

The forward and inverse problems in time-distance helioseismology

Jason Jackiewicz, Laurent Gizon

Max-Planck-Institut für Sonnensystemforschung, 37191 Katlenburg-Lindau, Germany

E-mail: jackiewicz@mps.mpg.de

Aaron C Birch

Colorado Research Associates, NWRA, 3380 Mitchell Lane, Boulder, CO 80301, USA

Abstract. Time-distance helioseismology is a set of tools for peering into the solar interior. In this paper we discuss and provide examples of the steps that go into current high-resolution time-distance helioseismic analyses. These steps include observations (cross covariances, travel times), modeling of the seismic wavefield for a weakly inhomogeneous solar model, and inversion of the travel times.

The discussion is framed in the context of studying quiet-Sun flows, although the extension to other solar perturbations is straightforward and analogous. The two-plus-one-dimensional (2+1D) inversion procedure implemented here produces maps of vector flows in the near-surface layers of the photosphere. We examine the flows obtained by compromising, or ‘trading off’, between different observation times, spatial resolutions, and noise levels. Also studied is the correlation of the flows at different depths and over different time intervals.

1. Introduction

Time-distance helioseismology [1] is the interpretation of travel-time data measured for helioseismic waves propagating between two points on the solar surface. The travel times contain information about inhomogeneities located near the path of the waves, such as flows, temperature, sound-speed and density perturbations, and the magnetic field. Extracting this information from the travel times requires (1) accurate forward modeling of wave propagation in a realistic solar model and (2) a robust inversion procedure incorporating the noise properties of the measurements.

Recently, we presented a fully-consistent time-distance inversion for near-surface vector flows using a novel 2+1D scheme (see [2] and [3]). The scheme is based on ridge filtering the travel times (described below) instead of the standard phase-speed filtering commonly carried out in time-distance helioseismology. It uses Born sensitivity kernels computed with the same travel-time definition with which the observations were made. Finally, the full noise properties were utilized in the analysis. In those papers, the methods were described and most of the new results were shown. In this paper we will give an overview of the time-distance technique and present some of the difficulties associated with performing high-resolution inversions, as well as some of the intermediate steps that have not been fully discussed elsewhere. In this way, its purpose is to serve as a general overview of time-distance helioseismology, but for a very specific problem.

For further references, a recent paper [4] by Gizon & Thompson provided a general overview of the procedures used in local helioseismology and many of the major obstacles and shortcomings these techniques face. An updated version with some new results is also given in this volume by Birch [5]. These two papers serve as good companions to an in-depth review paper [6] that summarizes the main methods, objectives, and results for all branches of local helioseismology, such as ring-diagram analysis [7], acoustic holography [8], and direct modeling [9]. Global helioseismology, which utilizes the frequencies of the normal modes of the Sun, has been thoroughly reviewed elsewhere [10].

This paper is organized as follows: In § 2 we will talk about the observations and data, in particular the ridge filtering that is carried out. § 3 briefly discusses the forward modeling that is done, and § 4 sketches the basics of the inversion method, including a look at the computational demands. Results of the procedure for flows are given in § 5, and conclusions are drawn in § 6.

2. Time-distance observations

We briefly summarize the basic steps of analyzing solar seismic data, starting from the Doppler measurements, to obtain specific types of travel-time measurements. In this discussion, we focus on the relatively recent technique of filtering the data not with phase-speed filters, but with ‘ridge’ filters.

2.1. Ridge filtering of the data

The first step in typical time-distance analysis is to filter a time series of Doppler images to retain particular oscillation modes. A standard practice in the past decade or so has been to choose mode sets, each with a common central wave phase speed, known as ‘phase-speed’ filtering (e.g. [11], among many). This filtering has been shown to increase signal/noise of the data, and roughly isolates modes that travel a similar path in the solar interior. Instead, we employ here what we now called ‘ridge’ filtering, whereby modes are selected that have a common radial order n , evident as the ‘ridges’ in a power spectrum. The individual ridges are retained and then used separately in further analysis. For instance, in this study we filter and retain the modes having $n = 0$ (surface-gravity f modes) and also those with $n = 1, 2, 3, 4$ (acoustic p modes).

Let $\phi(\mathbf{x}, t)$ be a time series of line-of-sight velocity data measured at the solar surface, where $\mathbf{x} = (x, y)$ is the horizontal position vector and t is time. The filtered data is obtained by multiplication of a ridge filter F_n for ridge n with the Fourier transform of the data:

$$\phi_n(\mathbf{k}, \omega) = F_n(\mathbf{k}, \omega)\phi(\mathbf{k}, \omega), \quad (1)$$

where \mathbf{k} is the horizontal wavevector, ω is the angular frequency, and $\phi(\mathbf{k}, \omega)$ is the three-dimensional Fourier transform of $\phi(\mathbf{r}, t)$. The filters are computed from studying a mean power spectrum. For each wavenumber, F_n is 1 along the width (in frequency) of ridge n and falls off exponentially to zero before encountering any significant mode power in adjacent ridges.

We apply equation (1) to a full-disk data set from MDI/SOHO spanning the dates January 20-28, 2002, using five different filters ($n = 0, 1, 2, 3, 4$). In Figure 1 we show an example power spectrum $|\phi_{n=1}(|\mathbf{k}|, \omega)|^2$ filtered along the p₁ ridge from a continuous six-hour time series within that data set. Note that wavenumbers smaller than $200/R_\odot$ are filtered out because they are dominated by solar convective noise.

2.2. Cross-covariance

The next step is to compute the cross-covariance function from the ridge-filtered Doppler velocities ϕ_n at two locations \mathbf{x}_1 and \mathbf{x}_2 on the surface that are separated by a distance $\Delta = \|\mathbf{x}_2 - \mathbf{x}_1\|$. For each radial order n , the cross-covariance is given by

$$C_n(\mathbf{x}_1, \mathbf{x}_2, t) = \int \phi_n(\mathbf{x}_1, t')\phi_n(\mathbf{x}_2, t' + t)dt', \quad (2)$$

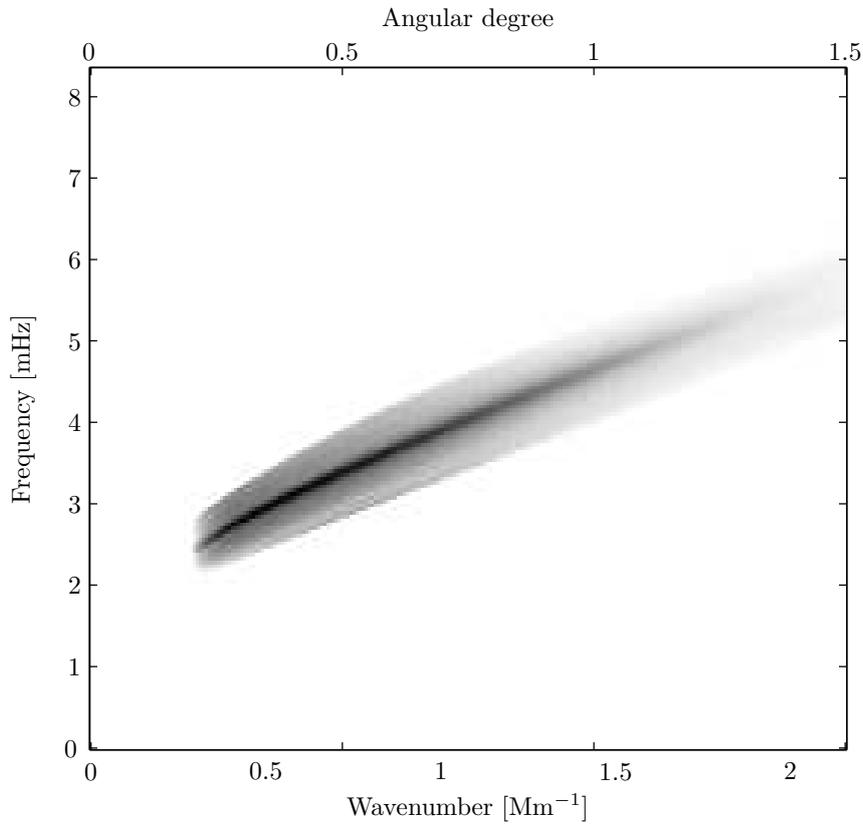


Figure 1. Example ridge-filtered power spectrum of full-disk MDI Doppler data. Shown here is the p_1 ridge. The power spectrum is averaged over the direction of the horizontal wavevector. The limits of the axes are the spatial and temporal Nyquist frequencies. The gray scale is logarithmic.

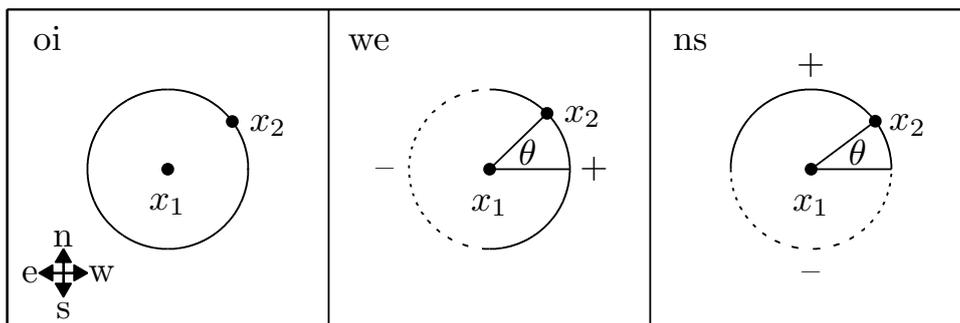


Figure 2. A sketch of the three point-to-annulus geometries considered here for the cross-covariance computation. The wavefield at the central point \mathbf{x}_1 is cross correlated with the wavefield at point \mathbf{x}_2 averaged over the respective annulus of radius Δ . The solid arcs denote a positive (+) weighting and the dashed arcs denote a negative (-) weighting. θ is the angle \mathbf{x}_2 makes with the x axis, which points in the east-west direction. The ‘we’ covariance is weighted by $\cos \theta$, and ‘ns’ is weighted by $\sin \theta$. The legend in the left panel illustrates the directional coordinate system used for the Sun.

where t is the time lag, positive for waves propagating from \mathbf{x}_1 to \mathbf{x}_2 and negative for waves propagating in the opposite direction. As is standard [11], various spatial averaging is done to C to reduce noise and to enhance sensitivity to different flow geometries. We choose three configurations for this averaging, illustrated in Figure 2. These are denoted by ‘oi’ (out minus in), ‘we’ (west minus east), and ‘ns’ (north minus south). The ‘oi’ covariances combine the wave signal at a point with the wave signal averaged over a concentric annulus about that point of radius Δ . The ‘we’ and ‘ns’ covariances are averages weighted by the $\cos \theta$ and $\sin \theta$, respectively, as depicted in Figure 2. The ‘oi’ covariances are mostly sensitive to the horizontal divergence of the flows, while the ‘we’ and ‘ns’ covariances are mostly sensitive to the x and y components of the flows respectively. We compute C for 20 different Δ , ranging from about 1.5 to 30 Mm. Along with the five ridges, three point-to-annulus geometries, and 28 six-hour long time series’, we compute 8400 cross-covariance functions.

2.3. Travel times

Travel times are computed from the cross covariances defined in equation (2) using the linearized definition developed in [12; 13]. Another commonly used method is to fit a Gaussian wavelet to the cross covariances using several fitting parameters. The expression used here for the travel time perturbation τ can be written as

$$\tau_i(\mathbf{r}) = \int W_i(t) [C_i(\mathbf{r}, t) - C_i^{\text{ref}}(\mathbf{r}, t)] dt, \quad (3)$$

where C^{ref} is a reference cross covariance derived from a solar model (as explained in the next section) and the index i denotes a particular choice of filtering, measurement type, and distance Δ . The quantity W is a weighting function proportional to the time derivative of C^{ref} [12]. This definition has proven to be quite robust for noisy data [14]. Each travel-time measurement obtained corresponds to a particular ridge, averaging geometry, and distance Δ . For this study (where we measure travel-time *differences*), perturbations to the travel times provide information about flows.

3. Forward modeling

Forward modeling consists of computing the wave field for a given solar model. The wave field is then connected to an observable, which in time-distance helioseismology is the travel time. In general, this is a very difficult problem and inevitably certain approximations are used. Here, we assume that the solar model is weakly inhomogeneous. The result of this approximation provides two important quantities: travel-time sensitivity kernels and the statistical noise properties of the travel times, which we now briefly discuss.

3.1. Sensitivity kernels

We compute the wave field in the first Born approximation, which establishes a linear relationship between the scattered wavefield and some small-amplitude perturbation in the Sun [12; 13; 15], which in our case are local flows. The flow perturbation \mathbf{u} can then be shown to be connected to the travel times through the relation

$$\tau_i = \iiint \mathbf{K}_i(\mathbf{r}) \cdot \mathbf{u}(\mathbf{r}) d^3\mathbf{r}, \quad (4)$$

where i may denote any type of travel-time measurement, ridge, or distance, \mathbf{K}_i the corresponding vector-valued travel-time sensitivity kernel, $\mathbf{r} = (\mathbf{x}, z)$ the three-dimensional position coordinate, and the quantity on the rhs is understood to be a convolution. In words, the function \mathbf{K}_i gives the sensitivity of a travel-time measurement τ_i (the observable) to a local

flow perturbation $\mathbf{u}(\mathbf{r})$. The two steps required to establish equation (4) are (1) to find a linear relationship between perturbations to the travel time due to changes in the cross covariance (cf. eq. [3]), and (2) to relate the cross covariance to small changes in some solar property, e.g., flows. Besides flows, where example kernels are derived and provided in [3; 16], time-distance Born kernels have been computed for source and damping perturbations [12], as well as sound-speed perturbations [15].

The kernels we use are point-to-annulus kernel functions, constructed in the same way as described in § 2.2. The kernels are obtained for each value of n . Because we have ridge-filtered kernels, it can be shown that they approximately separate in the horizontal and vertical directions [3]. This property motivates us to carry out a 2+1D inversion, as described in § 4.

3.2. Travel-time noise covariance

Due to the stochastic nature of the excitation mechanism of solar oscillations, different travel-time measurements may be correlated in space and time. Taking account of this correlation is essential for proper interpretation of the information contained in the travel times. Furthermore, the original noise gets propagated and amplified upon carrying out inversion procedures, and therefore must be accurately quantified.

In a detailed study [13], Gizon & Birch developed a way to model the noise covariance of quiet-Sun travel times based on utilizing the same model components that enter into the computation of the sensitivity kernels; specifically, the model cross-covariance function C^{ref} . From the recipe discussed in [13] we compute all of the appropriate noise-covariance matrices of the travel times. This allows us to assign errors to every flow map computed in the inversion. Example plots of the covariance matrices are shown in [2].

4. Inverse methods

All of the inputs for any standard inversion have now been computed. The goal is then, given the measurements, the kernels, and the noise, to find the three-dimensional flow field in the near-surface layers of the Sun.

4.1. Optimally localized averaging

We choose to implement a linear inversion technique known as subtractive optimally localized averaging (SOLA) [17]. It is different from another commonly used technique, regularized least squares (RLS), in that it does not involve any ‘fitting’ of the data.

The basic principle of the three-dimensional version of the inversion is to construct a quantity known as an ‘averaging’ kernel by combining the sensitivity kernels in a particular way using a set of inversion coefficients or ‘weights.’ The ideal averaging kernel for the flow component to be estimated should roughly resemble the shape of a ball located somewhere in the solar interior. If a set of coefficients is found that can do this, then the travel times can be combined properly in a way that gives the value of the flow in that local region of the ball. The diameter of the ball gives a good idea of the spatial resolution of the final flow estimation. Another consideration is to quantify the ‘cross talk’ among the flow components. The cross talk is the amount of sensitivity in the inversion to the components of the flow which are not being estimated. In general, the horizontal and vertical flows are not independent. Minimizing the cross talk is done in practice by imposing constraints on the integrals of the averaging kernels.

The coefficients are found by inverting a generally large matrix. An optimal set of coefficients is chosen by compromising between how ‘ball-like’ (and small) the averaging kernel is and the level of the noise. The more averaging that is done, the more the noise levels are increased. There is a certain amount of subjectivity when trying to find an appropriate balance between the two.

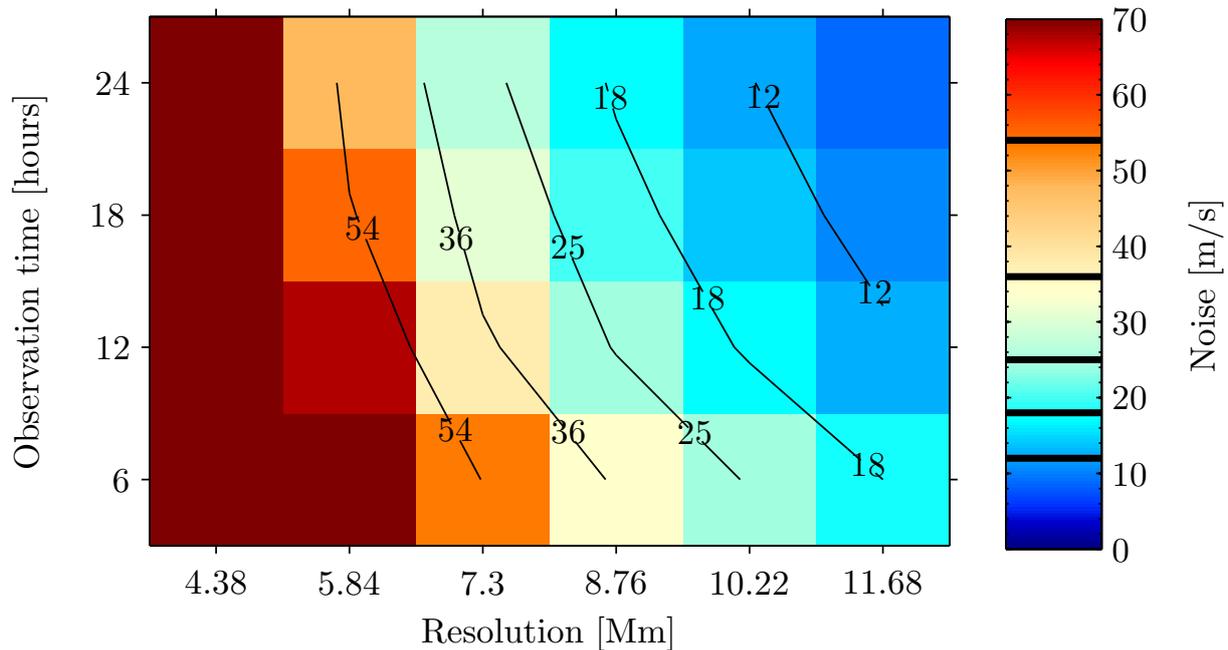


Figure 3. Noise estimation for an inversion 1 Mm beneath the surface for v_x for different resolutions and observation times T . The solid lines are contours of constant noise. Each block of the matrix represents the noise level of a corresponding flow map obtained from separate inversions. A subset of these flows maps are compared in Figure 4. Note that the contour levels are also marked on the color bar.

As mentioned in § 3, we carry out the inversion procedure in two separate steps due to the approximate separability of the ridge-filtered sensitivity kernels in the horizontal and vertical directions. In other words, we first do a 2D inversion in the horizontal coordinates to obtain an averaging kernel in the shape of a circle and constant in depth. This step is done for each ridge. The 1D inversion then combines the individual averaging kernels for each ridge to localize it about a target depth, thus completing the ball structure (see [3] for full details). One reason for doing this is that the trade-off decisions between resolution and noise are carried out separately and (in some way) independently.

4.2. Inversion decisions

As already stressed, an important component of the procedure is to decide on acceptable resolution and noise levels, given the available observations. Therefore, it is important to compute inversions for a range of resolutions and carefully examine the results. We remind the reader that this is done in the absence of any data and only the weights are needed (a nice aspect of OLA-type inversions). What one typically does is generate a trade-off curve (‘L’ curve).

Another parameter is the observation time T . It was shown in [13] that the noise in the travel times falls off as $T^{-1/2}$. One strategy is to settle for a level of noise one finds acceptable and then study the resolution and amount of data needed to be at or below this noise level. An example of a helpful tool for this decision is shown in Figure 3. This matrix is derived from inversions for v_x targeted to be sensitive at 1 Mm beneath the solar surface. To remain on a contour of constant error, one can achieve a high resolution by observing for a longer time, or, with a smaller amount of data available, one would then have to forfeit some resolution. In § 5 we will show flow maps corresponding to the inversions in this figure.

4.3. Computational concerns

For the 2D inversions explained above, an optimal set of weights is chosen for each ridge, distance, geometry, and resolution, for a total of 2400 sets. To obtain each set requires the inversion of a roughly 20 000 by 20 000 element matrix. Instead of direct matrix inversion, we have utilized the `mldivide` routine in `Matlab` that solves the system of equations $\mathbf{Ax} = \mathbf{b}$ using gaussian elimination. This ‘matrix left division’ operation, written as $x = \mathbf{A} \backslash \mathbf{b}$, is much more efficient and numerically accurate than solving $x = \mathbf{A}^{-1}\mathbf{b}$. The computation has been done on a dual processor machine with shared memory. The memory becomes an issue since about 10 – 20 Gb are needed to store the matrices. To compute all of the sets of weights, the computation time is about 2-3 weeks. The 1D inversion is quite fast comparatively. The final step, averaging the travel times to derive the flows using the computed weights, is a small fraction of the total computation time.

5. Results

We now present some results for quiet-Sun flows. Inversions are computed for the horizontal components of the flow, v_x and v_y , for several depths. For the vertical component, v_z , only a 2D inversion is implemented here (each ridge has some implied depth, that will for now serve as a proxy). A more complete quiet-Sun flow analysis is found in [3]: here, we are interested in some of the ideas not fully discussed in that paper. We focus on two results: (1) the differences in the maps themselves for various resolutions, observation times and noise, and (2) the spatial correlation of flows in the quiet Sun at different depths and times. Both of these are important for understanding how well the solar signal can be recovered, and to make sure that what is being measured is not just noise.

5.1. Influence of different resolutions and observing durations on flows with similar noise levels
The trade off between noise, resolution, and observation time from an inversion was studied in Figure 3. It is helpful to see how the different parameters are manifest in maps of quiet-Sun flows.

Figure 4 shows a comparison of example horizontal flow maps for different parameter values at a depth of 1 Mm. The arrows are derived from inverting for v_x and v_y . The $\sim 100 \text{ Mm}^2$ patch of Sun is relatively quiet as can be seen by the background color scale taken from line-of-sight MDI magnetograms that have been averaged over the observation time T . Each row of Figure 4 corresponds to two maps with very similar noise levels (roughly falling on one of the contour lines in the matrix of Figure 3), but obtained from inversions using different lengths of time series’ and having various resolutions. The spatial resolution of each map can be surmised from the gaussian averaging kernel plotted in the lower right of each panel, where the FWHM of the gaussian is outlined in black. In all plots supergranular structures are recognizable, with outflow speeds near 250 m s^{-1} , and sizes of about 30 Mm. Also evident is the magnetic flux concentration mostly at the boundaries of the supergranules.

For each row, the left panel has a low resolution with shorter T (corresponding to the bottom of a contour line in Figure 3), and the right panel has a higher resolution but larger T (corresponding to the top of a contour line in Figure 3).

As this inversion is for the very near surface of the Sun, the dominant modes contributing to the inversion are the f modes, whose wavelength is about 5 Mm. The inversions have resolutions above this value, except the bottom row. It seems from this plot that at a certain noise level there is not much difference in the flow estimations to small changes in the resolution or observation time. The correlations between left and right plots are all quite strong (> 0.75). It comes down to how much data one might have, and how well one would like to be able to examine fine structure.

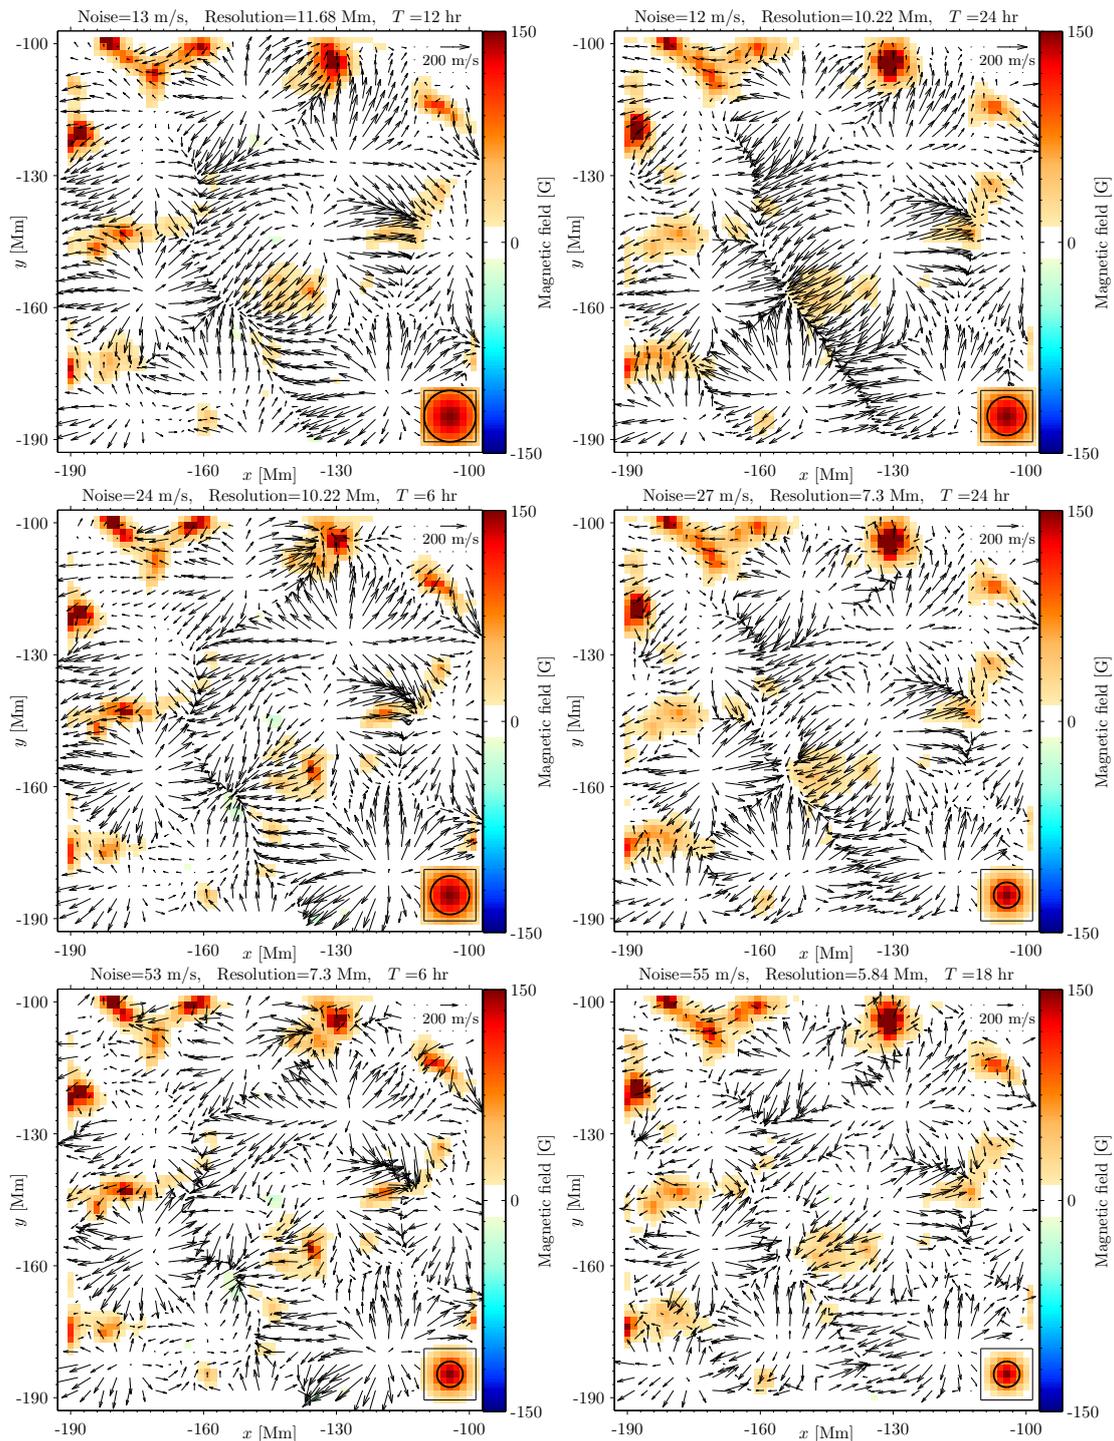


Figure 4. Comparison of horizontal flow maps at 1 Mm beneath the surface for different values of observation time (T) and resolution, but similar noise levels in each row. Each panel on the left is from an inversion with less data and higher resolution than the corresponding panel on the right. The background color scale is the line-of-sight truncated magnetic field. The quantity in the bottom right of each panel is the averaging kernel, whose FWHM is outlined by the circle. The correlation coefficient of the two maps in the top row is about 0.91, the middle row is 0.75, and the bottom row is 0.76.

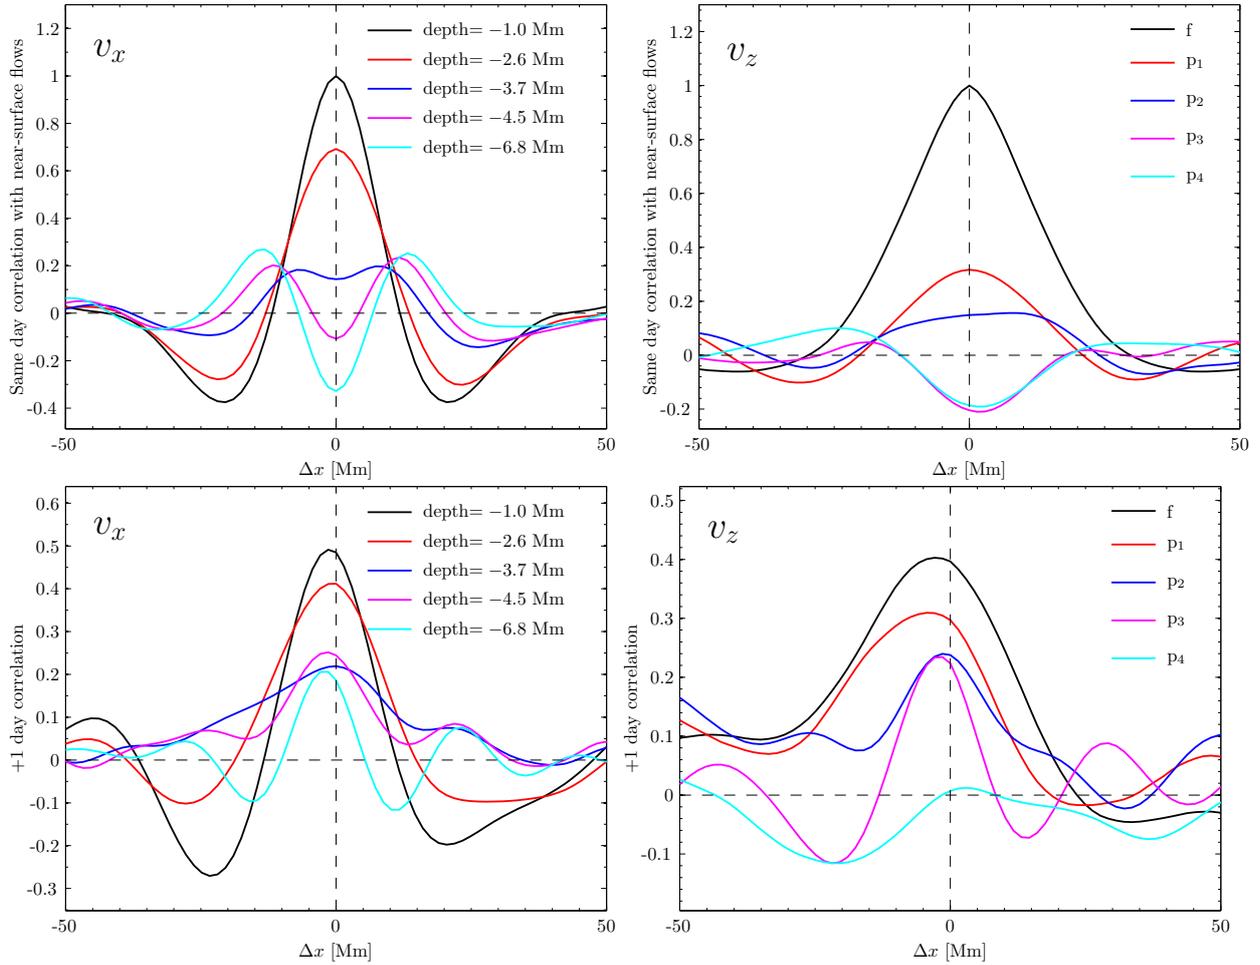


Figure 5. Correlation of supergranular flows over 24 hr for two different cases. Each panel shows cuts through spatial cross-correlation maps along $\Delta y = 0$. The figures in the top row show the correlation of the near-surface flows with those at different depths for the same day for v_x and v_z , respectively. The figures in the bottom row show the correlation of flows at a specific depth with the flows one day later at the same depth, again for v_x and v_z . The correlations have been obtained from averaging 7 continuous days of data.

5.2. Average flow correlations

To study average flow correlations over time and depth, we take the same quiet-Sun region as in Figure 4 and construct a series of 1 day flow maps (of v_x and v_z) for each of seven continuous days. A correlation ‘map’ in Δx and Δy is then computed by spatially cross correlating any two maps. The strongest signal contributing from these maps for 1 day of data should be due to supergranulation [13].

The above procedure is carried out for two cases. The first case is for studying ‘same day’ correlations, whereby we cross correlate the near-surface map with successively deeper layers for each day and study the 7-day average. For v_x , the near surface denotes a depth of 1 Mm. For v_z , we instead correlate the f-mode ridge flows with successive acoustic-ridge flows, since we have not yet fully implemented and tested the depth inversion for the vertical velocity. Slices through the resulting correlation maps along $\Delta y = 0$ are shown in the top panels of Figure 5. Since the supergranules are primarily responsible for the correlation, the width of the curves for the case

of v_x near the surface provides a rough estimate of their horizontal size, about 20–30 Mm. This roughly agrees with an inspection by eye of the structures in Figure 4. The negative sidelobes are due to neighboring supergranules. The correlation of v_x decreases with depth, and a small anti-correlation with the near surface occurs at about 5 Mm, which is not significant. As noted elsewhere, it is difficult to conclude anything regarding the supergranules at this depth due to the high noise levels [3; 18]. As for v_z , a decrease in the correlation with depth is also observed.

For the second case, a map from one day is cross correlated with a map at the same depth of the successive 24-hour period. If only noise that was being measured, correlations from one day to the next would vanish. However, this is not the case, as seen in the plots shown in the bottom two panels of Figure 5. Note that the correlation maxima are not at $\Delta x = 0$, most likely because supergranules (the strongest contributors to the signal) do not rotate at the tracking rate used for this analysis.

6. Conclusions

Time-distance helioseismology is a useful technique to probe the upper convection zone. In this paper we have shown results of a fully-consistent OLA analysis to infer near-surface horizontal and vertical flows at a high spatial resolution. Supergranulation is easily resolved and is the dominant signal in these maps. Future studies will aim at imaging not only flows, but also structural inhomogeneities (e.g. density, sound speed, first adiabatic exponent) and the magnetic field. Specific methods will have to be developed in order to properly model the perturbations caused by the magnetic field, which are not small in the near-surface layers.

References

- [1] Duvall Jr T L, Jefferies S M, Harvey J W and Pomerantz M A 1993 *Nature* **362** 430–432 (ADS http://adsabs.harvard.edu/cgi-bin/nph-bib_query?bibcode=1993Natur.362..430D&db_key=AST)
- [2] Jackiewicz J, Gizon L, Birch A C and Thompson M J 2007a *Astronomische Nachrichten* **328** 234–239 (Preprint [astro-ph/0702345](http://arxiv.org/abs/astro-ph/0702345))
- [3] Jackiewicz J, Gizon L and Birch A C 2008 *Solar Physics*, accepted
- [4] Gizon L and Thompson M J 2007 *Astronomische Nachrichten* **328** 204 (ADS <http://adsabs.harvard.edu/abs/2007AN....328..204G>)
- [5] Birch A C 2008 (*This volume*)
- [6] Gizon L and Birch A C 2005 *Living Reviews in Solar Physics* **2** 6 (ADS <http://adsabs.harvard.edu/abs/2005LRSP....2....6G>)
- [7] Hill F 1988 *Astrophys. J.* **333** 996–1013 (ADS <http://adsabs.harvard.edu/abs/1988ApJ...333..996H>)
- [8] Lindsey C and Braun D C 2000 *Solar Phys.* **192** 261–284 (ADS <http://adsabs.harvard.edu/abs/2000SoPh..192..261L>)
- [9] Woodard M F 2002 *Astrophys. J.* **565** 634–639 (ADS http://adsabs.harvard.edu/cgi-bin/nph-bib_query?bibcode=2002ApJ...565..634W&db_key=AST)
- [10] Christensen-Dalsgaard J 2002 *Reviews of Modern Physics* **74** 1073–1129 (Preprint [arXiv:astro-ph/0207403](http://arxiv.org/abs/astro-ph/0207403))
- [11] Duvall Jr T L, Kosovichev A G, Scherrer P H, Bogart R S, Bush R I, de Forest C, Hoeksema J T, Schou J, Saba J L R, Tarbell T D, Title A M, Wolfson C J and Milford P N 1997 *Solar Phys.* **170** 63–73 (ADS http://adsabs.harvard.edu/cgi-bin/nph-bib_query?bibcode=1997SoPh..170..63D&db_key=AST)
- [12] Gizon L and Birch A C 2002 *Astrophys. J.* **571** 966–986 (ADS http://adsabs.harvard.edu/cgi-bin/nph-bib_query?bibcode=2002ApJ...571..966G&db_key=AST)
- [13] Gizon L and Birch A C 2004 *Astrophys. J.* **614** 472–489 (ADS http://adsabs.harvard.edu/cgi-bin/nph-bib_query?bibcode=2004ApJ...614..472G&db_key=AST)

- [14] Roth M, Gizon L and Beck J G 2007 *Astronomische Nachrichten* **328** 215–222 (ADS <http://adsabs.harvard.edu/abs/2007AN....328..215R>)
- [15] Birch A C, Kosovichev A G and Duvall Jr T L 2004 *Astrophys. J.* **608** 580–600 (ADS http://adsabs.harvard.edu/cgi-bin/nph-bib_query?bibcode=2004ApJ...608..580B&db_key=AST)
- [16] Birch A C and Gizon L 2007 *Astronomische Nachrichten* **328** 228 (ADS http://adsabs.harvard.edu/cgi-bin/nph-bib_query?bibcode=2007AN....328..228B&db_key=AST)
- [17] Pijpers F P and Thompson M J 1992 *Astron. Astrophys.* **262** L33–L36 (ADS <http://adsabs.harvard.edu/abs/1992A26A...262L..33P>)
- [18] Braun D C, Birch A C, Benson D, Stein R F and Nordlund Å 2007 *Astrophys. J.* **669** 1395–1405 (Preprint [arXiv:0708.0214](http://arxiv.org/abs/0708.0214))