and an acceptance threshold τ on the peak correlation that filters out weak detections.

The Horn–Schunck method. The Horn–Schunck method is a classical approach to estimating optical flow. It is based on the assumption that the apparent velocity of brightness patterns in an image sequence varies smoothly almost everywhere. The key assumptions are brightness constancy (the intensity of a moving object remains constant between frames) and a smoothness constraint (the optical flow field varies smoothly over the image). Let $I(x, y, t)$

represent the image intensity at position (x, y) and time t. Assuming small motion between frames, the brightness constancy constraint reads

$$\frac{\partial I}{\partial x}u + \frac{\partial I}{\partial y}v + \frac{\partial I}{\partial t} = 0, \tag{2}$$

where u and v are the horizontal and vertical components of the optical flow. To make the problem well-posed, Horn and Schunck added a regularization term that enforces smoothness:

$$E(u, v) = \iint \left[(I_x u + I_y v + I_t)^2 + \alpha^2 \left(\|\nabla u\|^2 + \|\nabla v\|^2 \right) \right] dx\, dy, \tag{3}$$

where α is a regularization parameter controlling smoothness and ∇u, ∇v are the gradients of the flow field. The Euler–Lagrange equations are used to minimize the energy functional; this leads to an iterative Jacobi update scheme where each flow component is refined using local averages and gradient information. The two control parameters that the user must select are α and the number of Jacobi iterations N_{it}.

Lucas–Kanade optical flow. The Lucas–Kanade method assumes that the optical flow is constant within a small window of side W. It solves the optical flow equation

$$I_x u + I_y v + I_t = 0 \tag{4}$$

by performing a least-squares fit in this local region. Minimizing the squared residual over a window of pixels leads to the normal equations

$$\begin{bmatrix} \sum I_x^2 & \sum I_x I_y \\ \sum I_x I_y & \sum I_y^2 \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = - \begin{bmatrix} \sum I_x I_t \\ \sum I_y I_t \end{bmatrix}, \tag{5}$$

where the sums run over all pixels in the local window. The system is well-conditioned only where the structure tensor on the left-hand side has both eigenvalues above some threshold $\lambda_{\min}$. The spatial derivatives I_x, I_y are evaluated with a Sobel kernel of size k. The three control parameters are therefore $(W, k, \lambda_{\min})$.

Farnebäck: polynomial expansion of neighborhoods. The Farnebäck optical-flow algorithm [Farnebäck, 2002] is a dense optical-flow method that estimates motion between two images by modeling pixel neighborhoods as polynomial expansions. Unlike sparse methods (e.g., Lucas–Kanade), which track specific features, Farnebäck's algorithm computes flow vectors for every pixel, offering a full motion field. The core idea is to approximate the local neighborhood of each pixel using a quadratic polynomial, allowing a smooth representation of image intensity variations. Specifically, the algorithm models an image patch around a pixel $\mathbf{x}$ as

$$f(\mathbf{x}) \approx \mathbf{x}^T \mathbf{A} \mathbf{x} + \mathbf{b}^T \mathbf{x} + c, \tag{6}$$

where $\mathbf{A}$ is a symmetric matrix representing the second-order terms, $\mathbf{b}$ is a vector for the linear terms, and c is a constant. This polynomial expansion is computed for both frames; the motion is estimated by comparing the polynomial models. The method captures both large and small displacements through a multi-level pyramid. The final output is a dense flow field $\mathbf{v}(x, y) = (u(x, y), v(x, y))$. The user-selectable parameters are the polynomial neighborhood size n (`poly_n`), the Gaussian weighting σ used to compute the polynomial coefficients (`poly_sigma`), the averaging window W_F (`winsize`), the pyramid scale s (`pyr_scale`) and the number of refinement iterations.

2 Parameter fitting: analytic principles and AIA-based calibration

The parameter values below are calibrated on the X1.5-class flare that began at 13:50 UT on 10 May 2022, originating from NOAA active region AR 13006, located near the center of the solar disk as seen from Earth. The event triggered significant ionospheric disturbances, evident in variations of Total Electron Content (TEC) measured using GNSS data [Lòpez-Urias et al., 2023], and was the subject of a detailed shock-wave analysis by [Rigney et al., 2024], who tracked the propagation of a large CBF through a coronal channel in the low solar atmosphere and correlated its movement with the production of radio bursts. Beyond its scientific interest, the event was chosen for its processing complexity, which makes it valuable for fine-tuning algorithm setups and parameters. Throughout this section we use ten Level-1 AIA 193 Å frames from 13:48 to 13:57 UT to measure the empirical inputs that drive the parameter fitting.

Rather than a rigorous derivation, what follows is a worked example: for each algorithm we combine a simple analytic (scale-matching) argument with a handful of quantities measured directly on the AIA running-difference frames of the 10 May 2022 event, to arrive at a practical starting point for the four algorithms' control parameters. These relations give parameter estimates rather than unique optimal values, and the point of walking through them is methodological – to show that this kind of reasoning can save an algorithm from being tuned blind, and give it something sensible to start from – rather than to establish a definitive configuration for every dataset. We summarize the empirical inputs measured on AIA first, then walk through each algorithm in turn.

2.1 Empirical inputs measured from AIA 193 Å

Running-difference cubes $\mathrm{RD}[t] = I[t+1] - I[t]$ are formed in raw $\mathrm{DN\,s^{-1}}$ (each frame normalized by FITS EXPTIME for photometric consistency, which stores the actual shutter open time of that frame (in seconds)). The eruption centre is identified as the brightest RD pixel near AR 13006; a quiet-corona patch with the lowest RD standard deviation provides a noise reference. The strongest rim azimuthal sector (southern flank) yields the radial profile of the wave rim. Figure 1 illustrates the three measurements that drive the subsequent parameter fitting.

These measurements give: an AIA 193 Å plate scale $p = 435\,\mathrm{km\,px^{-1}}$; an AIA cadence $\Delta t = 12\,\mathrm{s}$; a wave-rim FWHM (southern flank) $w = 11\,\mathrm{px}$; a

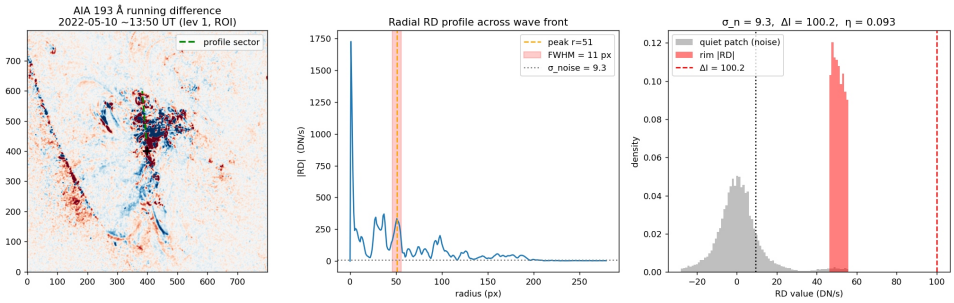

Fig. 1. Empirical statistics measured on AIA 193 Å running difference frames of the 10 May 2022 X1.5 event (∼13:50 UT). *Left:* running-difference ROI with the strong-rim azimuthal sector (dashed green) along which the radial profile is sampled. *Middle:* radial profile of |RD| across the wave front; the rim FWHM is $w = 11\,\mathrm{px}$ (orange band), well inside the 10–30 px target band. The horizontal dotted line marks the quiet-corona noise floor $\sigma_n = 9.3\,\mathrm{DN\,s^{-1}}$ ($1.4826 \times \mathrm{MAD}$). *Right:* histograms of |RD| for the quiet patch (gray) and the rim pixels (red); the rim contrast $\Delta I = 100.2\,\mathrm{DN\,s^{-1}}$ (75th percentile of rim |RD|) and the dimensionless noise/signal ratio $\eta = \sigma_n/\Delta I = 0.093$ follow directly.

rim contrast (75th percentile of rim |RD|) $\Delta I \approx 100\,\mathrm{DN\,s^{-1}}$; a quiet-corona noise level ($1.4826 \times \mathrm{MAD}$) $\sigma_n \approx 9\,\mathrm{DN\,s^{-1}}$, giving a dimensionless noise-to-signal ratio $\eta = \sigma_n/\Delta I \approx 0.09$; and a rim gradient magnitude (75th percentile) $\|\nabla I\| \approx 337\,\mathrm{DN\,s^{-1}\,px^{-1}}$. Combined with the wave-speed range of 500–1000 $\mathrm{km\,s^{-1}}$ reported for this event by [Rigney et al., 2024], these imply an inter-frame displacement budget of $\Delta = v\Delta t/p \approx 13.8$–$27.6$ px.

The measured rim FWHM $w \approx 11$ px falls inside the 10–30 px target band declared above, which is at least internally consistent with that choice. Each of the recommended values below draws on a subset of $\{w, \Delta I, \sigma_n, \eta, \|\nabla I\|, \Delta\}$ together with a simple analytic scale-matching relation.

2.2 Parameter choices: a worked example

Scale-matching principles. Each of the parameter choices below reduces to matching one of two algorithm scales to one of two data scales, as an example of the kind of reasoning that can give a starting point rather than a blind guess:

1. **Spatial scale.** The algorithm's intrinsic spatial parameter (window, kernel, σ, smoothing length) must bracket the rim half-width $w/2$ (to resolve the edge) and stay below the feature radius $r = L/2$ (to remain coherent across the feature).
2. **Displacement scale.** Any parameter that controls the maximum trackable displacement (LK pyramid level, Farnebäck winsize × pyr_scale, FLCT window) must absorb $\Delta = v\Delta t/p$, the per-frame pixel shift of the wave.

Lucas–Kanade: window, kernel, eigenvalue threshold

Window size W. The LK aperture must hold one edge of the rim while keeping $\vec{v}$ approximately constant inside the window. We use $W \gtrsim w/2$ as a practical criterion for resolving the observed rim, while $W \lesssim r$ is used as a practical criterion for retaining local flow coherence across the feature. With $w = 11$ px, the first criterion gives $W \gtrsim 5.5$ px. We therefore select the smallest odd window satisfying this condition, $W = 7$.

Sobel kernel k. The Sobel kernel must sample the rim profile without blurring it across more than one rim width. Setting $k \approx w/2$, with $w = 11\,\mathrm{px}$, and rounding to the nearest odd integer gives $k = 5$.

Eigenvalue threshold $\lambda_{\min}$. For the normalized image $I_{\mathrm{norm}} = I/S$, the characteristic gradient across a rim of width w and intensity contrast ΔI is estimated as

$$|\nabla I_{\mathrm{norm}}| \sim \frac{\Delta I/S}{w}. \tag{7}$$

This width-averaged estimate is a deliberately conservative proxy for the local gradient: it spreads the total rim contrast ΔI over the full rim width w, rather than using the sharper pointwise gradient $\|\nabla I\| \approx 337$ DN s^{-1} px^{-1} measured directly above. The conservative, width-averaged proxy yields a smaller, safer threshold that is less sensitive to the single steepest pixel along the rim and reduces the risk of rejecting valid, moderately sloped windows. The corresponding characteristic scale of the structure-tensor eigenvalues is therefore

$$\lambda_{\mathrm{char}} \sim W^2 \left(\frac{\Delta I/S}{w}\right)^2. \tag{8}$$

For $W = 7$, $\Delta I = 100$ DN s^{-1}, $w = 11$ px, and $S = 10^4$ DN s^{-1}, Eq. (8) gives

$$\lambda_{\mathrm{char}} \sim 7^2 \left(\frac{100/10^4}{11}\right)^2 \simeq 4.05 \times 10^{-5}. \tag{9}$$

We therefore use $\lambda_{\min} = 5 \times 10^{-5}$ as a practical threshold for rejecting poorly conditioned Lucas–Kanade windows. The estimate in Eq. (8) is a characteristic scale for threshold selection rather than an exact prediction of the smaller structure-tensor eigenvalue.

Farnebäck: polynomial size, smoothing, pyramid

Polynomial neighborhood n. The polynomial model in Eq. (6) must span at least the rim half-width, requiring $n - 1 \gtrsim w/2$. With $w = 11\,\mathrm{px}$ the smallest odd integer satisfying this rule is $n = 7$.

Polynomial smoothing σ. We adopt the following practical relation between the polynomial neighborhood size and its associated Gaussian smoothing,

$$\sigma \simeq 0.3(n-1) - 0.8, \tag{10}$$

which gives $\sigma = 1.0$ for $n = 7$. We therefore use $n = 7$ and $\sigma = 1.0$ for the polynomial expansion. This relation is used here as a practical parameter-selection rule rather than as a formal requirement of the Farnebäck method, and it involves no AIA quantity.

Averaging window W_F. The averaging window must encompass at least one rim width and, after pyramid downsampling, must be able to absorb the maximum displacement $\Delta_{\max}$. With $w = 11\,\mathrm{px}$ and $\Delta_{\max} \approx 27.6\,\mathrm{px}$, an odd window of $W_F = 15$ satisfies both constraints at the base level.

Pyramid scale s. The pyramid must satisfy $W_F\, s^{-L} \gtrsim \Delta_{\max}$ over L levels. With $W_F = 15$ and $\Delta_{\max} \approx 27.6\,\mathrm{px}$, a single coarse level $L = 1$ at $s = 0.3$ already gives an effective window of $\sim 50\,\mathrm{px}$, which as a scale-matching estimate appears large enough to cover the displacement budget.

Iterations. Gauss–Seidel updates of the displacement field converge to within $\sim 1\%$ in ≈ 10 sweeps for the structure-tensor conditioning typical of our problem. No AIA quantity enters this estimate.

FLCT: window, weighting, threshold

Correlation window W_{FLCT}. The window must cover the feature on its largest scale; for a Gaussian profile of FWHM equal to the feature size $L_{\max}$ one requires $W \gtrsim 1.5\, L_{\max}/2.355$. For $L_{\max} \approx 30\,\mathrm{px}$ (the upper bound of the target range), this gives $W \gtrsim 19.1$; the smallest odd window satisfying this bound is $W = 21$. This bound is set by the chosen target scale and does not consume an AIA-measured quantity.

Gaussian weighting. The Gaussian apodization $G(x, y)$ in Eq. (1) is mandatory to suppress wrap-around artefacts in the FFT-based correlation; without it spectral leakage spreads the correlation peak and biases the displacement.

Correlation threshold τ. For Gaussian noise of relative amplitude $\eta = \sigma_n/\Delta I$, the expected correlation peak is $\langle C \rangle \approx 1/(1+\eta^2)$. With $\eta = 0.09$ this predicts $\tau \approx 0.99$ for AIA alone. The value $\tau = 0.3$ adopted here is deliberately set instead by the *worst-instrument floor* in our wider processing pipeline (LASCO C2/C3, where $\eta \sim 1$), so that the same threshold can be applied across instruments without prejudicing AIA detections.

Horn–Schunck: regularization, iterations

Regularization α. The Horn–Schunck regularization controls the degree of spatial smoothing of the estimated flow. As a heuristic scale estimate, we write the corresponding smoothing length as

$$\ell_{\mathrm{HS}} \sim \frac{\alpha}{\|\nabla I\|/S}, \tag{11}$$

where $S \approx 1 \times 10^4$ DN s^{-1} is the intensity scale used for the normalization. Using the measured gradient magnitude $\|\nabla I\| = 337$ DN s^{-1} px^{-1} gives

$$\frac{\|\nabla I\|}{S} = 0.0337 \text{ px}^{-1}. \tag{12}$$

Requiring the characteristic smoothing scale to be of order 15–22 px therefore gives the corresponding heuristic range

$$\alpha \sim \ell_{\mathrm{HS}} \frac{\|\nabla I\|}{S} \approx 0.51\text{–}0.74. \tag{13}$$

This estimate is used only to constrain a sensible parameter range; it is not an exact prediction of the effective smoothing length of the discrete Horn–Schunck solver, which also depends on the discretization and iteration scheme. We therefore adopt the somewhat broader practical range $\alpha = 0.3$–1.0 for the calculations below.

Iterations N_{it} . The number of Jacobi iterations is not determined uniquely by the spatial scale argument above, because convergence depends on the regularization strength, the local image gradients, the discretization, and the initial flow field. We therefore do not interpret the iteration count as an analytically derived physical scale. Instead, we use $N_{\mathrm{it}} = 500$ as a conservative fixed iteration budget for the present processing, chosen to allow the iterative solution to approach a stable flow field while keeping the computational cost practical.

The recommended values arrived at above are, for convenience: Lucas–Kanade window $W = 7$, Sobel kernel $k = 5$, eigenvalue threshold $\lambda_{\min} = 5 \times 10^{-5}$; Farnebäck pyramid scale (`pyr_scale`) = 0.3, window size (`winsize`) = 15, 10 iterations, polynomial neighborhood (`poly_n`) = 7, polynomial smoothing (`poly_sigma`) = 1.0; FLCT correlation window size = 21, correlation threshold $\tau = 0.3$, with Gaussian weighting applied; Horn–Schunck regularization $\alpha = 0.3$–1.0 and 500 iterations.

The only quantities measured directly on the AIA data throughout this section are the handful listed above (plate scale, cadence, rim width, contrast, noise level, gradient magnitude, and the resulting displacement budget); every parameter value above is then produced by combining a subset of those quantities with a simple analytic scale-matching relation. A few of them (the Farnebäck polynomial smoothing, the Farnebäck and Horn–Schunck iteration counts, and the FLCT window) involve no AIA quantity at all, and the FLCT threshold is driven by the noise floor of the noisier instrument in our wider pipeline rather than by AIA itself. The point of walking through both the analytic relations and the AIA-measured inputs together here is not to claim a uniquely optimal configuration, but to illustrate, with a concrete worked example, that this kind of reasoning can give each algorithm a sensible starting point instead of leaving it to be tuned blindly.

3 Velocity field estimation: 10 May 2022 X1.5 event

We apply the FLCT, Horn–Schunck, Lucas–Kanade, and Farnebäck algorithms to the 10 May 2022 dataset. Velocity fields are estimated from object masks generated by the Wavetrack framework on the AIA 193 Å running-difference images. For the purpose of pipeline integration, acceptable accuracy in correctly determining the shapes is one of the first priorities. The velocity magnitudes recovered by all four algorithms (Figs. 2–6) are generally consistent with the wave-speed range reported by [Rigney et al., 2024]; the comparison below focuses on the spatial coherence of each velocity field and on the statistical agreement between the recovered speed distribution and the literature values.

FLCT velocity field on the X1.5 event

The FLCT velocity field (Fig. 2) recovers the geometric evolution of the coronal bright front across the entire approx 7.5-minute sequence. At 13:57:04 UT the wave appears as a compact, asymmetric blob anchored on AR 13006, dominated by speeds in the 300–600 $\mathrm{km\,s^{-1}}$ range with isolated yellow pixels ($\gtrsim 800\,\mathrm{km\,s^{-1}}$) embedded in the core. Within $\sim$70 s the front has expanded into a clear C-shaped arc opening eastward (13:58:16 UT), with both flanks lit up in yellow while the apex remains in the red part of the colour scale. The arc continues to stretch and rotate, becoming elongated and tilted toward the disk centre by 14:00:16 UT, and finally decays into a thinner crescent by 14:04:16 UT in which only the leading southern flank still carries high-speed pixels. This visual narrative is qualitatively consistent with the "flanks faster than apex" behaviour expected for an obliquely propagating wave.

Quantitatively, the aggregate speed histogram (Fig. 3) is markedly bimodal. The first lobe, between $\sim$50 and $\sim$400 $\mathrm{km\,s^{-1}}$, samples the interior pixels of the bright front and has a median speed of $\widetilde{|V|} = 295\,\mathrm{km\,s^{-1}}$ and a mean of $\langle |V| \rangle = 355\,\mathrm{km\,s^{-1}}$. The second lobe, between $\sim$500 and $\sim$700 $\mathrm{km\,s^{-1}}$, samples the rim and the flanks; it extends up to the 75th-percentile speed of 506 $\mathrm{km\,s^{-1}}$ and brackets the $\sim$700 $\mathrm{km\,s^{-1}}$ FLCT-derived flank speed reported by [Rigney et al., 2024] (marked in Fig. 3). The 95th percentile reaches 924 $\mathrm{km\,s^{-1}}$ and the tail extends to $\sim$1200 $\mathrm{km\,s^{-1}}$, exceeding the Rigney–24 FLCT-derived value; this tail could plausibly include a faster, shock-associated component, although we do not verify this independently here.

Examined frame by frame, the per-frame median is stable at $\sim$260–330 $\mathrm{km\,s^{-1}}$ while the 95th percentile decreases monotonically from $\sim$980 $\mathrm{km\,s^{-1}}$ early in the sequence to $\sim$700–800 $\mathrm{km\,s^{-1}}$ towards the end, exposing the deceleration of the front already visible as the fading high-speed tail in the speed-vs-time spectrogram of the same event.

The statistical agreement with [Rigney et al., 2024] is broadly consistent with both the FLCT outputs and the calibration approach of Sec. 2: the noise-to-contrast ratio measured on the AIA frames, $\eta = \sigma_n/\Delta I \approx 0.09$, is about a factor of three below the LASCO-derived correlation floor $\tau = 0.3$ used here, so the threshold is conservative for AIA in absolute terms, and we did not visually identify spurious high-speed pixels in the quiet-corona regions of

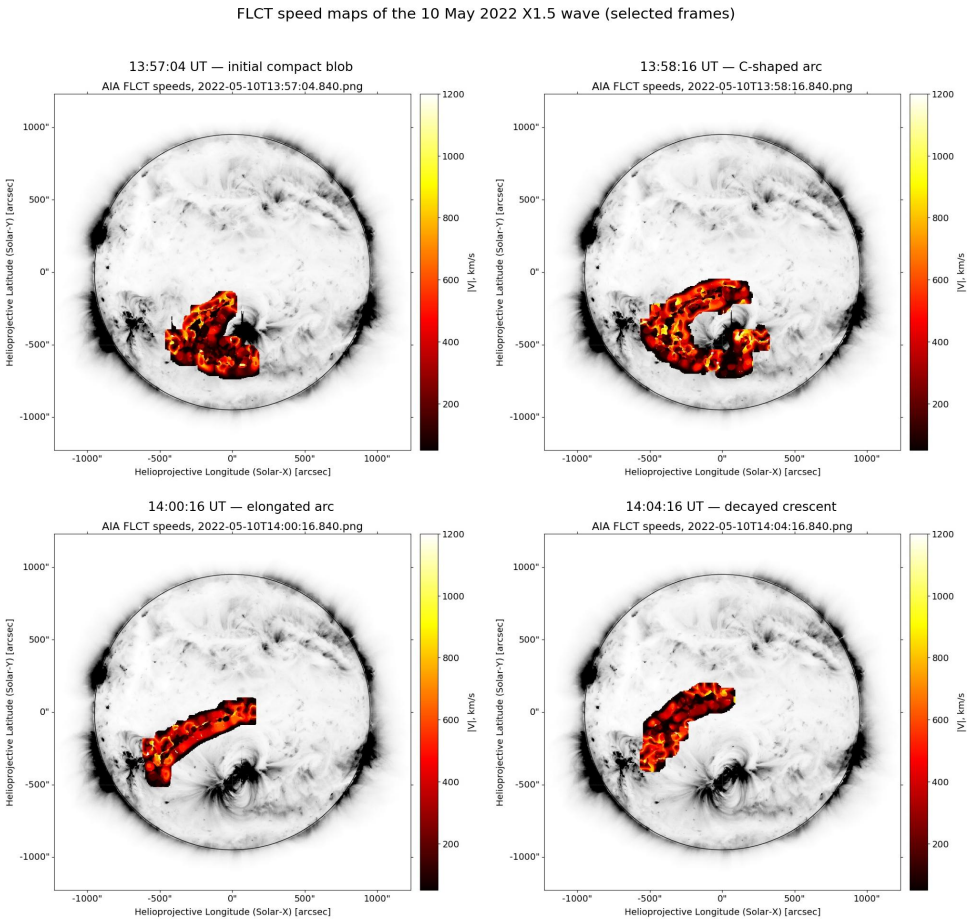

Fig. 2. Selected FLCT speed maps of the 10 May 2022 X1.5 wave, computed on AIA 193 Å running-difference frames. The four panels capture the lifecycle of the coronal bright front: a compact blob over AR 13006 at 13:57:04 UT, a C-shaped horseshoe arc opening eastward at 13:58:16 UT, an elongated arc rotated toward the disk centre at 14:00:16 UT, and a decayed crescent at 14:04:16 UT. In all four panels the rim and the southern/northern flanks (yellow, $\gtrsim 700\,\mathrm{km\,s^{-1}}$) move faster than the apex (red, ~ 300–$500\,\mathrm{km\,s^{-1}}$), consistent with an oblique-shock geometry.

Fig. 2 (though this was not quantified). Conversely, the choice of the Gaussian weighting kernel (mandatory for FLCT's FFT-based correlation, to suppress wrap-around artefacts at the patch edges) is what allows the rim pixels to be tracked despite their proximity to the bright-front boundary; without the Gaussian taper, the correlation peaks at the rim would be biased toward the static background. The FLCT result of Fig. 2 therefore illustrates the same parameter-scaling logic discussed for the HS/LK/FB outputs below, applied here to a Fourier-correlation framework.

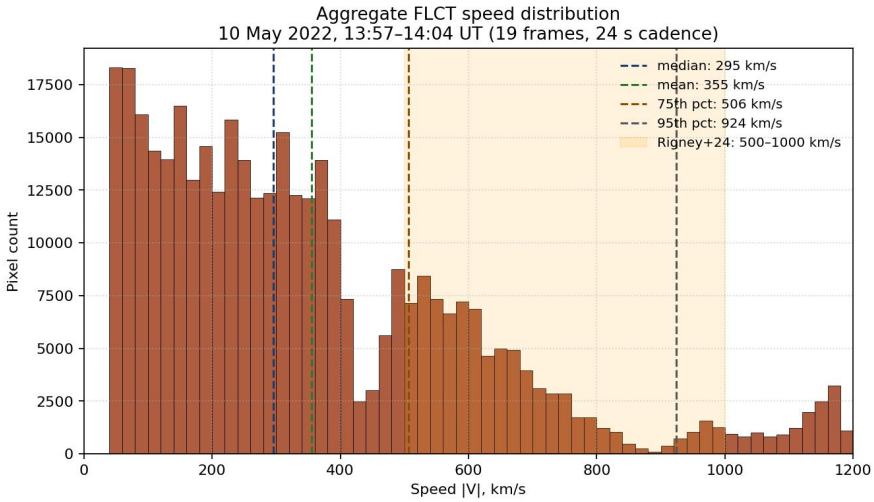

Fig. 3. Aggregate distribution of FLCT speed magnitudes for the 10 May 2022 X1.5 event, pooled over all 19 frames between 13:57 and 14:04 UT (380,595 valid pixels in total, 24 s cadence). The distribution is bimodal: a dominant interior bulk between ∼50 and ∼400 km s^{-1} representing the bright-front body, and a secondary component between ∼500 and ∼700 km s^{-1} tracing the rim and the flanks. The high-speed tail extends to ∼1200 km s^{-1}. The marked line indicates the ∼700 km s^{-1} FLCT-derived flank speed reported by [Rigney et al., 2024], which falls within the secondary peak.

Optical-flow velocity fields on the X1.5 event

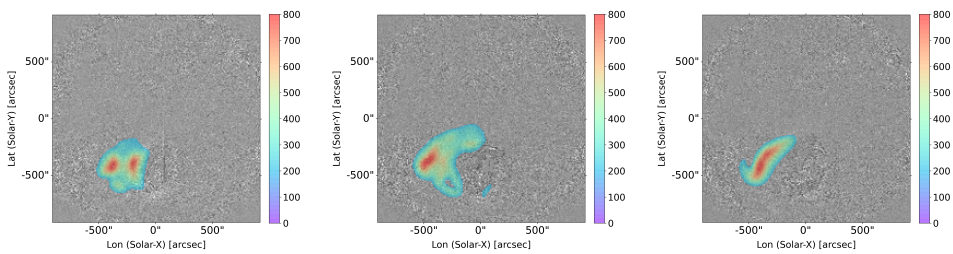

Fig. 4. Shock-wave propagation calculated using the Horn–Schunck optical-flow algorithm with the parameters derived in Sec. 2. The sequence shows coherent propagation, with a smooth and continuous shock-wave advancement over time in this example.

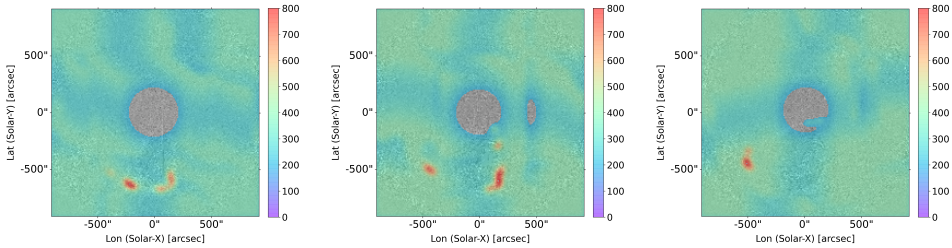

Fig. 5. Shock-wave propagation calculated using the Lucas–Kanade optical-flow algorithm with the parameters derived in Sec. 2. The sequence depicts a more fragmented and less consistent wave propagation.

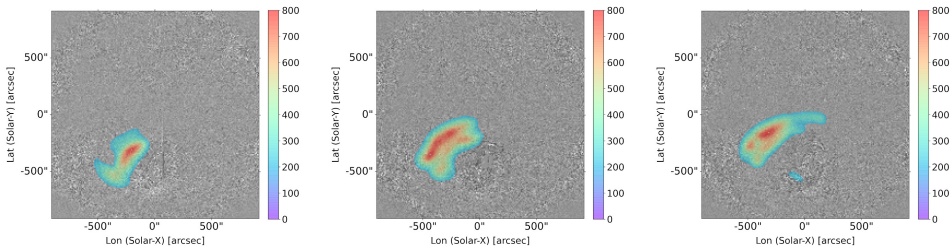

Fig. 6. Shock-wave propagation calculated using the Farnebäck optical-flow algorithm with the parameters derived in Sec. 2. The images show a well-defined shock-wave progression, comparable in coherence to the Horn–Schunck algorithm in this example, with a continuous wavefront.

The Horn–Schunck algorithm (Fig. 4) appears to resolve the shock wave's propagation reasonably well in this example, providing a smooth, continuous wavefront. This is consistent with the parameter choice discussed in Sec. 2: the regularization $\alpha = 0.3$–1.0 was set so that the diffusion length $\ell_{\mathrm{HS}} \approx \alpha/(\|\nabla I\|/S)$, with $S \approx 1 \times 10^4\,\mathrm{DN\,s^{-1}}$ as in Sec. 2, matches the feature scale, and the 500 Jacobi iterations are intended to be enough to propagate information across the feature radius. The result is a globally coherent field, consistent with what Horn–Schunck is designed to enforce; we did not notice obvious artefacts or discontinuities in this example.

Previously we had mixed results on the precision of the Lucas–Kanade algorithm [Stepanyuk and Kozarev, 2024] applied to ground-based and remote-instrument data. For this event, the algorithm (Fig. 5) exhibits notably poorer performance in this example: the shock-wave propagation appears fragmented, with notable discontinuities and inconsistencies in wavefront shapes across the

sequence. This outcome is also consistent with the parameter scaling derived above. The Lucas–Kanade window $W = 7$, set by $W \approx w/2$ with the AIA-measured rim width $w = 11$ px, resolves the rim edge but is much smaller than the feature radius $r \approx 15$–22 px. LK therefore enforces *only* local coherence and cannot, by construction, propagate flow information across regions of low gradient between rim segments. The local window-based approximation also struggles with the regions of large displacement implied by $\Delta \approx 14$–28 px per 12-second cadence – larger than the window itself.

The Farnebäck algorithm (Fig. 6) performs comparably well to Horn–Schunck in this example. The shock waves appear well defined and maintain continuity and coherence across frames. The pyramid scale $s = 0.3$ at $L = 1$ gives an effective window of ~ 50 px, which is intended to be large enough to absorb the AIA-measured displacement budget $\Delta_{\max} \approx 27.6$ px and allow the polynomial expansion of Eq. (6) to capture the nonlinear intensity variations across the front. Farnebäck therefore combines the local accuracy of a polynomial fit with the longer-range tracking ability of a multi-level pyramid – representing the shock-wave propagation in this example without the global oversmoothing risk that can affect Horn–Schunck-type regularization, or the local-window limitation seen in Lucas–Kanade here.

In summary, the scale-matching arguments of Sec. 2 are consistent with the relative performance of the four algorithms observed on the X1.5 event: FLCT, Horn–Schunck and Farnebäck all possess a mechanism (Fourier patch correlation with a Gaussian taper, global regularization, pyramid coverage) to bridge the displacement budget $\Delta_{\max} \gg W$ in the AIA data, whereas Lucas–Kanade does not. FLCT additionally returns a speed distribution whose secondary peak (500–700 $\mathrm{km\,s^{-1}}$) brackets the $\sim$700 $\mathrm{km\,s^{-1}}$ FLCT-derived flank speed of [Rigney et al., 2024], while its 95th percentile (924 $\mathrm{km\,s^{-1}}$) exceeds it, plausibly tracing a faster shock-associated component; this is at least qualitatively consistent with the calibration used here, though it remains a single-event comparison rather than an independent validation. The Lucas–Kanade algorithm's local method, by contrast, appears less practical in this configuration for on-disk shock-wave propagation than Horn–Schunck's global regularization, Farnebäck's polynomial-pyramid approach, or FLCT's Fourier correlation, which in this example produced visually clearer and more spatially coherent results.

4 Limitations and Outlook for Future Work

While the present study provides a conceptual framework and a transparent, physics-informed methodology for tuning optical-flow and local-correlation algorithms to moderately sized coronal bright fronts, two important limitations should be noted. First, both the parameter derivation above and the performance demonstration rely on a single, well-observed X1.5-class flare event (10 May 2022). Although this event is scientifically rich and representative of complex on-disk shock-wave propagation, the recommended parameter set has not yet been tested across a statistically meaningful sample of events. Its generalizability to quieter EUV waves, limb or off-disk geometries, lower signal-to-noise regimes, or data from other instruments (e.g., LASCO C2/C3 where the conservative correlation threshold $\tau = 0.3$ was chosen) therefore remains an open

question. The provided estimates are order-of-magnitude constraints rather than exact derivations.

These limitations point directly to several follow-up studies that would strengthen and extend the current concept. A natural next step is to perform a multi-event validation campaign using a carefully curated catalog of EUV waves (e.g., 20–50 events spanning quiet-Sun, active-region, and limb cases across multiple solar cycles and instruments). This would allow statistical assessment of parameter robustness and, if needed, the development of adaptive or event-specific tuning rules.

Quantitative metrics—such as velocity error relative to manual tracking, wavefront continuity scores, divergence/curl statistics, synthetic-data benchmarks, or cross-comparison with independent kinematic methods—on a larger datasets would be required for applications in pipelines.

Conclusion

Tuning of velocity estimation algorithms depends strongly on the spatial scale and motion intensity of the structures of interest. FLCT, and to some extent Lucas-Kanade when properly windowed and thresholded, appear effective in this example for moderate-scale coherent motions. Farneback provides dense and efficient estimations but requires careful adjustment of smoothing parameters to avoid over-generalization. Horn-Schunck appears reasonably robust in this example but requires balancing between structure-preserving and magnitude-sensitive regimes.

We discussed a practical set of algorithm parameters—arrived at by combining analytic reasoning with empirical measurements—and illustrated the performance of velocity-field estimation techniques for on-disk shock-wave propagation during the X1.5-class event of 10 May 2022. In practical applications, adaptive parameter-selection strategies are likely to deliver the most consistent results across varying motion fields.

Acknowledgments

This work is supported by the Bulgarian National Science Fund, VIHREN program, under contract KP-06-DV-8/18.12.2019 (MOSAIICS project)

References

Byrne, J., Long, D., Gallagher, P., Bloomfield, D., Maloney, S., McAteer, R., Morgan, H., and Habbal, S. R. (2013). Improved methods for determining the kinematics of coronal mass ejections and coronal waves. *Astronomy & Astrophysics*, 557:A96.

Farnebäck, G. (2002). *Polynomial Expansion for Orientation and Motion Estimation.* PhD thesis, Linköping University. Dissertation No. 790, ISBN 91-7373-475-6.

Fisher, G. H. and Welsch, B. T. (2008). FLCT: A Fast, Efficient Method for Performing Local Correlation Tracking. In Howe, R., Komm, R. W., Balasubramaniam, K. S., and Petrie, G. J. D., editors, *Subsurface and Atmospheric Influences on Solar Activity*, volume 383 of *Astronomical Society of the Pacific Conference Series*, page 373.

Horn, B. K. and Schunck, B. G. (1981). Determining optical flow. *Artificial Intelligence*, 17(1-3):185–203.

Kozarev, K. A., Davey, A., Kendrick, A., Hammer, M., and Keith, C. (2017). The coronal analysis of shocks and waves (cashew) framework. *Journal of Space Weather and Space Climate*, 7:A32.

Long, D. M., DeLuca, E. E., and Gallagher, P. T. (2011). The Wave Properties of Coronal Bright Fronts Observed Using SDO/AIA. *The Astrophysical Journal Letters*, 741:L21.

Lòpez-Urias, C., Vazquez-Becerra, G. E., Nayak, K., and López-Montes, R. (2023). Analysis of ionospheric disturbances during x-class solar flares (2021–2022) using gnss data and wavelet analysis. *Remote Sensing*, 15(18).

Lucas, B. D. and Kanade, T. (1981). An iterative image registration technique with an application to stereo vision. In *Proceedings of Imaging Understanding Workshop*, pages 121–130.

Rigney, J., Gallagher, P. T., Ramsay, G., Doyle, J. G., Long, D. M., Stepanyuk, O., and Kozarev, K. (2024). Tracking the motion of a shock along a channel in the low solar corona. *"Astron. Astrophys."*, 684:L7.

Stepanyuk, O. and Kozarev, K. (2024). Multi-instrument observations and tracking of a coronal mass ejection front from low to middle corona. *Journal of Space Weather and Space Climate*, 14:20.

Stepanyuk, O. and Kozarev, K. (2025). Segmentation and tracking of eruptive solar phenomena with convolutional neural networks. *Journal of Geophysical Research - Machine Learning and Computation*.

Stepanyuk, O., Kozarev, K. A., and Nedal, M. (2024). Multi-instrument observations of the dynamics of a solar coronal shock wave in the low and middle corona. *Journal of Space Weather and Space Climate*, 14:2.