



OPEN ACCESS

EDITED BY

Dipankar Banerjee,
Indian Institute of Astrophysics
(IIA), India

REVIEWED BY

Tanmoy Samanta,
Indian Institute of Astrophysics
(IIA), India
Sudarshan Bharadwaj,
Independent Researcher, Odenton,
United States

*CORRESPONDENCE

Shahin Jafarzadeh,
✉ shahin.jafarzadeh@qub.ac.uk

RECEIVED 31 March 2026

REVISED 04 May 2026

ACCEPTED 18 May 2026

PUBLISHED 17 June 2026

CITATION

Jafarzadeh S, Jess DB, Stangalini M,
Keys PH, Chambers G, Grant SDT,
Berretti M and Duckenfield TJ (2026)
Adaptive multi-line fitting for stable
line-core intensity and Doppler velocity.
Front. Astron. Space Sci. 13:1843842.
doi: 10.3389/fspas.2026.1843842

COPYRIGHT

© 2026 Jafarzadeh, Jess, Stangalini,
Keys, Chambers, Grant, Berretti and
Duckenfield. This is an open-access
article distributed under the terms of
the [Creative Commons Attribution
License \(CC BY\)](#). The use, distribution or
reproduction in other forums is
permitted, provided the original
author(s) and the copyright owner(s) are
credited and that the original
publication in this journal is cited, in
accordance with accepted academic
practice. No use, distribution or
reproduction is permitted which does
not comply with these terms.

Adaptive multi-line fitting for stable line-core intensity and Doppler velocity

Shahin Jafarzadeh^{1*}, David B. Jess^{1,2}, Marco Stangalini³,
Peter H. Keys¹, Glen Chambers¹, Samuel D. T. Grant¹,
Michele Berretti^{4,5} and Timothy J. Duckenfield¹

¹Astrophysics Research Centre, School of Mathematics and Physics, Queen's University Belfast, Belfast, United Kingdom, ²Department of Physics and Astronomy, California State University Northridge, Northridge, CA, United States, ³ASI Italian Space Agency, Rome, Italy, ⁴University of Trento, Trento, Italy, ⁵Department of Physics, University of Rome Tor Vergata, Rome, Italy

Next-generation solar spectrographs increasingly record dense wavelength windows in which tens to hundreds of spectral lines are sampled at each spatial location and time step. This expands the scope for multi-line, multi-height diagnostics of magnetohydrodynamic (MHD) motions, but also raises a practical challenge: deriving stable line-core intensity and line-of-sight (LOS) velocity time series when profiles evolve rapidly, become asymmetric, blend, or develop multi-lobed cores. Common fast estimators (e.g., parabolic minima, centre-of-gravity measures, Fourier-phase proxies, and fixed-window symmetric fits) can perform well for simple, isolated absorption lines, yet can intermittently misidentify the core in crowded or morphologically complex cases. Even infrequent mis-tracking can leave step-like artefacts that redistribute power and bias spectral, phase, and coherence measures used in wave and dynamics analyses. We introduce LineFit, a fully reproducible adaptive multi-line fitting approach tailored to dense-window spectroscopy from facilities such as DKIST, Sunrise III, and integral-field instruments. LineFit models each line locally with bounded non-linear least-squares fits to a Voigt-family profile, including an asymmetric-Voigt option to accommodate unequal wing broadening, and incorporates close-pair ownership control together with conservative, per-line window adaptation and split-core-aware handling. Using a synthetic near-UV time series with unambiguous ground truth, we benchmark LineFit against four widely used fast baselines and assess both instantaneous centre errors and downstream time-series diagnostics. Several fast methods remain competitive for many lines, whereas LineFit is most robust in key stress cases involving intermittently split-core profiles and correspondingly yields power spectra that agree most closely with the truth. We also demonstrate a proof-of-principle that benchmarks hybrid acceleration of the LineFit software via supervised emulation, offering at least three orders-of-magnitude improvement in processing time, with 87% of validation samples agreeing with LineFit to within 0.1 pm in recovered line-centre position. All code and notebooks required to reproduce the testbed, figures, and results are publicly available, providing a portable benchmark and a practical basis for stable multi-line time-series extraction, with a clear route towards hybrid acceleration via supervised emulation.

KEYWORDS

methods, MHD, oscillations, spectroscopy, the sun, waves

1 Introduction

New flagship facilities and modern spectroscopic designs are shifting the practical limiting factor in wave and dynamics studies of the solar atmosphere (Jess et al., 2023b). Instruments now routinely provide high spatial resolution, high cadence, and dense spectral sampling, increasingly in configurations that deliver many spectra per exposure (e.g., integral-field units (IFUs) and multiplexed/slit-based spectrographs; Iglesias and Feller, 2019; Katsukawa et al., 2020; Rast et al., 2021; Feller et al., 2025; Solanki et al., 2026). In this regime, dense spectral windows containing tens to hundreds of lines can be recorded for every spatial sample and time step, enabling multi-line constraints on magnetohydrodynamic (MHD) motions across atmospheric layers. At the same time, they can create an underappreciated bottleneck: converting raw spectra into science-grade, stable line-core intensity and line-of-sight (LOS) velocity time series at scale. This bottleneck is particularly relevant for next-generation datasets in which the information content is high but profile morphology is not static in the highly dynamic solar atmosphere, and where accurate time-series extraction must remain reliable both for well-behaved profiles and during episodes of strong profile distortion.

The scientific motivation for reliable line-core time series is broad. Doppler shifts and line-core intensity variations are key observables in studies of MHD waves and oscillations, shocks, flows, and small-scale transients, and they underpin frequency-phase-coherence analyses, wave-mode identification, and multi-height coupling inferences (e.g., Tziotziou et al., 2007; De Moortel and Nakariakov, 2012; Jess et al., 2015; Stangalini et al., 2021; 2022; Grant et al., 2022; Calchetti et al., 2023; Bate et al., 2024; Berretti et al., 2025; Duckenfield et al., 2025). In dense-window spectroscopy, their value can increase because multiple diagnostics can be extracted simultaneously from different lines sampling different formation regimes. However, the same dense windows may also intensify algorithmic challenges, particularly if blends become common, line widths vary strongly across the field of view and in time, asymmetries appear and disappear, or some lines develop multi-lobed cores or self-reversals. In such conditions, even occasional misidentification of the line core can be scientifically consequential because it can introduce discontinuities that inject broadband artefacts, bias power spectra, and degrade phase and coherence measurements.

A wide range of techniques are used in practice to infer LOS velocities and related atmospheric parameters. At one end, fast proxies estimate a representative line position from the observed intensity profile within a fixed or semi-fixed window: e.g., parabolic interpolation around the minimum, centre-of-gravity (COG) estimators of the absorption depression, bisectors, Fourier/phase-based centroids (Fourier-tachometer style proxies), polynomial core fits (including weighted variants), or symmetric parametric fits (Gaussian/Lorentzian) in a prescribed window (e.g., Uitenbroek, 2003; Couvidat et al., 2012; Beckers and Brown, 2013; Gilchrist-Millar et al., 2021). These approaches can be accurate and efficient for isolated, near-symmetric profiles with stable continua, and they often remain the only feasible option when computational cost limits what can be processed. Their limitations are also well known: if neighbouring features or blends fall within the chosen wavelength window, they can bias methods that

compute the centre from the full windowed profile (e.g., COG or Fourier-phase centroids); asymmetric or evolving wings can systematically shift COG-like measures; and multi-lobed cores can cause the inferred “centre” to hop between competing minima or lobes, producing discontinuities in time series that are difficult to identify and correct in subsequent analysis (Bouchy et al., 2001; Figueira et al., 2013; Figueira, 2018).

At the other end, spectropolarimetric inversions infer atmospheric parameters by fitting Stokes profiles under radiative-transfer constraints (Ruiz Cobo and del Toro Iniesta, 1992). Inversions (and related forward-modelling frameworks) provide a physically interpretable route to LOS velocities (and magnetic/thermodynamic structure) and can be essential when magnetic diagnostics are measured. They also come with substantial computational cost, model assumptions, and practical sensitivity to initialisation, parameterisation, and data quality. Inversion families span Milne–Eddington (e.g., Lagg et al., 2004; Lites et al., 2007; Borrero et al., 2011) and depth-dependent approaches (e.g., Frutiger et al., 2000; Socas-Navarro et al., 2015; Milić and van Noort, 2018), LTE¹ and non-LTE formalisms, and increasingly multi-line strategies that exploit complementary constraints from multiple transitions (e.g., de la Cruz Rodríguez et al., 2019; Riethmüller and Solanki, 2019; Ruiz Cobo et al., 2022; Hölken et al., 2026). In practice, the choice of inversion method and configuration depends on the scientific objective, the line(s) observed, and the cadence/volume constraints, and it is often not feasible to run full inversions for every pixel and time step in very large field-of-view and high-cadence datasets. As a result, many workflows still rely on line-core time-series extraction as an intermediate, scalable product even when inversions are applied to subsets or representative regions.

Between these extremes, machine-learning (ML) approaches and emulators are increasingly used to accelerate profile interpretation and/or produce robust proxies. Examples include supervised tools designed to fit or classify chromospheric line profiles in an automated manner (e.g., the MCALF framework; MacBride et al., 2021), as well as broader efforts to emulate expensive inversions or to enable robust, scalable classification of spectropolarimetric profile morphologies for large surveys and statistical comparisons (Asensio Ramos and Díaz Baso, 2019; Jarolim et al., 2025; Campbell et al., 2025). While ML strategies can deliver dramatic speed-ups, their reliability depends on training coverage, careful validation, and explicit handling of out-of-distribution profiles—challenges that become important precisely in the highly dynamic regimes, such as those for wave and transient studies. Irrespective of whether future workflows rely on parametric fitting, ML, or hybrids, well-characterised reference extractions provide a necessary anchor for validation, uncertainty assessment, and quality control in these dynamic regimes.

In this paper we focus on a complementary but widely needed product: robust, automated extraction of line-core intensity and LOS velocity time series across many lines in dense windows, suitable for subsequent wave and dynamics analysis. Our goal is not to replace inversions, but to provide an adaptive, reproducible approach that reduces failure modes of common fast estimators in crowded and

¹ Local Thermodynamic Equilibrium.

evolving spectra, and that scales to multi-line use cases. As such, this method-focused contribution is intended to be immediately usable by observers working with, e.g., Daniel K. Inouye Solar Telescope (DKIST; [Tritschler et al., 2016](#)) and Sunrise III balloon-borne solar observatory ([Korpi-Lagg et al., 2025](#)) spectroscopy, and Fibre Resolved OpticAl and Near-Ultraviolet Czerny-Turner Imaging Spectropolarimeter (FRANCIS; [Jess et al., 2023a](#)), and it is readily transferable to other slit and IFU spectrographs that deliver (or will deliver) time-dependent line profiles (e.g., the Interface Region Imaging Spectrograph (IRIS), the Multi-slit Solar Explorer (MUSE), and Solar-C; [De Pontieu et al., 2014, 2020](#); [Shimizu et al., 2021](#)), while also remaining applicable to single-line time-series studies. We present **LineFit**, an adaptive multi-line fitting methodology built around four design principles: (i) line-guided seeding near expected centres (with explicit handling of close blended pairs), (ii) conservative per-line adaptive windowing with hard safety bounds, (iii) iterative centre refinement for stability under profile evolution, and (iv) optional wing-weighted or wing-masked fitting to suppress lobe-hopping when the core becomes complex. We benchmark the method using a realistic synthetic near-UV testbed with unambiguous ground truth, compare against a number of widely used fast baselines, and quantify not only per-time-step centre errors but also the impact on standard time-series wave products. Finally, we briefly note how such a robust fitting baseline could support scalable processing strategies, including hybrid acceleration via supervised emulation.

2 Materials and methods

This section describes the validation dataset and the comparison baselines used throughout the paper. We adopt a synthetic, dense-window spectral time series because it provides unambiguous ground truth for the line-centre evolution and enables controlled stress tests (blends, asymmetries, close pairs, and split/reversal morphologies) that are difficult to isolate and quantify in real observations. The goal is not to replace real-data demonstrations, but to establish a reproducible performance benchmark where centre errors and time-series artefacts can be measured directly. All routines used to generate the synthetic dataset, run the estimators, and reproduce all figures presented in this work are publicly available in a GitHub repository described in [Section 2.3](#).

2.1 Synthetic multi-line testbed

We validate and benchmark the proposed adaptive multi-line fitting methodology using a realistic synthetic near-UV spectral time series designed to mimic modern dense-window spectroscopy, where tens of absorption lines coexist within a narrow wavelength interval, and line profiles evolve significantly in time. The testbed consists of $N_{\text{lines}} = 10$ diagnostic lines sampled on a common wavelength grid $\lambda \in [\lambda_{\min}, \lambda_{\max}]$, spanning a total window width $\Delta\lambda \equiv \lambda_{\max} - \lambda_{\min} \approx 0.36$ nm around a central wavelength of approximately 299.90 nm, with $N_{\lambda} = 2400$ spectral samples (uniform spacing $d\lambda = (\lambda_{\max} - \lambda_{\min})/(N_{\lambda} - 1) \approx 1.5 \times 10^{-4}$ nm). The cadence is $\Delta t = 10$ s and the duration is $N_t = 300$ time steps ($T = 3000$ s). Each time step contains a full spectrum $I(\lambda, t)$ in which (i) line depths and widths vary smoothly in time, (ii) asymmetric

profiles exhibit time-dependent core skewness, (iii) a close blended pair has controlled separation, and (iv) one line intermittently develops a split (two absorption lobes) with a central emission-reversal component. In addition, weak blending contaminants are injected near selected lines to mimic crowding in dense windows.

All line profiles are synthesised from Voigt-based primitives. For a nominal rest wavelength $\lambda_{0,i}$, the time-dependent Doppler-shifted line-centre wavelength is,

$$\lambda_i(t) = \lambda_{0,i} \left(1 + \frac{v_i(t)}{c} \right),$$

where $\lambda_i(t)$ is the ground-truth line-centre wavelength of line i at time t , $v_i(t)$ is the imposed LOS velocity driver and c is the speed of light. The absorption component of each line is represented by a normalised Voigt shape with time-dependent depth and widths. Asymmetric profiles are produced by a controlled, time-dependent core skewness (i.e., the deepest point is displaced relative to the nominal Doppler-shifted centre), so that local-minimum based estimators face a realistic ambiguity in “what constitutes the core”. The split+reversal morphology is produced by transitioning between a single absorption core and a two-lobe split profile with time-dependent separation $\Delta\lambda_{\text{split}}(t)$, and by adding a Gaussian emission component localised near the nominal line centre. Importantly, the central emission feature is not required to exceed the continuum; throughout this work, “emission reversal” refers to a local core brightening relative to the surrounding absorption depression, i.e., a split/multi-lobed structure in the core region.

2.1.1 Time-dependent evolution and truth products

Temporal evolution is imposed through stochastic-but-smooth processes that approximate non-stationary solar dynamics. For each line, the LOS velocity $v_i(t)$ is generated as a multi-frequency signal composed of three to five sinusoidal components with random phases, time-varying amplitudes, small time-dependent frequency drift, low-frequency jitter, and intermittent smoothed impulses. Component frequencies are drawn within 2–9 mHz. To avoid artificial discontinuities, no hard clipping is applied to $v_i(t)$; instead, each $v_i(t)$ is smoothly rescaled by a single multiplicative factor so that $\max|v_i(t)|$ does not exceed a prescribed per-line cap. In the configuration used here, caps span 1.2–3.5 km s⁻¹ across the ten lines. Line depth and width parameters evolve via first-order autoregressive [AR(1)] drivers, and the split+reversal line includes additional stochastic modulation of the splitting strength and core emission amplitude.

The synthetic FITS product stores, in addition to the spectrum cube $I(\lambda, t)$, the rest wavelengths $\lambda_{0,i}$, the ground-truth time-dependent line-centre wavelengths $\lambda_i(t)$, and the truth velocities $v_i(t)$. These truth products enable direct, per-time-step centre residuals for any estimator, defined as

$$\Delta\lambda_i(t) = \hat{\lambda}_i(t) - \lambda_i(t),$$

where $\hat{\lambda}_i(t)$ is the estimated line-centre wavelength returned by the method.

2.1.2 Noise and instrumental effects

To approximate realistic measurement conditions, the synthetic spectra include (i) a weak continuum gradient across the window, implemented in continuum-normalised (dimensionless) intensity units as

$$I_{\text{cont}}(\lambda) = 1 + s \frac{\lambda - \lambda_{\text{mid}}}{\Delta\lambda},$$

where $I_{\text{cont}}(\lambda)$ is the continuum level as a function of wavelength, s is the dimensionless gradient amplitude, $\lambda_{\text{mid}} \equiv (\lambda_{\text{min}} + \lambda_{\text{max}})/2$ is the midpoint of the simulated wavelength window, and $\Delta\lambda \equiv \lambda_{\text{max}} - \lambda_{\text{min}}$ is the full window width. In the benchmark presented here we use $s = 0.02$, (ii) a low-amplitude quasi-periodic ripple to mimic residual fringing/throughput structure (amplitude 1.5×10^{-3} and wavelength period of 0.035 nm), and (iii) additive Gaussian noise whose amplitude increases in deep absorption cores (“core-boosted” noise). In the configuration used here, the continuum noise standard deviation is $\sigma_{\text{cont}} = 2.1 \times 10^{-3}$ (relative intensity units) and the noise is increased in deep cores by a scaling boost factor of 2.0.

Optionally, a spectral point-spread function (PSF) convolution can be applied along the wavelength axis; however, we do not include PSF convolution in the benchmark presented here in order to isolate algorithmic centre-tracking behaviour from instrument-dependent spectral broadening and to keep the ground-truth morphology and blend structure as directly comparable as possible across methods (the PSF option remains available for sensitivity tests and instrument-specific tuning).

Other observational effects, such as scattered light, residual flat-field errors, spatially correlated noise, or imperfect wavelength calibration, are not explicitly included in the baseline benchmark presented here. These could modify the apparent core depth, local contrast, or wing shape, and therefore affect the absolute performance of any centre-tracking method. In general, methods that use constrained local fitting and conservative window selection are expected to be less sensitive to such complications than unconstrained fast estimators, although quantifying this in detail is beyond the scope of the present benchmark. In the case of LineFit (described in Section 2.2), the bounded fitting, conservative window adaptation, and morphology-aware handling are specifically intended to reduce sensitivity to such complications by preventing large excursions toward unrelated structures or unstable competing minima.

These perturbations are also visible in Figure 1 as (i) a slight large-scale tilt of the local continuum across the full spectral window, produced by the weak continuum gradient, (ii) low-amplitude quasi-periodic waviness between and around lines, produced by the ripple term, and (iii) small local intensity irregularities that become more noticeable in deeper cores, produced by the core-boosted Gaussian noise.

2.2 Adaptive multi-line profile fitting: LineFit

The proposed method (LineFit) is designed for dense spectral windows with many lines and evolving morphologies, where a single fixed fitting configuration is insufficient. In the implementation used here, each spectrum is first sanitised for non-finite samples (interpolated along wavelength index where needed),

then robustly normalised to the local min–max range, lightly smoothed (optional), and intensity-inverted so that absorption cores become peaks². LineFit then enforces stable centre tracking via three coupled stages:

- i. *Coarse seeding near expected centres (including close-pair ownership control)*. Initial seeds are generated in the intensity-inverted spectrum by searching locally around each expected rest wavelength. Candidate peaks are identified via peak finding and scored using a (dominant) distance-to-prior term (preference for proximity to the expected centre) together with a (minor) peak-height term (a small preference for deeper, better-defined absorption cores, i.e., stronger peaks in the intensity-inverted spectrum). This means that shallow secondary valleys or unrelated nearby structures are not usually selected unless they remain sufficiently close to the expected centre, and for declared close pairs the explicit valley-based ownership split further prevents neighbouring lines from capturing each other’s minima. Seeds are further guarded by a maximum allowed centre offset from the expected wavelength, preventing spurious matches to unrelated structures in crowded windows. For declared close pairs, “ownership” is enforced by splitting the shared search region at the inter-line valley, i.e., assigning each member of the pair its own side of the valley, so that the two members cannot swap under moderate profile evolution.
- ii. *Per-line adaptive windowing with hard safety bounds and morphology gating*. Each line i starts from a per-line baseline half-window (or a global default), but the effective half-window is constrained by user-provided per-line minimum and maximum half-windows. Window adaptation is conservative: it expands only when justified by local measurements of profile width and complexity, and it can contract to protect the core when the measured depression width indicates that a smaller window is adequate (limited by the per-line minimum half-window). Two morphology-specific gates, i.e., decision rules that activate specialised handling only when a given profile morphology is detected, are applied in this paper: (a) for lines declared as potential self-reversal candidates, a split-core detector is evaluated time-step by time-step in the intensity-inverted profile; specifically, the detector searches for two sufficiently prominent local peaks separated by an inter-peak valley that drops below a prescribed fraction of the weaker peak, thereby distinguishing a genuine split-core structure from small fluctuations or noise. If such a morphology is detected, reversal handling is activated only for that time step, forcing the fit window to the permitted maximum so both lobes are contained. Operationally, this detector identifies a two-lobed core structure with a sufficiently deep inter-lobe valley in the intensity-inverted profile, so that the fit can be stabilised around the overall split-core morphology rather than jumping to one of the competing lobes; (b) for

² Here “inverted” refers only to an intensity transformation (absorption minima mapped to peaks). It does not refer to Stokes inversions or atmospheric inversions.




FIGURE 1

Example synthetic spectrum and line-centre recovery. The blue curve shows the normalised spectrum; red dashed segments indicate the fitted model within each adaptive window. Vertical markers show the recovered centres (green dashed) and the truth centres (black dotted). Numbered circles label the ten diagnostic lines used throughout the testbed.

non-reversal lines, optional emission-bump and secondary-peak guards are used to avoid unsafe window growth in the presence of nearby competing structures (e.g., blends), thereby reducing the risk of window leakage. In practice, conservative thresholds are preferred: false positives mainly reduce computational efficiency by triggering safer handling more often than necessary, whereas false negatives are more consequential because they allow a genuinely complex profile to be treated too simply.

- iii. *Iterative refinement of centres with asymmetric handling.* Given the current window, a parameterised Voigt-family line model (Voigt or asymmetric Voigt) is fitted via bounded non-linear least squares optimisation in intensity-inverted space, with the fitted parameters defining the local core shape, width, amplitude, and centre. We adopt a Voigt-family model because it provides a flexible yet compact description that can accommodate Gaussian-like broadening (e.g., thermal/instrumental) together with Lorentzian-like wings (e.g., damping/pressure broadening), while remaining stable under bounded fitting. Purely Gaussian or purely Lorentzian fits can systematically misrepresent either the core or the wings for many photospheric/chromospheric lines. Our aim here is not to claim that the Voigt family provides a unique physical description of every observed profile, but rather to adopt a compact and sufficiently flexible parametric form that yields stable centre tracking across a wide range of morphologies. In this context, the centre position is the primary product of interest, and modest mismatch in the detailed wing shape is generally less important than avoiding unstable centre shifts between competing minima or lobes. The choice of a Voigt-family model should therefore be viewed as a pragmatic compromise between flexibility, stability, and computational cost, rather than as a full radiative-transfer description of the line formation process. The fitted centre is used to recentre the window and the fit is iterated for a fixed number of iterations or until convergence. Wing asymmetry is handled explicitly through an asymmetric Voigt parameterisation, allowing the fit to accommodate unequal wing broadening while maintaining bounded parameters. The parameter bounds are imposed as practical local constraints to keep the solution within a plausible regime for the current line and fitting window, rather than being derived from an external library of fixed Voigt parameters for each transition. In the

present implementation, the choice between a Voigt and an asymmetric-Voigt profile is set by the fitting configuration, i.e., the algorithm does not automatically try both models and select between them at each time step. The refinement therefore contains two conceptually distinct levels: an outer loop in which the fitted centre is used to recentre the window and repeat the fit if needed, and an inner bounded non-linear least-squares optimisation that solves for the profile parameters within a given window. To further suppress discontinuities associated with lobe-hopping or unstable asymmetry solutions, the implementation includes (for asymmetric-Voigt fits) a guarded fall-back refit in which the asymmetry-parameter bounds are tightened when the initial solution lies close to parameter bounds or window edges. This fall-back is disabled for time steps where a split-core morphology has been detected, so that genuine multi-lobed behaviour is not artificially suppressed.

A schematic summary of this workflow is provided in Figure 2. The primary outputs are the fitted line-centre wavelengths $\hat{\lambda}_i(t)$, line-core intensity samples $I(\hat{\lambda}_i, t)$ evaluated at the fitted centre, and LOS velocities $\hat{v}_i(t) = c(\hat{\lambda}_i - \lambda_{0,i})/\lambda_{0,i}$ computed relative to the rest wavelengths. For transparency and optional regression testing, the implementation can optionally print diagnostic information, including the half-window used for each line and time step (and, when enabled, fitted parameter values).

Figure 1 illustrates a representative spectrum from the synthetic time series together with the fitted model segments and the recovered line centres. The ground-truth centres for the selected time step are overplotted, enabling direct visual verification of centre recovery across the spectral window. The corresponding per-line residuals at that time step demonstrate sub-picometre accuracy for the adaptive fitting approach under these controlled conditions, with absolute centre residuals of (1) 0.01 pm (0.007 km s⁻¹), (2) 0.10 pm (0.103 km s⁻¹), (3) 0.05 pm (0.049 km s⁻¹), (4) 0.01 pm (0.008 km s⁻¹), (5) 0.03 pm (0.033 km s⁻¹), (6) 0.03 pm (0.029 km s⁻¹), (7) 0.01 pm (0.005 km s⁻¹), (8) 0.02 pm (0.018 km s⁻¹), (9) 0.01 pm (0.009 km s⁻¹), and (10) 0.02 pm (0.019 km s⁻¹).

In practical use, the most dataset-dependent settings are the expected rest wavelengths and the per-line minimum/maximum half-window bounds, which should reflect the spectral sampling, typical line widths, and the degree of crowding in the observed window. By contrast, the overall logic of the method–line-guided




FIGURE 2

Schematic overview of the LineFit workflow. The method proceeds through three coupled stages: (i) coarse seeding, including local search near the expected wavelength together with offset guarding and close-pair ownership control; (ii) adaptive windowing, in which per-line windows are conservatively adjusted within prescribed bounds and morphology checks activate safer handling when needed; and (iii) iterative bounded fitting, in which Voigt-family models are fitted, recentered, and iterated to convergence, yielding the recovered line centre, LOS velocity, and line-core intensity.

seeding, close-pair ownership control, bounded fitting, iterative recentering, and morphology-aware handling—is portable and does not require substantial redesign between datasets. In this sense, LineFit is intended to be tuned mainly at the level of conservative per-line bounds, rather than through highly case-specific heuristics.

2.3 Reproducibility and code availability

The line-fitting implementation (LineFit) is distributed as part of the workflow repository `WaLSAlib`³, which collects practical routines that connect calibrated data to analysis-ready time series for wave and dynamics studies. All Python scripts and notebooks required to reproduce the synthetic testbed, all figures in this paper, and the method comparisons are provided under `examples/LineFit_Frontiers` in the same repository.

2.4 Other line-centre estimators

We compare the proposed adaptive fitting approach against four commonly used fast estimators. All baseline methods are applied independently per line, using a per-line half-window in index space: we use the same nominal per-line half-window settings as for LineFit across all baselines for a fair comparison; LineFit may adapt within its prescribed bounds, while the baselines keep the window fixed. For each method, the output is a line-centre estimate $\hat{\lambda}_i(t)$ and residuals are evaluated relative to the ground-truth line-centre wavelength $\lambda_i(t)$.

2.4.1 Parabolic minimum (quadratic interpolation)

The parabolic-minimum estimator identifies the local intensity minimum within a prescribed search window and fits a quadratic polynomial through the three samples surrounding that minimum. The vertex of the quadratic provides a sub-sample line-centre estimate. This method is computationally inexpensive and often

robust for isolated, near-symmetric lines, but it can be biased or destabilised when the profile becomes asymmetric, multi-lobed, or strongly blended.

2.4.2 Centre-of-gravity (COG) of the absorption depression

The COG estimator is computed from the absorption depression,

$$D(\lambda, t) = 1 - \frac{I(\lambda, t)}{I_{\text{cont},i}(t)},$$

where $I_{\text{cont},i}(t)$ is a local continuum estimate for line i at time t , obtained robustly from the top fraction of samples within the fitting window (here the top 20% in each window). The centre is then,

$$\hat{\lambda}_{\text{COG}}(t) = \frac{\sum_{\lambda} \lambda D(\lambda, t)}{\sum_{\lambda} D(\lambda, t)},$$

where the sums are taken over the wavelength samples within the fitting window, and $D \geq 0$ enforced by clipping. COG is inexpensive and often stable for moderately deformed profiles, but it can be sensitive to blends, window definition (i.e., window “leakage”), and changes in wing asymmetry because it integrates over the full window.

2.4.3 Fourier-phase centroid ($k = 1$ phase)

The Fourier-phase estimator computes a reference-free centroid from the phase of the first Fourier component ($k = 1$) of an apodised (Hanning) depression profile within the fitting window. Conceptually related to Fourier-tachometer style Doppler proxies, this estimator can be efficient and less sensitive to pixel-scale noise, but can fail or become biased when the profile contains multiple competing structures within the window (e.g., split cores or strong blends).

2.4.4 Weighted polynomial core fit (Poly6w)

The weighted polynomial method fits a high-order polynomial (here order 6) to a narrow core region around the local

³ <https://GitHub.com/WaLSAteam/WaLSAlib>




FIGURE 3

Instantaneous centre accuracy comparison for a representative synthetic time step ($t = 0$ s). Lower panel: spectrum with recovered centres from LineFit and four fast baselines overplotted, alongside the truth centres. Upper panel: normalised absolute centre residuals ($|\Delta\lambda_i|/\max|\Delta\lambda_i|$) across methods for each line; text annotations show the corresponding absolute maximum $|\Delta\lambda_i|$ per line (pm).

minimum, using weights that emphasise samples closest to the core (inverse-distance weighting). The centre is obtained by locating the minimum of the fitted polynomial within the fit interval (e.g., via stationary points of the derivative, with interval-endpoint checks). This can outperform simple parabolic interpolation for smooth symmetric cores, but remains vulnerable to multi-lobed cores and to mis-centring in crowded windows.

3 Analysis and results

We quantify estimator performance in two complementary ways. First, we evaluate instantaneous line-centre errors at a representative time step (Figure 3), where the truth centres are known exactly. Second, we assess time-series fidelity over the full 50 min sequence (300 time steps) by comparing recovered LOS-velocity time series against the truth $v_i(t)$ in Figure 4, then by propagating the resulting time series into standard wave diagnostics (Figure 5).

3.1 Single-time-step centre accuracy in the spectral window

Figure 3 compares five line-centre estimators for one time step (at $t = 0$ s) within a window containing ten lines, including a close blended pair (lines 4–5) and a split-core/emission-reversal candidate (line 6). The lower panel shows the spectrum and the recovered centres overplotted alongside the truth centres, while the upper panel summarises the absolute residuals $|\Delta\lambda_i| = |\hat{\lambda}_i - \lambda_i|$ for each line, normalised by the maximum $|\Delta\lambda_i|$ across methods for that line to highlight relative method performance line-by-line. The label above each line group gives the absolute maximum $|\Delta\lambda_i|$ in picometers across methods for that line.

For this example time step, LineFit achieves sub-picometre centre recovery for all ten lines, with a maximum absolute residual of 0.103 pm across the set. The fast baseline estimators show small errors for many isolated lines, but exhibit substantially larger

excursions for the split-core line (line 6), reaching picometre-level offsets ($|\Delta\lambda_i| \approx 1.0$ – 2.2 pm depending on method; $|\Delta\lambda_i| = 0.03$ pm for LineFit). This behaviour is consistent with the known failure modes of minimum- and core-centroid proxies under multi-lobed cores: the inferred “centre” can shift towards one lobe or towards a blended feature when the core morphology becomes non-single-valued within the window. By contrast, in LineFit the morphology-aware handling of the split-core case is intended precisely to avoid such lobe-hopping, so that the recovered centre remains tied to the overall core structure rather than to one of the instantaneous lobes.

3.2 Time-series fidelity: velocity residuals over 50 min

Instantaneous accuracy does not guarantee time-series stability: rare centre jumps can inject broadband power and bias wave diagnostics. We therefore compute, for each line and each method, the LOS velocity time series following,

$$\hat{v}_i(t) = c \frac{\hat{\lambda}_i(t) - \lambda_{0,i}}{\lambda_{0,i}},$$

and evaluate the residuals against the truth $v_i(t)$ over all 300 time steps. As a summary metric we use the per-line Root Mean Square Error (RMSE),

$$\text{RMSE}_i = \sqrt{\langle (\hat{v}_i(t) - v_i(t))^2 \rangle_t},$$

where $\langle \cdot \rangle_t$ denotes averaging over time samples; in practice, the RMSE is computed over time samples where both truth and estimate are finite.

Figures 4a–c shows representative velocity time-series comparisons for three lines spanning increasing profile complexity: an isolated symmetric line (line 1), an asymmetric/blended line (line 2), and the split-core/emission-reversal line (line 6). For the isolated case, all methods track the truth reasonably well, typically showing the closest agreement and the smallest residual structure. For the asymmetric/blended case, the dispersion between methods increases, consistent with sensitivity to evolving asymmetry and





FIGURE 4
 Velocity time-series fidelity over the full synthetic sequence. Rows (a–c) examples for an isolated symmetric line (line 1), an asymmetric/blended line (line 2), and a split-core/emission-reversal line (line 6). Each row shows truth and one recovered method per sub-panel to aid visual comparison. Residual-versus-truth time-series plots for the same representative lines. (d) per-line RMSE of $\hat{v}(t)$ against the truth for all ten lines and all methods.



to blends or window leakage (when a fixed window includes neighbouring structure); LineFit and COG broadly capture the dominant variation, whereas FFT-phase and, more prominently, the parabolic minimum and Poly6w exhibit larger deviations at several intervals.

The most pronounced differences occur for the split-core/emission-reversal line (line 6). Here LineFit remains stable and closely follows the truth throughout the sequence, while all four fast estimators exhibit substantially larger mismatches, including episodes of divergence associated with time steps where the core morphology becomes strongly multi-lobed (see the spectrum evolution in the online movie⁴, which can be regenerated from the public repository). This behaviour is reflected directly in the per-line RMSE summary (Figure 4d): for line 6, LineFit yields RMSE =

0.067 km s⁻¹, whereas the fast methods range from RMSE ≈ 0.58 km s⁻¹ (COG) to ≈ 1.09 km s⁻¹ (parabolic minimum), i.e., nearly an order-of-magnitude degradation in time-series fidelity for this complex-core case.

Across all ten lines, LineFit provides the best overall performance in this testbed (mean RMSE 0.094 km s⁻¹; median 0.083 km s⁻¹), with COG second-best (mean RMSE 0.121 km s⁻¹). The parabolic minimum, FFT-phase, and Poly6w methods show larger mean RMSE values (0.30, 0.21, and 0.26 km s⁻¹, respectively), driven primarily by their failures on the split-core line and, to a lesser extent, by sensitivity to evolving asymmetry and blends. We emphasise that these baseline implementations are intentionally kept simple and uniform across lines to provide a controlled comparison; more aggressive method-specific tuning and case-dependent heuristics may improve performance on some profiles, at the expense of additional parameter complexity and reduced portability.

⁴ https://WaLSA.team/LineFit/synthetic_spectrum_timeseries

Whether a given centre error is scientifically acceptable depends primarily on how it propagates into the derived LOS-velocity time series and wave diagnostics. In this benchmark, the simpler estimators are often adequate for well-behaved lines, but become problematic in the morphologically complex stress cases that motivate LineFit.

3.3 Impact on wave diagnostics: refined global wavelet spectra

To quantify how centre-extraction differences propagate into wave inference, we compute Refined Global Wavelet Spectra (RGWS) from each recovered velocity time series using the WaLSAtools⁵ (Jafarzadeh et al., 2025b; a) implementation. Briefly, RGWS is the time-integrated Morlet wavelet power refined by retaining only statistically significant power (95% confidence) and excluding regions within the cone of influence (subject to edge effects). All velocity signals are linearly detrended and apodised using a Tukey window ($\alpha = 0.1$) prior to the wavelet analysis.

Figure 5 compares RGWS for two representative lines: a blended line that is one component of the close pair and exhibits relatively strong depth variability (line 4) and the split-core/emission-reversal line (line 6). For line 4 (which has the second-largest RMSE across the five methods, though still modest), all methods recover consistent peak locations and broadly similar spectral structure, indicating that for moderately complex but single-core profiles the baseline estimators can preserve the dominant oscillation content. However, even in this regime LineFit and COG provide the closest match to the truth spectrum in both peak power and overall spectral shape, while FFT-phase and Poly6w show slightly larger deviations and the parabolic method performs worst in this example.

For line 6, the differences become qualitatively important. LineFit reproduces both the locations and amplitudes of the truth RGWS peaks, with only minor residual mismatch at the highest-frequency components. In contrast, the four fast estimators miss or distort key spectral content: while some peak locations are recovered, their associated power is often biased, and the power enhancement near ~ 3.2 mHz is not reproduced. At higher frequencies (e.g., near ~ 10 mHz) the spectra show shifted peaks and additional power bias. This outcome is consistent with the time-domain behaviour in Figure 4 and demonstrates that centre instabilities in complex-core time steps can translate directly into missing or mislocated oscillation power, thereby biasing subsequent wave analyses, including phase/coherence and mode-identification diagnostics.

4 Discussion and conclusion

We have presented an adaptive multi-line fitting methodology (LineFit) aimed at producing stable line-core intensity and LOS-velocity time series from dense spectral windows in which line profiles evolve rapidly and nonlinearly. The motivation is practical but scientifically consequential: discontinuities and occasional

misidentification of the line core can inject artefacts into time-series diagnostics, bias dominant-frequency estimates, and possibly degrade phase/coherence measurements that underpin MHD-wave inferences. This challenge is becoming increasingly prominent as modern solar spectrographs and IFUs (e.g., DKIST/DL-NIRSP and ViSP, Sunrise III/SUSI and SCIP, and FRANCIS) routinely deliver crowded windows containing many lines per spatial sample and per time step.

To isolate performance in a controlled and fully reproducible way, we benchmarked LineFit using a realistic synthetic near-UV time series with unambiguous ground truth for the time-dependent line centres. The synthetic testbed was constructed to explicitly stress the main failure modes encountered in practice: evolving asymmetries, close blends, and intermittent split-core (reversal) morphologies in which a core brightening appears relative to the surrounding absorption depression. Against four widely used fast estimators (parabolic minimum, COG of absorption depression, Fourier-phase centroid, and a weighted polynomial core fit), LineFit consistently recovered accurate centres in representative time steps (Figure 3), while also producing the most robust LOS-velocity time series in the presence of multi-lobed cores (Figure 4). Across the full time series, several of the fast methods performed competitively for many isolated or moderately deformed lines, and in our testbed the COG estimator in particular was often among the best-performing baselines for simple absorption profiles. However, in the key stress case (an intermittently split-core line with strong core-structure variability) LineFit was the only method that consistently maintained agreement with the truth time series and avoided large excursions in the recovered velocity.

A central point of this work is that centre-tracking stability should be assessed not only by per-time-step residuals but also by the downstream wave products used in MHD studies. When the derived velocity contains sporadic mis-tracking events, the resulting time series can acquire artificial power and distort spectral diagnostics. This is clearly demonstrated by the RGWS: for a representative blended but single-core line all methods broadly reproduce the main power enhancements, whereas for the split-core case the fast estimators systematically miss or misplace key spectral peaks and misestimate power compared with the truth, while LineFit preserves both the peak structure and the power distribution, with only minor differences at a few frequencies (Figure 5). These differences are precisely the kinds of discrepancies that matter for multi-height wave and dynamics studies across the solar atmosphere.

The synthetic benchmark provides the key advantage that the true line centres and velocities are known exactly, allowing direct measurement of centre residuals, velocity residuals, and propagated wave-diagnostic errors. In real observations, such absolute residuals are generally not accessible, and performance must instead be assessed through internal consistency, temporal stability, resistance to centre-hopping, and the plausibility and continuity of downstream products such as velocity spectra and wavelet diagnostics. We therefore view the synthetic tests as a controlled validation benchmark, while the practical value of LineFit in real data is demonstrated primarily by its stable behaviour in dense-window applications where complex and evolving profiles are present. Additional real-data effects such as

⁵ <https://GitHub.com/WaLSAteam/WaLSAtools>

scattered light, residual flat-field structure, correlated noise, and wavelength-calibration uncertainties may modify the absolute error scale, but they do not change the core motivation of the method, namely, to reduce unstable centre tracking in precisely those cases where such instabilities would most strongly bias time-series diagnostics.

While the paper focuses on a synthetic benchmark, LineFit has already been applied to dense-spectral-window observations in real-data studies and contexts: it has been used in the analysis of Sunrise III/SUSI near-UV spectra containing more than 100 spectral lines for active-region wave diagnostics (Jafarzadeh et al., 2026b; a), and it has been used to extract line-core parameters in a FRANCIS integral-field analysis (Chambers et al., 2026). In such real-data applications, the most relevant indicators are not absolute residuals to a known truth, but whether the extracted line-core products remain temporally stable, physically plausible, and sufficiently robust for downstream wave and dynamics analysis. These applications demonstrate that the method is practical for novel, information-rich datasets and that it integrates naturally into wave-analysis workflows. We note, however, that the occurrence frequency of multi-lobed or reversal-like profiles in real data is strongly dataset-, line-, and target-dependent; such morphologies occur often enough in dynamic solar spectroscopy to warrant explicit handling when stable time-series extraction is required. At the same time, we emphasise that fast estimators remain valuable: for large-volume datasets and for lines that remain single-lobed and isolated, methods such as COG or Fourier-phase proxies can be attractive due to their computational efficiency, and they are used in contemporary analyses where their assumptions are well matched to the data (e.g., Poly6w in Lagg et al., 2026; FFT-phase approaches in Grant et al., 2026). Our results therefore support a pragmatic view: robust adaptive fitting is most critical for the subset of lines, locations, and times where the profiles become complex—often the same regimes that carry the strongest dynamical signatures—while fast estimators can provide efficient performance elsewhere, provided their limitations are recognised and quality control is applied.

All code required to reproduce the synthetic testbed and figures (including the LineFit implementation) is available via the WaLSALib repository (see Section 2.3). The current implementation is intentionally modular and can be extended to include additional profile families, morphology-specific quality flags, and instrument- or line-dependent constraints as required by particular datasets. A natural next step is to export additional shape diagnostics (e.g., proxy widths, asymmetry measures, and core-structure metrics); although width estimation is non-trivial for deep lines whose wings do not return to the local continuum, it can be handled using robust, locally normalised definitions and reported with quality flags.

Looking forward, the increasing scale of next-generation, high-cadence slit-based spectroscopic and IFU datasets makes acceleration strategies increasingly important. As a proof-of-principle, we trained a lightweight supervised convolutional neural network (CNN; Lecun et al., 1998; Goodfellow et al., 2016) emulator on the synthetic testbed, using LineFit outputs as training labels (ground truth was withheld from training and used only for verification). The emulator predicts the residual line-centre shift relative to a robust coarse centre computed from each

spectrum, reducing the dynamic range the network must learn and decoupling coarse localisation from fine-structure refinement. A hybrid inference mode routes only morphologically complex profiles—identified by a split-core/multi-peak morphology flag and a Monte Carlo dropout uncertainty estimate—to LineFit as a fallback, while the emulator handles the remainder. On the reserved validation set (i.e., the final 20% of the time series, strictly withheld from training), the emulator achieves a median absolute centre error of 0.049 pm and a 95th-percentile error of 0.245 pm relative to the ground truth, with 87% of validation samples agreeing with LineFit to within 0.1 pm. The hybrid fallback rate is 2.1% of all samples (applied only to the stress-case line), and wave diagnostics computed from the emulator and hybrid outputs are well preserved relative to the LineFit reference. The hybrid speed-up relative to running LineFit across the full sequence is approximately three orders of magnitude on the hardware used here, with the emulator component alone approaching five orders of magnitude; the operationally relevant hybrid figure accounts for the LineFit fallback cost on the small flagged fraction (specific numbers and timing methodology are given in Supplementary Section S2). Such a hybrid strategy can preserve the physical transparency and traceability of parametric fitting while achieving the throughput required for very large field-of-view and high-cadence time-series datasets. This proof-of-principle is intentionally limited in scope, and full development—including larger training sets, cross-validation, and instrument-specific tuning—is left for future work. More broadly, we expect that robust, adaptive extraction of line-core products will remain a critical enabling step for wave and dynamics studies using dense-window spectroscopy, particularly when multi-line diagnostics are used to infer energy transport and mode coupling across atmospheric layers.

Data availability statement

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession number(s) can be found below: <https://github.com/WaLSAteam/WaLSALib>.

Author contributions

SJ: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Project administration, Resources, Software, Supervision, Validation, Visualization, Writing – original draft, Writing – review and editing. DJ: Conceptualization, Formal Analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – review and editing. MS: Conceptualization, Formal Analysis, Investigation, Methodology, Software, Validation, Writing – review and editing. PK: Conceptualization, Formal Analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – review and editing. GC: Conceptualization, Formal Analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – review and editing. SG: Formal Analysis, Software, Validation, Visualization, Writing – review and editing. MB: Formal Analysis, Software, Validation, Writing – review and

editing, TD: Formal Analysis, Software, Validation, Writing – review and editing.

Funding

The author(s) declared that financial support was received for this work and/or its publication. UK Science and Technology Facilities Council (STFC) consolidated grants ST/T00021X/1 and ST/X000923/1. UK STFC PATT Travel Grant UKRI372. Leverhulme Trust Research Project Grant RPG-2019-371. UK Space Agency National Space Technology Programme grant SSc-009. Research Council of Norway project no. 262622. The Royal Society award no. Hooke18b/SCTM. NASA grants 19-HSODS-004 and 21-SMDSS21-0047. International Space Science Institute (ISSI) Team 502.

Acknowledgements

SJ, DJ, and GC wish to thank the UK Science and Technology Facilities Council (STFC) for the consolidated grants ST/T00021X/1 and ST/X000923/1, alongside the PATT Travel Grant UKRI372. DJ acknowledges support from the Leverhulme Trust via the Research Project Grant RPG-2019-371. DBJ and SDTG also acknowledge funding from the UK Space Agency via the National Space Technology Programme (grant SSc-009). SJ acknowledges support from the Rosseland Centre for Solar Physics, University of Oslo, Norway and the Max Planck Institute for Solar System Research, Germany. MB acknowledges that this publication was produced while attending the PhD programme in PhD in Space Science and Technology at the University of Trento, Cycle XXXIX, with the support of a scholarship financed by the Ministerial Decree no. 118 of 2nd March 2023, based on the NRRP - funded by the European Union - NextGenerationEU - Mission 4 “Education and Research”, Component 1 “Enhancement of the offer of educational services: from nurseries to universities” - Investment 4.1 “Extension of the number of research doctorates and innovative doctorates for public administration and cultural heritage” - CUP E66E23000110001. We wish to acknowledge scientific discussions with the Waves in the Lower Solar Atmosphere (WaLSA; www.WaLSA.team) team, which has been supported by the Research Council of Norway (project no. 262622), The Royal Society (award no.

Hooke18b/SCTM; Jess et al., 2021), and the International Space Science Institute (ISSI) in Bern through ISSI International Team project 502 “WaLSA: Waves in the Lower Solar Atmosphere at High Resolution”.

Conflict of interest

The author(s) declared that this work was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

The author PK declared that they were an editorial board member of Frontiers at the time of submission. This had no impact on the peer review process and the final decision.

Generative AI statement

The author(s) declared that generative AI was not used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher’s note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fspas.2026.1843842/full#supplementary-material>

References

- Asensio Ramos, A., and Díaz Baso, C. J. (2019). Stokes inversion based on convolutional neural networks. *A&A* 626, A102. doi:10.1051/0004-6361/201935628626A.102A
- Bate, W., Jess, D. B., Grant, S. D. T., Hillier, A., Skirvin, S. J., Van Doorselaere, T., et al. (2024). Unveiling the true nature of plasma dynamics from the reference frame of a superpenumbral fibril. *ApJ* 970, 66. doi:10.3847/1538-4357/ad4d9797066B
- Beckers, J. M., and Brown, T. M. (2013). The history of the fourier tachometer. In *Fifty Years of Seismology of the Sun and Stars*, eds K. Jain, S. C. Tripathy, F. Hill, J. W. Leibacher, and A. A. Pevtsov vol. 478 (San Francisco, CA: Astronomical Society of the Pacific), 93B.
- Berretti, M., Stangalini, M., Verth, G., Fedun, V., Jafarzadeh, S., Jess, D. B., et al. (2025). Umbral oscillations in the photosphere: a comprehensive statistical study. *A&A* 697, A156. doi:10.1051/0004-6361/202453176697A.156B
- Borrero, J. M., Tomczyk, S., Kubo, M., Socas-Navarro, H., Schou, J., Couvidat, S., et al. (2011). VFISV: very fast inversion of the stokes vector for the helioseismic and magnetic imager. *Sol. Phys.* 273, 267–293. doi:10.1007/s11207-010-9515-6
- Bouchy, F., Pepe, F., and Queloz, D. (2001). Fundamental photon noise limit to radial velocity measurements. *A&A* 374, 733–739. doi:10.1051/0004-6361:20010730374.733B
- Calchetti, D., Stangalini, M., Jafarzadeh, S., Valori, G., Albert, K., Albelo Jorge, N., et al. (2023). Spectropolarimetric investigation of magnetohydrodynamic wave modes in the photosphere: first results from PHI on board solar orbiter. *A&A* 674, A109. doi:10.1051/0004-6361/202245826674A.109C
- Campbell, R. J., Mathioudakis, M., Quintero Noda, C., Keys, P. H., and Orozco Suárez, D. (2025). Application of deep learning to the classification of stokes profiles: from the quiet sun to sunspots. *ApJ* 988, 9. doi:10.3847/1538-4357/addb49988.9C
- Chambers, G., Jess, D. B., Jafarzadeh, S., Berretti, M., Grant, S. D. T., Stangalini, M., et al. (2026). Standing oscillations in a resonant sunspot atmosphere captured by integral field spectroscopy. *Front. Astronomy Space Sci. Prep.*
- Couvidat, S., Rajaguru, S. P., Wachter, R., Sankarasubramanian, K., Schou, J., and Scherrer, P. H. (2012). Line-of-Sight observables algorithms for the helioseismic and magnetic imager (HMI) instrument tested with interferometric bidimensional

- spectrometer (IBIS) observations. *Sol. Phys.* 278, 217–240. doi:10.1007/s11207-011-9927-y
- de la Cruz Rodríguez, J., Leenaarts, J., Danilovic, S., and Uitenbroek, H. (2019). STiC: a multiatom non-LTE PRD inversion code for full-stokes solar observations. *A&A* 623, A74. doi:10.1051/0004-6361/201834464623A.74D
- De Moortel, I., and Nakariakov, V. M. (2012). Magnetohydrodynamic waves and coronal seismology: an overview of recent results. *Philosophical Trans. R. Soc. Lond. Ser. A* 370, 3193–3216. doi:10.1098/rsta.2011.0640
- De Pontieu, B., Title, A. M., Lemen, J. R., Kushner, G. D., Akin, D. J., Allard, B., et al. (2014). The interface region imaging spectrograph (IRIS). *Sol. Phys.* 289, 2733–2779. doi:10.1007/s11207-014-0485-y
- De Pontieu, B., Martínez-Sykora, J., Testa, P., Winebarger, A. R., Daw, A., Hansteen, V., et al. (2020). The multi-slit approach to coronal spectroscopy with the multi-slit solar explorer (MUSE). *ApJ* 888, 3. doi:10.3847/1538-4357/ab5b03888.3D
- Duckenfield, T. J., Jess, D. B., Morton, R. J., and Jafarzadeh, S. (2025). Determining the polarization of a coronal standing kink oscillation using spectral imaging techniques with CoMP. *ApJ* 982, 202. doi:10.3847/1538-4357/adb8d6982.202D
- Feller, A., Gandorfer, A., Grauf, B., Hölken, J., Iglesias, F. A., Korpi-Lagg, A., et al. (2025). The sunrise ultraviolet spectropolarimeter and imager: instrument description. *Sol. Phys.* 300, 65. doi:10.1007/s11207-025-02471-765F
- Figueira, P. (2018). “Deriving high-precision radial velocities,” in *Asteroseismology and Exoplanets: Listening to the Stars and Searching for New Worlds*. Editors T. L. Campante, N. C. Santos, and M. J. P. F. G. Monteiro, 49, 181. doi:10.1007/978-3-319-59315-9_1049.181F
- Figueira, P., Santos, N. C., Pepe, F., Lovis, C., and Nardetto, N. (2013). Line-profile variations in radial-velocity measurements. Two alternative indicators for planetary searches. *A&A* 557, A93. doi:10.1051/0004-6361/201220779557A.93F
- Frutiger, C., Solanki, S. K., Fligge, M., and Bruls, J. H. M. J. (2000). Properties of the solar granulation obtained from the inversion of low spatial resolution spectra. *A&A* 358, 1109–1121.
- Gilchrist-Millar, C. A., Jess, D. B., Grant, S. D. T., Keys, P. H., Beck, C., Jafarzadeh, S., et al. (2021). Magnetoacoustic wave energy dissipation in the atmosphere of solar pores. *Philosophical Trans. R. Soc. Lond. Ser. A* 379, 20200172. doi:10.1098/rsta.2020.0172
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press. Available online at: <http://www.deeplearningbook.org> (Accessed March 10, 2026).
- Grant, S. D. T., Jess, D. B., Stangalini, M., Jafarzadeh, S., Fedun, V., Verth, G., et al. (2022). The propagation of coherent waves across multiple solar magnetic pores. *ApJ* 938, 143. doi:10.3847/1538-4357/ac91ca938.143G
- Grant, S. D. T., Berretti, M., Chambers, G., Jess, D. B., Stangalini, M., Jafarzadeh, S., et al. (2026). Probing oscillations in the lower solar atmosphere with the francis integral field unit. *Front. Astronomy Space Sci. Prep.*
- Hölken, J., van Noort, M., Solanki, S. K., Feller, A., Doerr, H.-P., Cao, W., et al. (2026). Toward solar many-line inversions of high-resolution spectropolarimetric data. *A&A* 705, A220. doi:10.1051/0004-6361/202556890705A.220H
- Iglesias, F. A., and Feller, A. (2019). Instrumentation for solar spectropolarimetry: state of the art and prospects. *Opt. Eng.* 58, 082417. doi:10.1117/1.OE.58.8.082417
- Jafarzadeh, S., Jess, D. B., Stangalini, M., Grant, S. D. T., Higham, J. E., Pessah, M. E., et al. (2025a). *Walsatools - Wave Analysis Tools*. doi:10.5281/zenodo.17569951
- Jafarzadeh, S., Jess, D. B., Stangalini, M., Grant, S. D. T., Higham, J. E., Pessah, M. E., et al. (2025b). Wave analysis tools. *Nat. Rev. Methods Prim.* 5, 21. doi:10.1038/s43586-025-00392-0521J
- Jafarzadeh, S., Jess, D. B., Stangalini, M., Keys, P. H., Grant, S. D. T., Duckenfield, T. J., et al. (2026a). Magnetically structured oscillatory power along an active-region transect in Near-UV sunrise iii/SUSI spectroscopy. *ApJ Prep.*
- Jafarzadeh, S., Jess, D. B., Stangalini, M., Morton, R. J., Felipe, T., Berretti, M., et al. (2026b). Multi-line wave signatures in a sunspot from Near-UV sunrise iii/SUSI observations. *ApJ Prep.*
- Jarolim, R., Molnar, M. E., Tremblay, B., Centeno, R., and Rempel, M. (2025). PINN ME: a physics-informed neural network framework for accurate milne-eddington inversions of solar magnetic fields. *ApJ* 985, L7. doi:10.48550/arXiv.2502.13924985L7J
- Jess, D. B., Morton, R. J., Verth, G., Fedun, V., Grant, S. D. T., and Giagkiozis, I. (2015). Multiwavelength studies of MHD waves in the solar chromosphere. An overview of recent results. *Space Sci. Rev.* 190, 103–161. doi:10.1007/s11214-015-0141-3
- Jess, D. B., Keys, P. H., Stangalini, M., and Jafarzadeh, S. (2021). High-resolution wave dynamics in the lower solar atmosphere. *Philosophical Trans. R. Soc. Lond. Ser. A* 379, 20200169. doi:10.1098/rsta.2020.0169
- Jess, D. B., Grant, S. D. T., Bate, W., Liu, J., Jafarzadeh, S., Keys, P. H., et al. (2023a). The fibre resolved Optical and near-ultraviolet czerny-turner imaging spectropolarimeter (FRANCIS). *Sol. Phys.* 298, 146. doi:10.1007/s11207-023-02237-z
- Jess, D. B., Jafarzadeh, S., Keys, P. H., Stangalini, M., Verth, G., and Grant, S. D. T. (2023b). Waves in the lower solar atmosphere: the dawn of next-generation solar telescopes. *Living Rev. Sol. Phys.* 20, 1. doi:10.1007/s41116-022-00035-620.1J
- Katsukawa, Y., del Toro Iniesta, J. C., Solanki, S. K., Kubo, M., Hara, H., Shimizu, T., et al. (2020). “Sunrise chromospheric infrared SpectroPolarimeter (SCIP) for sunrise III: system design and capability,” in *Ground-Based and Airborne Instrumentation for Astronomy VIII*. Editors C. J. Evans, J. J. Bryant, and K. Motohara, 114470Y. doi:10.1117/12.2561223
- Korpi-Lagg, A., Gandorfer, A., Solanki, S. K., del Toro Iniesta, J. C., Katsukawa, Y., Bernasconi, P., et al. (2025). SUNRISE III: overview of observatory and instruments. *Sol. Phys.* 300, 75. doi:10.1007/s11207-025-02485-175K
- Lagg, A., Woch, J., Krupp, N., and Solanki, S. K. (2004). Retrieval of the full magnetic vector with the he I multiplet at 1083 nm. Maps of an emerging flux region. *A&A* 414, 1109–1120. doi:10.1051/0004-6361:20031643414.1109L
- Lagg, A., Smitha, H. N., Solanki, S. K., Cameron, R., Schou, J., Riethmüller, T. L., et al. (2026). Direct measurement of the height dependence of *p*-mode propagation in the lower solar atmosphere with sunrise. *ApJL*.
- Lecun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *IEEE Proc.* 86, 2278–2324. doi:10.1109/5.726791
- Lites, B., Casini, R., Garcia, J., and Socas-Navarro, H. (2007). A suite of community tools for spectro-polarimetric analysis. *Mem. Soc. Astron. Ital.* 78, 148.
- MacBride, C. D., Jess, D. B., Grant, S. D. T., Khomenko, E., Keys, P. H., and Stangalini, M. (2021). Accurately constraining velocity information from spectral imaging observations using machine learning techniques. *Philosophical Trans. R. Soc. Lond. Ser. A* 379, 20200171. doi:10.1098/rsta.2020.0171
- Milić, I., and van Noort, M. (2018). Spectropolarimetric NLTE inversion code SNAPI. *A&A* 617, A24. doi:10.1051/0004-6361/201833382617A.24M
- Rast, M. P., Bello González, N., Bellot Rubio, L., Cao, W., Cauzzi, G., Deluca, E., et al. (2021). Critical science plan for the daniel K. inouye solar telescope (DKIST). *Sol. Phys.* 296, 70. doi:10.1007/s11207-021-01789-270R
- Riethmüller, T. L., and Solanki, S. K. (2019). The potential of many-line inversions of photospheric spectropolarimetric data in the visible and near UV. *A&A* 622, A36. doi:10.1051/0004-6361/201833379622A.36R
- Ruiz Cobo, B., and del Toro Iniesta, J. C. (1992). Inversion of stokes profiles. *ApJ* 398, 375. doi:10.1086/171862398.375R
- Ruiz Cobo, B., Quintero Noda, C., Gafeira, R., Uitenbroek, H., Orozco Suárez, D., and Páez Mañá, E. (2022). DeSIRE: departure coefficient aided stokes inversion based on response functions. *A&A* 660, A37. doi:10.1051/0004-6361/202140877660A.37R
- Shimizu, T., Imada, S., Kawate, T., Suematsu, Y., Hara, H., Tsuzuki, T., et al. (2021). “The Solar-C (EUVST) mission: the latest status,” in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Editors J.-W. A. den Herder, S. Nikzad, and K. Nakazawa, 114440N. doi:10.1117/12.2560887
- Socas-Navarro, H., de la Cruz Rodríguez, J., Asensio Ramos, A., Trujillo Bueno, J., and Ruiz Cobo, B. (2015). An open-source, massively parallel code for non-LTE synthesis and inversion of spectral lines and Zeeman-induced stokes profiles. *A&A* 577, A7. doi:10.1051/0004-6361/201424860577A7S
- Solanki, S. K., Smitha, H. N., Lagg, A., Gandorfer, A., del Toro Iniesta, J. C., Katsukawa, Y., et al. (2026). Sunrise iii: instrument, mission, data, and first results. *ApJL*.
- Stangalini, M., Jess, D. B., Verth, G., Fedun, V., Fleck, B., Jafarzadeh, S., et al. (2021). A novel approach to identify resonant MHD wave modes in solar pores and sunspot umbrae: *b* – *ω* analysis. *A&A* 649, A169. doi:10.1051/0004-6361/202140429649A.169S
- Stangalini, M., Verth, G., Fedun, V., Aldhafaeri, A. A., Jess, D. B., Jafarzadeh, S., et al. (2022). Large scale coherent magnetohydrodynamic oscillations in a sunspot. *Nat. Commun.* 13, 479. doi:10.1038/s41467-022-28136-8
- Tritschler, A., Rimmele, T. R., Berukoff, S., Casini, R., Kuhn, J. R., Lin, H., et al. (2016). Daniel K. Inouye solar telescope: high-resolution observing of the dynamic sun. *Astron. Nachrichten* 337, 1064–1069. doi:10.1002/asna.201612434.337.1064T
- Tziotziou, K., Tsiropoulou, G., Mein, N., and Mein, P. (2007). Dual-line spectral and phase analysis of sunspot oscillations. *A&A* 463, 1153–1163. doi:10.1051/0004-6361:20066412463.1153T
- Uitenbroek, H. (2003). The accuracy of the center-of-gravity method for measuring velocity and magnetic field strength in the solar photosphere. *ApJ* 592, 1225–1233. doi:10.1086/375736592.1225U