Magnetic balltracking: Tracking the photospheric magnetic flux[★]

R. Attie and D. E. Innes

Max-Planck-Institut für Sonnensystemforschung, 37077 Göttingen, Germany
 e-mail: attie.raphael@gmail.com

Received 7 July 2014 / Accepted 16 December 2014

ABSTRACT

Context. One aspect of understanding the dynamics of the quiet Sun is to quantify the evolution of the flux within small-scale magnetic features. These features are routinely observed in the quiet photosphere and were given various names, such as pores, knots, magnetic patches.

Aims. This work presents a new algorithm for tracking the evolution of the broad variety of small-scale magnetic features in the photosphere, with a precision equal to the instrumental resolution.

Methods. We have developed a new technique to track the evolution of the individual magnetic features from magnetograms, called “magnetic balltracking”. It quantifies the flux of the tracked features, and it can track the footpoints of magnetic field lines inferred from magnetic field extrapolation. The algorithm can detect and quantify flux emergence, as well as flux cancellation.

Results. The capabilities of magnetic balltracking are demonstrated with the detection and the tracking of two cases of magnetic flux emergence that lead to the brightening of X-ray loops. The maximum emerged flux ranges from 10^{18} Mx to 10^{19} Mx (unsigned flux) when the X-ray loops are observed.

Key words. Sun: photosphere – Sun: magnetic fields

1. Introduction

Describing the evolution of the magnetic flux on the surface of the Sun is a key component for understanding the complex couplings involved in energetic events that release a considerable amount of energy into the solar atmosphere and beyond. For the quiet photosphere, tracking algorithms already exist that describe the motion and some physical properties of the magnetic features. Each of them have different limitations. Because of that, the most complete description of the photospheric magnetic flux cannot be achieved by using only one of these algorithms. Instead, combining different algorithms with different limitations, which are therefore complementary, gives a more complete description of the behaviour of the photospheric magnetic features.

Among the many algorithms that attempt to analyse the magnetic field, a consistent set of four algorithms have been developed and compared with each other in DeForest et al. (2007), where they are referred to as CURV (Hagenaar et al. 1999), MCAT (Parnell 2002), YAFTA (Welsch & Longcope 2003; DeForest et al. 2007; Stangalini 2014), and SWAMIS (DeForest et al. 2007; Lamb et al. 2008). Other algorithms designed to track magnetic features exist, such as those described in Démoulin & Berger (2003) and Welsch et al. (2004), and they focus on deriving averaged velocity flows in magnetised regions.

The present paper describes a new algorithm using a completely different paradigm than the algorithms mentioned above. It is also designed to track and label small magnetic features in the quiet photosphere, as small as can be seen, and it measures the magnetic flux they carry. This algorithm does not pretend to replace these well-tested algorithms, but it aims at providing another alternative that simply broadens the spectrum of science applications that can benefit from feature tracking techniques in

magnetograms, more specifically at the smallest scales (a few Mm or less).

This technique uses the same paradigm as the one used in a known algorithm called balltracking (Potts et al. 2004; Attie et al. 2009), which is designed to track granules in continuum images. In this technique, numerical spheres or “balls” are released onto the image. Their position is known at any given time, and they settle in the local intensity minima that move with the granules so we can track the group motions of the latter. However, because of the very different nature of the data provided by magnetograms, we needed to write a variant of this algorithm, which we refer to as “magnetic balltracking”.

After a brief summary of the balltracking paradigm in Sect. 2, the magnetic balltracking is explained in four main phases in Sect. 3. In Sects. 4 and 5 we discuss a few technical aspects of the algorithm. Its practical application is detailed in Sect. 6 and applied in Sect. 7 to a case study of flux emergence associated with the brightening of X-ray loops.

2. The balltracking paradigm

The full derivation of the equations of motion that animate the balls within balltracking can be found in Potts et al. (2004). Here we schematise the balltracking paradigm in the context of tracking magnetic features. The general principle is sketched in Fig. 1, where the forces and the integration of the position of one ball is represented across two consecutive data surfaces at the times t_1 and t_2 . These data surfaces are, in fact, pre-processed magnetograms (more on this in Sect. 3.1).

The three forces governing the motion of each ball can simply be expressed with the vectorial equation (vectors noted in bold letters):

$$m\dot{\mathbf{v}} = \sum_i \mathbf{f}_i(d_i) - m\mathbf{g}\mathbf{e}_z - \alpha\mathbf{v} \quad (1)$$

[★] Movies associated to Figs. 6 and 18 are available in electronic form at <http://www.aanda.org>

Fig. 1. General concept of the balltracking algorithm for one ball. The blue squares represent the pixels drawn at a height equal to the rescaled value of the unsigned flux density so that maxima in flux density appear as local minima. Only a few pixels are represented.

where v is the velocity of the ball and $\dot{v}$ its time derivative. Here, $\sum_i f_i(d_i)$ is the sum over all the elementary buoyancy forces exerted by the data point at the i th pixel at the penetration depth d_i . The user-defined gravity field g sets the maximum acceleration that a ball can possibly reach. It is oriented downward, in a 3D cartesian frame of reference where e_z is a unit vector pointing upward and α is a damping coefficient that controls the stability of the tracking. For simplicity, the mass m of the balls can be set to unity. In Fig. 1, the ball is moving with a magnetic feature, where the intensity of the magnetic flux density is reversed and rescaled so it can be seen as a geometric height. In such a picture, the local maximum of unsigned flux density is, in fact, a local minimum. While floating on the data surface represented by the blue waves, the ball settles in the local minimum, while the latter moves with the group motion. In this framework, we are only interested in using the final position of the balls at the end of the integration of Eq. (1) in each magnetogram.

3. The magnetic balltracking algorithm

3.1. Phase 1: preprocessing of the magnetograms

The first step of magnetic balltracking is to make the magnetic features “trackable” by the balls. The main change from the original balltracking is to account for the signed values of the magnetograms, and for the more contrasted maps, spanning typically over more than 2 orders of magnitude in the quiet Sun (from a few G to hundreds of G). So we have to rescale the magnetograms in order to rescale the magnetic features vertically into

“holes” of reasonable depth, from either positively or negatively signed magnetic flux density, and allow the balls to settle inside them. This is the purpose of the following preprocessing (see also Figs. 2 and 3).

To start with, let us consider the absolute value of the flux density in a single magnetogram as a scalar field $|B_z(x, y)|$, like the one in Fig. 2 (top). It is first reverse-scaled (non-linearly) into

$$B_z^*(x, y) = \max(\sqrt{|B_z(x, y)|}) - \sqrt{|B_z(x, y)|}.$$

Next, B_z^* is offset by its mean value and normalised to its standard deviation $\sigma_{B_z^*}$ taken across the whole frame:

$$B_{zn}^*(x, y) = \frac{B_z^*(x, y) - \langle B_z^*(x, y) \rangle}{\sigma_{B_z^*}}$$

where B_{zn}^* is displayed in Fig. 2 (bottom), with an intensity spanning over a few units, which is of the order of the horizontal size of the magnetic features that we want to track. This rescaled signal can be seen as a geometrical height. In a 3D plot (Fig. 3), the magnetic features look like holes into which the balls can settle easily. This particular choice of rescaling (against for instance a linear rescaling) is further justified in Sect. 4.

3.2. Phase 2: initialisation

Once the magnetograms are rescaled, the balls are initially positioned at the pixels whose absolute intensity in the original (not

Fig. 2. Preprocessing of the magnetic balltracking. *Top*: initial magnetogram B_z , calibrated into Gauss units, scaled between -40 G and $+40$ G. *Bottom*: B_{zn}^* obtained after rescaling B_z .

rescaled) magnetogram is above a given threshold. This initialisation is illustrated in Fig. 4 (top) where the balls’ centres are plotted on the original magnetograms.

With the magnetograms from the Narrowband Filter Imager of the Solar Optical Telescope (NFI/SOT) onboard Hinode (Tsuneta et al. 2008), the threshold is usually set between 5 G and 20 G. Regardless, these values should be chosen so one does not track random noise. In addition, this saves some computational time by reducing the number of balls that is much smaller than the total number of pixels. In the magnetic balltracking, we do not make assumptions on the size of the magnetic features, and the minimum length between the balls’ centres, within each magnetic feature, is 1 px at any time. Nonetheless, in a practical case of instrumental data, the minimum ball spacing should account for the full-width-at-half-maximum (FWHM) of the point spread function (PSF) of the instrument. Should the pixel scale of the imager be smaller than the FWHM, the minimum spacing between the ball should be set to the nearest integer value. This will optimise the total number of balls, and thus the computing time.

Once the balls are positioned on the magnetic features, or more precisely, within the “magnetic holes”, the polarity of each true feature is retrieved from the signed intensity of the original magnetogram, at the pixels mapped to the coordinates of each ball’s centre. This polarity is stored, and is a constant associated with each ball. It is referred to as the initial “ball polarity”.

Next, a few integration steps, typically 10 to 20 depending on the size of the features, are performed between the first and second frame, so the balls have time to converge down into the local minima. For instance, in Fig. 1, such a local minima in the first magnetogram (Magnetogram 1) would be found at X_{final}^1 .

Because a segmentation algorithm will be used on the tracked magnetic features (more on that later), it is not necessary to have several balls within the same feature. If several balls have converged to the same local minima, only one ball is kept. After this stage, it is still possible to have one large magnetic

Fig. 3. *Top*: B_{zn}^* in 3D, using the same data as in the bottom of Fig. 2, with the intensity used as a geometrical height. The colourmap is scaled exactly as the intensity. *Bottom*: same as the *top* figure, with balls that have settled in the “magnetic holes” after a few integration steps. The position of these balls on the original 2D magnetograms are shown in Fig. 4.

feature being tracked by several balls, if for instance the feature has several local minima. This is illustrated in Fig. 4 (bottom). Note that this significantly reduces the number of balls between the first (top) and the next frame (bottom), and consequently the computational time.

3.3. Phase 3: main tracking phase

After the initialisation, the next frames are loaded, and the balls track the local minima (the “magnetic holes”) like they do within intergranular lanes with balltracking (Potts et al. 2004).

In the original magnetograms, the local minima of the rescaled data correspond, respectively, to the local maxima and minima of the signed intensity (positive and negative, respectively) of the magnetic flux density. At any time, the position of each ball is known, and each ball is labelled with a unique number, referred to as the “ball number”. The positions can be plotted on-the-fly, so that one can check by eye the quality of the tracking. An example of the main tracking phase is visible in the snapshots of Fig. 5, with three ball numbers identifying three distinct magnetic features that are being tracked.

When a magnetic feature is moving too rapidly, the balls do not have time to settle in the local minima. In this situation, at worst, the balls may be delayed by a few frames, and several integration steps between each frame are necessary to make sure that the balls do not get lost. This gives them more time to “catch up” on the fastest features. Typically, for magnetograms taken at cadences of up to 3 min, 10 to 20 intermediate integration steps of the equation of motion are used between each frame. This may or may not be the same number of intermediate steps used for the initialisation phase, these are indeed two independent “tuning” parameters that depend on the resolution of the data and the time sampling rate of the instrument. Typically, the

Fig. 4. Initialisation of the magnetic balltracking. *Top*: the values in the original magnetogram B_z are used to dispatch the balls on pixels above 20 G (red crosses). *Bottom*: new positions of the balls, after integrating the equations of motion, on the same magnetogram.

number of integration steps is proportional to the former, and inversely-proportional to the latter. Indeed, at higher resolution, the features are relatively wider, and the balls must travel over a relatively greater number of pixels. At increasing time sampling rates, the features evolve less rapidly between two consecutive magnetograms, and so the balls need less time to catch up on the motions of the local minima.

For large connected magnetic features, the shape of the magnetic features looks like an extended surface full of holes, where each hole can be filled with a ball. See for example the white magnetic patch in Fig. 4 (bottom) around the coordinates (60, 120) and the balls in the 3D view in Fig. 3 (bottom). This method is suitable for tracking clustered features that are made of several fragments, such as the ones in Fig. 5.

At each tracking step, the polarity under the current position of the ball's centre is compared to the ball polarity. The tracking of a given ball ends as soon as the polarity of the current pixel is reversed with respect to the ball polarity.

This strategy has several advantages. Indeed, to “see” a reversed polarity, a ball needs either to keep tracking down to the noise level until the current pixel polarity reverses (which

Fig. 5. Magnetic balltracking of a small field of quiet Sun, at different time steps. The numbers bound to the crosses (red) are the labels that are unique to each ball (i.e. the ball number).

is known by looking it up automatically in the original magnetograms), or it needs to encounter a magnetic feature of opposite polarity. This defines two conditions for the end of the tracking of a given ball. They are described below separately.

Condition 1: tracking down to the noise level

Tracking down to the noise level makes the algorithm use the true sensitivity of the instrument. Indeed, even if the initialisation step uses a threshold, the features are tracked until they cannot be detected by the instrument, i.e. to values below the threshold. Should we ever need to track the faintest magnetic features from the beginning, the threshold may simply be lowered down to the noise level, which has the only consequence of increasing the computing time (more balls will be added). When the ball is floating over random noise, the sign of the intensity in the original magnetograms eventually reverses. In this case the tracking of this ball ends. The stability of the ball within noisy data is set by the damping coefficient α (Eq. (1)). Typical values are between 0.1 and 1, depending on the time sampling rate and spatial resolution. In term of damping time (defined as $T_d = 1/\alpha$), this is equivalent to values between 1 and 10, in units of time interval

between frames. The damping force is necessary not only in the presence of noise, in which case the tracking is, in fact, more resilient with damping than without, but it is also necessary for the general stability of the code, regardless of the noise. This is illustrated in more detail in Sect. 3.5.

Condition 2: no crossing of opposite polarity

Another issue we had to solve is how to prevent the ball from crossing a feature with positive flux to a neighbouring one with negative flux. This would occur for example if the ball is in a local minimum, and if it has kept enough momentum to reach another close-by local minimum of opposite polarity at the next integration step. This can also occur, if the close-by local minimum moves quickly towards the ball, like for example in the case of the two footpoints of a loop being submerged.

This problem can be thought of as a “numerical tunnel effect”, in the sense that at one time a ball is in a magnetic hole, and at the next, it has crossed a barrier and lies within the second magnetic hole. The solution is as follows: because the true polarity of each pixel is known from the original magnetogram, the local minima are always associated with a polarity, which is compared to the ball polarity (see Sect. 3.2) at the end of each integration step. Should a ball lie in a magnetic hole with a polarity opposite to the ball polarity, the tracking of this ball ends.

Conditions 1 and 2 are checked independently for each ball. If either of these two conditions is fulfilled for a given ball, its tracking ends, so that the lifetime of a ball corresponds to the lifetime of the magnetic feature tracked so far. This does not end the whole algorithm, which continues as long as other balls remain. Note also that in Phase 3, whenever there are overlapping balls, i.e., their centres are positioned on the same pixel, only the oldest ball is kept, which allows the right estimation of the features lifetime (more details on this in Phase 4).

Note also that the magnetic features may exist at very irregular places (see Fig. 4, top), and consequently, the balls end up scattered over a rather irregular grid. Therefore, contrary to the original balltracking algorithm, the conversion of the velocity, initially calculated in a Lagrangian frame of reference, into a Eulerian frame of reference cannot happen. In other words, this method tracks and follows the individual motions of the magnetic features. It is not designed to output a “flow field”, with the values of the velocity known at any given time at fixed positions, on a regular grid, like in, e.g., Démoulin & Berger (2003) and Welsch et al. (2004). This is simply because there are parts of the magnetograms where the magnetic flux is not detected, and thus, there is no velocity defined there. At best, Eulerian flow fields can only be derived locally, in regions where there are enough magnetic features that provide a more “reasonable” sampling.

3.4. Phase 4: detection of emerging flux

Note that this phase is an optional module in the algorithm. When used, it runs simultaneously with Phase 3. During this phase, the algorithm permanently scans for new pixels that would rise above a given “detection threshold”, which must be at least greater than the noise level of the instrument. It may, or may not be equal to the threshold defined in Phase 2. When new pixels “rise” above this threshold, new balls are added on these areas so that emerging flux can be tracked. This phase needs another tunable parameter, referred to as the “detection spacing”, which sets the minimum distance between the new balls and the

ones already present. Any pixel whose intensity rises above the threshold must satisfy the condition of being at a distance larger than this parameter, in which case a new ball is put there. Note that this makes it also possible to follow large features whose size varies significantly over time. Indeed, the large features are the ones with many pixels above the detection threshold. If the distance between these pixels and those that have balls already, is greater than the detection spacing, new balls will eventually be put there too. Consequently, if a given feature, initially small and populated with one ball, grows over time (for instance, as a result of flux emergence), it may be populated by more than one ball. Note that the detection spacing also defines the resolution with which emerging flux is detected. The subsequent tracking of the new balls is exactly the same as the other balls.

An example is given with the ball 5274 in Fig. 5 (starting from frame 5, near the upper right corner, seen as a small red cross). It is tracking emerging flux that, before frame 5, was below the initialisation threshold of 5 G. When the magnetic feature emerges above 5 G, this new ball locks onto it and tracks it until the last frame. This is allowed because the nearest balls tracking other features are at a distance greater than the detection spacing. So this phase makes the algorithm useful not only to track the flux visible from the start, but also to track the emerging flux, and little pieces of the largest features that may (or may not) fragment over time.

In the first movie attached to the online version of this paper, one can compare the effects of different values of the detection spacing (at a given detection threshold of 5 G), and visualise Phases 3 and 4 of magnetic balltracking. A snapshot of this movie is shown in Fig. 6. Three panels are shown. With these data, the pixel size is 0.2 arcsec and the resolution is about 0.3 arcsec px⁻¹ (Chae et al. 2007). In the left-hand panel, the detection of the emerging is deactivated. In the middle panel, the detection spacing is 15 px, which is still a bit “loose”, and one misses a few emerging features. Finally, in the right-hand panel, the detection spacing is set to 5 px which is a rather “aggressive” detection. The latter may be harder to follow because the finer the detection grid, the more balls are involved, and the more “crowded” these movies get. Nonetheless, it allows for more flux to be detected. A detection spacing of 10 px is typical. With a pixel size of 0.2 arcsec, which is about 6 to 7 seven times the size of the resolution element (0.3 arcsec px⁻¹). One can also see several balls gathering from the boundaries toward the centre of the largest features. This is due to new balls, added near the edges of the large features (Phase 4), converging toward the same local minima. Like in Phase 3, only the oldest ball is kept whenever, and wherever they overlap. This prevents “young” balls, i.e., the ones added during the flux detection, from replacing the older ones. Otherwise, we could not extract meaningful lifetimes during Phase 3.

Note that it is also possible to track emerging flux by simply reversing the timeline of the data series and recording submerging flux. In such a case, as far as the algorithm is concerned, flux removal is indistinguishable from flux emergence. Depending on the situation, this may be more suitable than using the module described in Phase 4.

3.5. Notes on the damping force

Without damping, the balls have too much inertia and “free-fall” toward the local minima. With too great a speed, they can move passed the local minimum. The damping coefficient prevents this “overshoot”, and damps the subsequent oscillations of the positions when the balls “jiggle” around the local minimum. This is

Fig. 6. Snapshot comparing the effects of the detection spacing, at a given detection threshold of 5 G. The red crosses are plotted at the positions of the ball's centre. *Left*: no flux detection. *Middle*: loose detection spacing at 15 px. *Right*: aggressive detection spacing at 5 px. A [movie](#) showing the temporal evolution is available in the online edition.

Fig. 7. Tracking of the local minimum of a synthetic Gaussian surface. The balls are projected on the x -axis and z -axis at three different numbers of integration steps. Step 1 (yellow), 15 (red), 100 (green). The *left column* uses the same damping time for the vertical and horizontal axis. The *right column* shows the tracking with a horizontal damping time (T_{dh}) different from the vertical damping time (T_{dz}).

illustrated in Fig. 7 using a synthetic, rescaled Gaussian surface, used as a magnetic feature that has a maximum strength of 50 G, and a local minimum at $x = 11$ px. The FWHM of the original gaussian curve (i.e. before rescaling) is ~ 7 px. On the left column, three values of damping times are used ($T_d = 1, 5$, and 10).

On the right column, we separated the damping times into a horizontal (T_{dh}) and a vertical (T_{dz}) damping time that are associated respectively with the horizontal and vertical components of the damping force, whose usage are recommended in Potts et al. (2004) (although in the case of tracking granules). The balls are plotted at the integration step 1, 15, and 100. Note the overshoot seen with the red ball at $T_d = 5$ and $T_d = 10$, as well as at $[T_{dh} = 5; T_{dz} = 5]$ and $[T_{dh} = 5; T_{dz} = 10]$. The red ball has moved passed the local minimum and falls back to it at a later time (green). The overshoot is not seen with a shorter damping time in both directions $T_d = 1$ and a shorter vertical damping time $[T_{dh} = 5; T_{dz} = 0.5]$.

The x -coordinates of the ball in each case are plotted in Fig. 8. The top and bottom panel (resp.) correspond to the displacements the balls in the left and right column (resp.) in Fig. 7. The overshoot corresponds to $x > 11$. Note how the oscillations are reduced more efficiently with a decreasing damping time. At $T_d = 1$, there is indeed no overshoot, and the ball settles in the local minimum after 30 steps. Yet it may be more efficient to use $[T_{dh} = 5; T_{dz} = 0.5]$, which induces a small overshoot, while within 25 steps the ball is less than 1 px from the local minimum. $[T_{dh} = 5; T_{dz} = 10]$ is a limit case where the damping force on the vertical axis is too small such that the tracking does not converge (the ball is not visible in Fig. 7 due to its z -coordinate that is off the grid). If efficiency is not an issue, the use of a unique damping time $T_d = 1$ is sufficient, as long as the convergence criterion is fulfilled (see later in Sect. 4). However, for an optimum stability on the vertical axis, and based on many experiences (not shown here), we recommend the use of T_{dz} within $[0.5; 1]$, and T_{dh} within $[1; 10]$, which may vary according to the data characteristics: resolution, pixel size, cadence, noise level, and sizes of the features.

4. Motivations for the rescaling method and convergence criterium

The choice of the rescaling method of the magnetograms is driven by two objectives: (i) a successful tracking of the features;

Fig. 8. Coordinate on the x -axis of the ball’s centre during the tracking of Fig. 7. The *top and bottom panels* correspond to the motions of the ball in the left and right column of Fig. 7.

and (ii) a minimisation of the number of integration steps that are necessary to achieve (i). In this section we demonstrate the advantages of the choice of rescaling defined in Sect. 3.1.

4.1. Formal definition of the rescaling methods

Our “generic” rescaling method formalises as follows:

$$B_z^*(x, y) = \max(|B_z(x, y)|^\gamma) - |B_z(x, y)|^\gamma \quad (2)$$

$$B_z^m(x, y) = \frac{B_z^*(x, y) - \langle B_z^*(x, y) \rangle}{\beta} \quad (3)$$

in which we have introduced a normalisation factor β , and γ : linear rescaling is equivalent to $\gamma = 1$ and non-linear rescaling equivalent to $\gamma \neq 0$ and 1. Then we compare the following three scaling methods:

1. $[\gamma = 1; \beta = 1]$: no rescaling is actually performed.
2. $[\gamma = 1; \beta = \sigma_{B_z^*}]$: linear rescaling.
3. $[\gamma = 0.5; \beta = \sigma_{B_z^*}]$: non-linear rescaling.

Method 3 is the method defined in Sect. 3.1. $\sigma_{B_z^*}$ was originally defined as the standard deviation of $B_z^*(x, y)$, it implicitly depends on γ (Eq. (2)). Note that subtracting by the mean value in Eq. (3) only changes the zero-point of the data surface and does not affect the results. It is used here to conveniently define the mean value of the data as the origin of the vertical axis.

4.2. Convergence criterium

The positions of the balls at the end of 100 integration steps are shown in Fig. 9 for the three scaling methods 1, 2, 3, respectively,

Fig. 9. Final positions of the balls’ centre, after 100 integration steps, for the three scaling methods. The original positions are the ones shown in Fig. 4 (top, red crosses).

in blue dots, green and red crosses, respectively. The yellow contours are set at the initialisation threshold of 10 G. The initial positions are the ones plotted in Fig. 4 (top). In the current example, a successful tracking has the necessary (but not sufficient) condition that all the final positions are within the yellow contours in Fig. 9. Yet the blue dots end up outside the contours (i.e., outside the magnetic feature they are meant to track), whereas the others are all within the contours, which proves that the scaling method 1 does not result in a successful tracking, and that Methods 2 and 3 are more appropriate.

In Fig. 10 we have plotted three examples of the traveled distance (in pixels) of three different balls’ centres, using the linear (green dotted line) and the non-linear (red continuous line) scaling. The chosen balls, respectively from top to bottom, correspond to the ones pointed to by the cyan arrows in Fig. 9, respectively from left to right. The oscillations are due to the balls overshooting to either side of the local minima before they can finally settle. The oscillations are damped within one to two times the horizontal damping time (here it is set to 4).

The time that the balls take to reach their final position defines a convergence factor represented in Fig. 11. The “final position” is set as the position reached by the ball after a sufficiently high number, here set to 100. The convergence factor is then defined as the ratio of the number of balls that have reached their final position at a given integration step, to the total number of balls. When the convergence factor reaches 1, the convergence criterium is satisfied, and this sets the number of initialisation steps. The scaling method 1 converges after about 95 integration

Fig. 10. Distance from the initial position, for the three balls (resp. *top to bottom*) pointed to by the cyan arrows in Fig. 9 (resp. from left to right).

Fig. 11. Convergence factor against the number of integration steps, for the three scaling methods.

steps, the other two methods converge in less than 40 integration steps, with the non-linear rescaling (red continuous line) converging more rapidly by a few integration steps. With the scaling methods 2 and 3, 90% of the balls have converged within 10 to 20 steps, and all the balls have converged within 40 steps (the convergence factor equals 1). 20 initial steps is typical but a “trial-run” like the one here must be done in each case study to determine the optimum value for a given data set. Here, the optimum value would be 40.

4.3. Comparisons between the rescaling methods 2 and 3

At some places, like in the yellow rectangle (top right) of Fig. 9, the green crosses are not overlaid by the red ones, and instead, a

Fig. 12. Close-up in the region of the yellow rectangle in Fig. 9. *Left:* linear rescaling (Method 2). *Right:* non-linear rescaling (Method 3). The small cyan dots are the initial positions. The coloured circles are the final positions. The yellow lines are mapping the final position to the corresponding initial positions. They are not the actual trajectories. *Bottom:* 3D data surface viewed from the left Y -axis.

few of them ended up at different positions within a few pixels. This first suggests that whereas the linear rescaling (Method 2) has converged according to Fig. 11, the balls, in fact, have not converged in a local minimum, contrary to Method 3 where the red crosses suggest that the balls have unambiguously settled in the same local minimum. Figure 12 is a close-up in the yellow rectangle of Fig. 9 viewed from the top (top panels) and from the side (along the left Y -axis) of the 3D rescaled data surface (bottom panels). The coloured circles are plotted at the final positions of the balls’ centre. One sees that there is indeed only one local minimum in this area, but the balls in Method 2 (left) did not settle in it. Some of them are standing still in the middle of the steep slope (bottom left), which is not what one might expect in a real “natural” situation. This is a typical example of a linear rescaling being less appropriate than a non-linear rescaling with $\gamma = 0.5$ (right). With a linear rescaling, the strongest magnetic features (here above 100 G) still have a very steep slope (bottom left), while the slope is more gradual with a non-linear rescaling (bottom right).

The reason for the very steep slope resulting in an inconsistent tracking is related to the number of data points that are actually in contact with the balls, and therefore, contributing to the total force f_i (see also Eq. (1) and the blue squares in Fig. 1). Indeed, for maximum efficiency, the 3D grid defined by the balls, and which samples the data surface, uses regularly spaced values with a grid size of 1 px. For a data point to be considered by the algorithm, it needs to be located within one ball radius from the ball’s centre. The data at these ball grid points are interpolated. Note that nearest neighbour or linear interpolation make no difference on the final positions (the former is preferred for efficiency). At a given ball radius, the more of these “contact points”, the better, as more contact points means that

Fig. 13. Average number of contact points per ball, against the number of integration steps. *Top:* average over the FOV of Fig. 9 for the three scaling methods. *Bottom:* average over the balls in Fig. 12 for the linear and non-linear scaling methods.

more of the magnetic feature actual topology is accounted for, and thus the more coherent the tracking is. The number of contact points involved in the tracking in the three different scaling cases is shown in Fig. 13. In the top panel, we use the number of contact points averaged over all the balls involved in Fig. 9. The top panel again shows that in the case of no rescaling (blue dashed line), there is on average merely 1 contact point, which explains why the scaling method 1 fails. Methods 2 and 3 differ only slightly on average, by only 1 contact point. Locally however (bottom panel), in the cases corresponding to the close-up in Fig. 12, this difference is more significant. Before the convergence criteria is reached (i.e. before the integration step 40), the number of contact points varies from 1 to less than 6 with a linear rescaling, whereas it varies from 1 to more than 8 with a non-linear rescaling. Furthermore, with the latter method, this number increases more rapidly: it equals 8 while it is only 2 with the linear rescaling, which can only result in an insufficient force to push the balls down the slope. Put simply, with too few contact points, the balls can stall.

Thus, Method 3 has two advantages over Method 2: a better handling of the steeper slope for the magnetic features with field strength of the order of 10^2 G, and a relatively higher efficiency (a few integration steps less are needed). The tracking of faint magnetic features is not affected by these effects. Visual inspections show no difference in the final positions between Methods 2 and 3 when tracking the smallest and weakest features close to the initialisation threshold. In addition, no case where Method 3 is less accurate than Method 2 were found, and therefore, the non-linear rescaling is the most appropriate choice among the three tested methods. Note that other values of γ within $[0; 1]$, but relatively close to 0.5, are possible, and they

Fig. 14. Final positions using three different ball radii: $R_s = 2$ px (small red crosses), $R_s = 3$ px (medium-sized green crosses), and $R_s = 4$ px (large blue crosses). “Nb” is the number of balls visible in the figure for each ball radius. The yellow arrows point at local minima in which only the balls with a 2-px-radius converged.

may be tuned to different values for optimisation purposes with quiet Sun features, with negligible changes on the final positions.

5. Effects of the balls radius

The appropriate choice of the ball radius relates to the size of the magnetic features that are to be tracked. The sphere radius must be greater than 1 px, regardless of how small the features are. At 1 px the number of contact points is too small to result in a coherent tracking, which stays dominated by high-frequency noise. Figure 14 shows the final positions of a tracking using three different radii (R_s), represented by coloured crosses of increasing sizes: $R_s = 2$ px, $R_s = 3$ px, and $R_s = 4$ px. “Nb” is the number of different local minima in which the balls have converged. It is here equal to the number of crosses of a given colour that are visible in the figure. Nb is respectively equal to 124, 117, and 103. So the number of different local minima in which the balls managed to settle decreases when the radius increases. The “missing” balls at $R_s = 3$ px, and $R_s = 4$ px have either “fallen off” the edges and/or passed “through” the smaller ones located near bigger local minima in which more of the bigger balls are pulled in more efficiently: the yellow arrows show examples of local minima resolved by the use of a 2 px-radius, but that are not resolved by the use of $R_s = 4$ px and/or $R_s = 3$ px. There is one place where a ball with only the 4-px-radius converges (near the coordinates

Fig. 15. Travelled distances in three cases, with three different ball radii: $R_s = 2$ px, $R_s = 3$ px, and $R_s = 4$ px. From top to bottom (resp.), the balls ending up at the location shown by the cyan arrows from left to right (resp.) in.

(100, 100)) and it is, in fact, a “false positive” (i.e., not an actual local minimum): the lack of resolution made the bigger grid (due to the greater radius) cover pixels on the neighbouring feature (black patch on the left), which averages the motion of the tracked white patch and of this nearby black patch. On the other hand, we do not find any “false positive” with the red crosses ($R_s = 2$ px), which makes it, here, the best choice.

Note that these test data are from NFI/SOT (Hinode). The data were binned onboard to a pixel size of ~ 0.2 arcsec while the spatial resolution of NFI magnetograms is $0.3 \text{ arcsec px}^{-1}$ (Chae et al. 2007), so a radius of 2 px or 3 px would be appropriate. Nonetheless, for a given initial ball spacing (defined as the minimum space between the balls at the initialisation phase), using a greater radius reduces the resolution of the tracking, and also increases the computing time and memory usage as a bigger ball encompasses more grid points.

Figure 15 is a comparison between the horizontal distances travelled by the balls at three locations chosen at random (from top to bottom, respectively) pointed to by the cyan arrows in Fig. 14 (from left to right, respectively). In each case, the three radii are tested. The distance has its origin at the initial position of the balls, and in each case, the three balls have the same initial position. Figure 15 shows there is less than one pixel of discrepancy in the final position between $R_s = 2$ px and $R_s = 3$ px. This discrepancy is greater with $R_s = 4$ px, about 2 px according to the bottom panel. Looking at other cases (not shown here) gives the same conclusion that $R_s = 2$ px and $R_s = 3$ px have the least discrepancy of less than one pixel, and that $R_s = 4$ px is less accurate. Note that because the size of the features, in pixels, also depends on the plate scale and on the resolution of the instrument, we cannot give ideal values that work in all situations. Such tests on subsets of large data with different radii are necessary to optimise this choice. In addition, one could combine the results of

different radii for multi-scale analyses. The small radii would be dedicated to the smaller features, while the greater radii would be tracking the broader patches.

The main output of magnetic balltracking is the series of final positions of the balls, for the purpose of tracking the time-dependent displacements of the local extrema of the flux. These positions are then used in a segmentation algorithm to integrate the flux over the area of the tracked features. This is detailed in the next section.

6. Segmentation of the magnetic features

As mentioned earlier, the magnetic features can be very clustered in the quiet Sun. Then quantifying the evolution of each bit of each magnetic structure, which can be near the limit of the instrumental resolution, is quite challenging. Magnetic balltracking facilitates this procedure. As explained in Sect. 3, this technique tracks the time-dependent positions of the local extrema of individual magnetic features. Thus the next step in describing their evolution is to integrate the magnetic flux of these features. An easy way to do this is by applying a “region extraction” algorithm. The technique is also known as “region growing”, and has many names and variations that depend on the scientific field in which this segmentation technique is applied. It is one of the basic algorithms detailed in textbooks of digital image processing (see for example Gonzalez & Woods 2008, Chap. 10, Sect. 10.4).

The “region growing” used here, consists in extracting the magnetic features that have been tracked, from the rest of the magnetograms, so we can easily integrate the intensity of the flux over the extracted area. We use the tracked positions (the final positions of the balls in each frame) as so-called “seed points”. For each position (or each “seed point”), the difference between the intensity of the neighboring pixels and the seed pixel is compared to a given threshold. If the comparison is true (in the logical sense), it is added to a list of pixels connected with each other, including the seed, which “grows” a region and thus segments the magnetic feature from the rest of the magnetogram. If the comparison is false, the pixel is not added to the list. The region stops growing when there are no more connected pixels. The output of the region-growing algorithm is a binary mask: an array of logical values, co-spatial with the magnetograms, where the pixels in the grown region are set to 1, and the other pixels are set to 0. These masks can directly be used to extract the features from the magnetogram in order to integrate the flux of the tracked features. Such extracted features are visible in Fig. 16, which used the tracked positions (i.e., the seeds) previously illustrated in Fig. 5 as the input of this segmentation.

As mentioned in the previous section, several balls may track the same wide magnetic features, with as many balls as there are local extrema in it. As the position of these balls are used as seed points, they will ultimately extract the same connected pixels, and output identical masks. Then we have to get rid of the duplicates, which we do by using a logical “or” (equivalent to a logical “union”). If the same masks of connected pixels are output for different balls, the logical “or” reduces them to one unique mask before integrating the flux. This makes sure that, when looping over the extracted regions to integrate their flux, we do not integrate it over the same region more than once.

Finally, each magnetogram is integrated over the extracted areas, which gives the flux carried by the tracked feature in each single frame. Repeating this for all the magnetograms provides the time-dependent flux of all the tracked features, including the emerging one if any has been detected during Phase 4. The

Fig. 16. Region of the magnetograms extracted using the region-growing algorithm onto the balltracked seed-points of Fig. 5.

ball numbers are used as the labels of the magnetic features, which makes it easy to select them individually. This way we can choose which one to extract with the region-growing algorithm, but it is also possible to simply take all of the tracked features for more global statistical analyses.

One limitation of this method is that it cannot grow a region using too low thresholds, otherwise, close but probably non-connected regions may be added to the list of “good” pixels. This leads to a wrong segmentation, and ultimately a biased estimation of the flux. This could probably be solved using a more sophisticated segmentation algorithm. For consistency, we use thresholds close, but not necessarily equal to the thresholds used in the previous phases of the magnetic balltracking (defined in Phase 2). They are between, typically, 5 G and 15 G. We found these values by “trial and error” and they turned out to be optimal in the very clustered flux of the quiet Sun. A direct unfortunate consequence is that no flux is integrated below these thresholds. Nonetheless, the actual value of the thresholds that we have used so far depend more on the instrumental detection limits rather than being an intrinsic limitation of the algorithm.

Difference with other algorithms

To identify the magnetic features, the four codes compared in DeForest et al. (2007) use a similarity metric in the (x, y, t) space (2 spatial dimensions and 1 time dimension) based on logical comparisons between the shape and position of the segmented areas. This is done after the segmentation in the region of interest. The features are then identified and labeled as magnetic entities at the end of this process. These are then the “tracked” magnetic features. The motion of the features may or may not be derived using centre of gravity of the extracted area.

In magnetic balltracking, the magnetic features are already labelled at the first phase, and in Phase 4 for the emerging ones, by the ball numbers that settle within their local maxima. Furthermore, the tracking is done before the discrimination of the magnetic fragments, using the balltracking paradigm. The latter occurs as a post-processing of the magnetic balltracking. Here we have used region growing as the post-processing method to build the masks that extract the magnetic fragment from the images, which is similar to the so-called “clumping” used in MCAT. Regardless, the magnetic balltracking can be combined with other methods of segmentation, as long as they can use the balls as seed points. Within magnetic features large enough to have more than one local extrema, and thus more than one balltracking the whole feature, the duplicated masks associated with each of the ball (used here as seed points) are unified into one single mask. Further discrimination within these large features cannot be achieved with this method, and depending on the science goals, one may revert to using a more discriminative method like the so-called “downhill” method in YAFTA and SWAMIS.

In the next section, we demonstrate the capabilities of magnetic balltracking in a case study of flux emergence using Michelson Doppler Image (MDI) data and co-spatial X-ray images from the X-Ray Telescope (XRT) on Hinode.

7. Magnetic balltracking on flux emergence

7.1. Observation of flux emergence

In what follows, all times are given in Universal Time (UT). The observations were made on September 26th, 2008, and consist of high-resolution, 1 min-cadence continuum images and magnetograms from MDI (Scherrer et al. 1995), and co-spatial images from XRT (Golub et al. 2007) in soft X-ray (pixel size of 1 arcsec), at ~ 30 s-cadence. The different time series last 4.5 h, between 15:00 until $\sim 19:30$. The MDI magnetograms were rigidly de-rotated using the local latitude at the centre of the field-of-view (FOV), which is consistent with the preprocessing recommendation in DeForest et al. (2007, Sect. 5.1). The time series of XRT images were calibrated and registered using the routines related to the XRT instrument in Solarsoft (xrt_prep.pro, xrt_jitter.pro). The co-alignment of the XRT images onto the MDI magnetograms was done using co-temporal XRT images and Extreme-ultraviolet Imaging Telescope (EIT) images (Delaboudinière et al. 1995) whose frame-of-reference is put into alignment with the MDI frame-of-reference using the header information. Accuracy of the latter was checked using limb fitting of the full solar disk. The XRT and EIT frame used for the co-alignment were taken near 19:00. We estimate the XRT and MDI frames to be co-spatial within less than 1 arcsec (< 1 Mm).

The flux emergence is associated with the rise of X-ray loops observed at the same location, and shown in Fig. 17 (pointed

Fig. 17. Flux emergence observed at two different places within the displayed field of view. Red/green contours are positive/negative flux (respectively). Thin unfilled contours are at 10 G, filled contours are at 50 G. The orange arrows point at emerging flux regions that are followed by the rise of X-ray loops. The blue contours are the contours of the supergranular boundaries.

to by the orange arrows). The snapshots are taken in a FOV of $60 \times 60 \text{ Mm}^2$. Balltracking was used to derive the flow fields, and the associated supergranular network lanes are drawn as blue contours. These are obtained with the algorithm of automatic recognition of supergranular cells from Potts & Diver (2008).

At 15:12, the flux is barely visible in the internetwork, until it emerges as a very fragmented, mixed-polarity flux after 16:00 (left arrows). This occurred quite close to the supergranular boundaries (blue lanes in Fig. 17, only 10 Mm away from it, which is consistent with previous observations of flux emergence (Wang 1988; Stangalini 2014). The clustered flux then drifts away, and is finally observed with one clearly visible X-ray loop, at 17:55. At 18:12 another X-ray loop seems to connect the negative-polarity footpoint from the left side of the network lane in the middle of the frame, to the positive one on the other side. A second, weaker emergence is seen near the bottom right part of the snapshots, starting at 16:42, and also gave rise to X-ray loops, visible at 19:05. The time scale of these emergences is a few hours. The amount of flux in the emerged bipoles that are observed at the footpoints of the X-ray loops is measured with the magnetic balltracking. The results are presented in the next section.

7.2. Results of the segmentation

Magnetic balltracking is performed on the FOV of the snapshots in Fig. 17 to track the features pointed to by the orange arrows. The results are used by the region-growing algorithm: as explained in Sect. 6, each ball can act as an identifier of a whole

Fig. 18. Snapshot of the region extraction of the emerging flux. The right column are the magnetograms. The left column are the masks of the extracted regions. The contours are at ± 20 G. A movie showing the temporal evolution is available in the online edition.

magnetic feature, using the ball numbers as unique labels. This makes it easier to detect and isolate only the emerging flux. To do so, we simply associate the number of the balls that were tracking these features to the time series of flux that is integrated by region-growing. Here, we selected only the flux that had increased at the end of the time series by twice the standard deviation. This is sufficient to isolate the emerging flux seen in Fig. 17. At this stage, all the emerging flux is not associated with bright X-ray emission. Finally, we use the ball numbers to select precisely the patches that are emerging underneath the X-ray loops. This last selection is not an automated process as we need to know where the X-ray loops are located, and to identify the ball numbers by eye.

One can see the result of this selection more precisely in the three snapshots shown in Fig. 18, and in the second movie (online supplemental). The left panel shows the magnetograms

with the positions of the balls (red crosses) that are tracking the emerging flux (see also Fig. 18). The small turquoise crosses show the positions of the balls used throughout the magnetic balltracking, excluding the ones added during Phase 4. The bigger red crosses correspond to the balls selected for the region extraction of the emerging flux. The blue and green contours outline the boundaries of the extracted masks that are used to spatially integrate the flux density in the magnetograms. These masks are shown as binary patches in the right panel (black for negative flux, white for positive flux). They are the main output of the region extraction algorithm (Sect. 6). The detection threshold was set to 20 G, which is near the noise level of the magnetograms in MDI. The pixels with an intensity below this value are ignored. It is possible to see the limitations of such a low threshold in the left panel of the movie, by considering the green contours on the left, near Frames 173–174 and again around Frame 213. The extracted region suddenly includes small nearby features for 1 to 2 frames. This illustrates possible errors when using region growing with a detection threshold too close to the noise level. Nonetheless, the short time scale (1 to 2 frames) of such extraction errors, compared to the lifetime of the correctly extracted features (here, an order of magnitude greater) makes it possible to filter them out, for instance with median filtering over a few time steps, or using Fourier filtering in the time-frequency domain, without significantly impairing the rest of the data. However, these techniques are not part of the magnetic balltracking and we will not discuss them further.

The integrated fluxes are plotted in Fig. 19. The flux in the first emergence (Fig. 17, left-side arrows, and Fig. 19, top) is balanced at the beginning (15:00), for ~40 min with a positive and a negative flux of a bit less than 2×10^{18} Mx. It is unbalanced for ~2.5 h, until it is balanced again at ~18:15 with an unsigned flux of $\sim 10^{19}$ Mx. The X-ray emission increases ~1 h after the emergence is first detected, by ~70% of the background intensity from $\sim 1200 \text{ DN s}^{-1}$ at 15:00 up to $\sim 3200 \text{ DN s}^{-1}$ after 18:30¹.

In the second case of emergence (Fig. 17, right-side arrows, and Fig. 19, bottom), the magnetic flux is about 50% weaker than in the first region, with a maximum positive and negative flux (respectively) between 3.5×10^{18} Mx and 4.5×10^{18} Mx. The flux is unbalanced during about 2 h, between ~16:40 and ~18:40, although the flux balance is not obvious afterwards. Like in the previous case, this flux emergence is followed by the rise of an X-ray loop. The X-ray emission increases by about 50% from a background level of 1000 DN s^{-1} at 15:00, up to $\sim 1500 \text{ DN s}^{-1}$ after 19:00 when the X-ray loop is visible (Fig. 17, bottom right panel). Note that there is an X-ray data gap between 18:36 and 19:02. The X-Ray gap is filled with the first value available after the gap (19:02), which is only an arbitrary cosmetic correction.

Comment on the results

The focus is given here on the current capability of the algorithm, regardless of the capability of the instrument. Therefore, for a better assessment of future improvements of the algorithm, this first example of a science application of magnetic balltracking is using data as close as possible to what is delivered by the instrument. Therefore, the magnetograms are not free of P-mode oscillations that may contaminate the Zeeman-related signals (DeForest et al. 2007). Here these oscillations directly affect the region-extraction algorithm by changing the apparent flux density and area of the extracted feature. The less the signal-to-noise ratio in the extracted area, the more visible the

Fig. 19. *Top:* evolution of the X-ray intensity (black line) and the magnetic flux in the region of emerging flux on the left part of the snapshots in Fig. 17. The colours of the curves of the flux are consistent with the contours. Red is positive flux, green is negative flux (in absolute value). *Bottom:* same as the top panel, for the second region of flux emergence in the bottom right of the snapshots in Fig. 17. There is an X-ray data gap between 18:40 and 19:02 that has been filled in by the first available value after the gap.

oscillations are. This is responsible for the relatively greater oscillations in the second case of the flux emergence in Fig. 19 (bottom), which we observed to be 50% weaker than in the first case (top). This patently shows that for actual statistical analyses with magnetic balltracking, the preprocessing guidelines in DeForest et al. (2007), which are not, per se, part of the magnetic balltracking, should be followed.

8. Summary and prospects

In this paper, we have presented our implementation of an efficient method called “magnetic balltracking” that tracks the magnetic features down to their finest scales. This new algorithm allows us to quantify the evolution of the magnetic flux. Applied on MDI data, it allowed us to detect, track, and quantify the evolution of emerging flux between 10^{18} Mx to 10^{19} Mx on a very fine scale of a few Mm, and that are followed by the rise of soft X-ray loops within a few hours.

Although we have presented here an application to estimate the flux over the tracked features, magnetic balltracking has a much broader range of possible applications. It could be applied to the tracking of the footpoints of magnetic field lines when performing extrapolations, the study of MHD waves in flux tubes, or the diffusion of internetwork elements whose emergence,

¹ DN s^{-1} : Data Number per second.

transport and disappearance can now be tracked in detail. In addition, the technique may be applied to the formation of Active Regions. Because magnetic balltracking can detect emerging flux and track the elements as they grow and move apart, it could complement other algorithms used in surveys of the solar cycle that use different algorithms to detect and track sunspots, e.g., [Watson et al. \(2009, 2011\)](#) and [Goel & Mathew \(2014\)](#).

Acknowledgements. This work was funded by the International Max Planck Research School in Göttingen (Germany), and by grant STFC/F002941/1 from the UK's Science and Technology Facilities Council, held at the School of Physics and Astronomy, University of Glasgow.

References

- Attie, R., Innes, D. E., & Potts, H. E. 2009, A&A, 493, L13
 Chae, J., Moon, Y., Park, Y., et al. 2007, PASJ, 59, 619
 DeForest, C. E., Hagenaar, H. J., Lamb, D. A., Parnell, C. E., & Welsch, B. T. 2007, ApJ, 666, 576
 Delaboudinière, J., Artzner, G. E., Brunaud, J., et al. 1995, Sol. Phys., 162, 291
 Démoulin, P., & Berger, M. A. 2003, Sol. Phys., 215, 203
 Goel, S., & Mathew, S. K. 2014, Sol. Phys., 289, 1413
 Golub, L., Deluca, E., Austin, G., et al. 2007, Sol. Phys., 243, 63
 Gonzalez, R., & Woods, R. 2008, Digital Image Processing, 3rd edn. (Pearson/Prentice Hall)
 Hagenaar, H. J., Schrijver, C. J., Title, A. M., & Shine, R. A. 1999, ApJ, 511, 932
 Lamb, D. A., DeForest, C. E., Hagenaar, H. J., Parnell, C. E., & Welsch, B. T. 2008, ApJ, 674, 520
 Parnell, C. E. 2002, MNRAS, 335, 389
 Potts, H. E., & Diver, D. A. 2008, Sol. Phys., 248, 263
 Potts, H. E., Barrett, R. K., & Diver, D. A. 2004, A&A, 424, 253
 Scherrer, P. H., Bogart, R. S., Bush, R. I., et al. 1995, Sol. Phys., 162, 129
 Stangalini, M. 2014, A&A, 561, L6
 Tsuneta, S., Ichimoto, K., Katsukawa, Y., et al. 2008, Sol. Phys., 249, 167
 Wang, H. 1988, Sol. Phys., 116, 1
 Watson, F., Fletcher, L., Dalla, S., & Marshall, S. 2009, Sol. Phys., 260, 5
 Watson, F. T., Fletcher, L., & Marshall, S. 2011, A&A, 533, A14
 Welsch, B. T., & Longcope, D. W. 2003, ApJ, 588, 620
 Welsch, B. T., Fisher, G. H., Abbett, W. P., & Regnier, S. 2004, ApJ, 610, 1148