

DeepVel: Deep learning for the estimation of horizontal velocities at the solar surface[★]

A. Asensio Ramos^{1,2}, I. S. Requerey^{1,2}, and N. Vitas^{1,2}

¹ Instituto de Astrofísica de Canarias, 38205 La Laguna, Tenerife, Spain
e-mail: aasensio@iac.es

² Departamento de Astrofísica, Universidad de La Laguna, 38205 La Laguna, Tenerife, Spain

Received 14 March 2017 / Accepted 21 April 2017

ABSTRACT

Many phenomena taking place in the solar photosphere are controlled by plasma motions. Although the line-of-sight component of the velocity can be estimated using the Doppler effect, we do not have direct spectroscopic access to the components that are perpendicular to the line of sight. These components are typically estimated using methods based on local correlation tracking. We have designed DeepVel, an end-to-end deep neural network that produces an estimation of the velocity at every single pixel, every time step, and at three different heights in the atmosphere from just two consecutive continuum images. We confront DeepVel with local correlation tracking, pointing out that they give very similar results in the time and spatially averaged cases. We use the network to study the evolution in height of the horizontal velocity field in fragmenting granules, supporting the buoyancy-braking mechanism for the formation of intergranular lanes in these granules. We also show that DeepVel can capture very small vortices, so that we can potentially expand the scaling cascade of vortices to very small sizes and durations.

Key words. Sun: granulation – Sun: photosphere – methods: observational – methods: data analysis

1. Introduction

Motions in the solar photosphere are fundamentally controlled by convection in a magnetized plasma. Magnetic field topology is controlled by plasma motions because gas pressure is much higher than magnetic pressure. As a consequence, many of the phenomena taking place in the photosphere are dominated by large-, medium-, and small-scale plasma motions. Among these phenomena, an incomplete list would include the emergence of the magnetic field thanks to convection; tangling of magnetic field lines, which eventually produces reconnection; convective collapse; and cancellation of magnetic fields.

Remotely sensing these three-dimensional velocities is important for the analysis of these events, ideally in combination with spectropolarimetric measurements to infer the magnetic field. The component along the line of sight (LOS) of the velocity can be extracted from spectroscopic observations due to the Doppler effect. However, the components of the velocity field in the plane perpendicular to the LOS cannot be diagnosed spectroscopically. Different algorithms have been used to trace horizontal flows at the solar surface from continuum images (November & Simon 1988; Strous 1995; Roudier et al. 1999; Potts et al. 2004) and also magnetograms (Kusano et al. 2002; Welsch et al. 2004; Longcope 2004; Schuck 2005, 2006; Georgoulis & LaBonte 2006). Among these methods the local correlation tracking (LCT; November & Simon 1988) is most used because of its simplicity and speed. In essence, LCT is a powerful cross-correlation technique for measuring the proper motions of granules. This technique correlates small local

windows in several consecutive images to find the best-match displacement. The tracking window is defined by a Gaussian function whose full width at half maximum (FWHM) is roughly the size of the features to be tracked. In addition, the spatially localized cross-correlation is commonly averaged in time to smooth the transition between consecutive images and reduce the noise induced by atmospheric distortion. All these methods can be considered to give estimations of the so-called optical flow, which is the vector field that needs to be applied to an image to be transformed into a different one. As such, the inferred vector field might not be strictly representative of the inherent horizontal velocity fields.

Given the widespread use of this cross-correlation technique, there have been some efforts to compare the horizontal velocity fields retrieved through LCT with simulated plasma velocities (Rieutord et al. 2001; Matloch et al. 2010; Verma et al. 2013; Yelles Chaouche et al. 2014; Louis et al. 2015). Current three-dimensional magnetohydrodynamical simulations are able to very well reproduce convection in a magnetized plasma, so one expects that the simulated velocities are a good representation of the real velocities in the Sun. These studies have revealed that granules are good tracers for large-scale persistent horizontal flows such as meso- and supergranular flows (e.g., Simon et al. 1988; Muller et al. 1992; De Rosa & Toomre 2004; Yelles Chaouche et al. 2011; Langfellner et al. 2015) or photospheric vortex flows (Brandt et al. 1988; Bonet et al. 2010; Vargas Domínguez et al. 2011; Requerey et al. 2017). The instantaneous velocity fields, which have been obtained by correlating two consecutive frames, also recover the overall morphological features of the flow, but they lack the fine structure observed in the simulated velocities (Louis et al. 2015; Yelles Chaouche et al. 2014). The correlation increases with the

[★] The movie attached to Fig. 3 is available at <http://www.aanda.org>

time average (Rieutord et al. 2001; Matloch et al. 2010) while the LCT-determined horizontal velocities continue to be underestimated roughly by a factor of three (Verma et al. 2013).

In this paper we propose an end-to-end deep learning approach for the estimation of horizontal velocity fields in the solar atmosphere based on a deep fully convolutional neural network. The neural network is trained on a set of simulated velocity fields. Our approach reveals a number of benefits that clearly overcome existing algorithms by a large margin: it is very fast, uses only two consecutive frames, and returns the velocity field in every pixel and for every time step. This is performed at the expense of a time consuming training that only needs to be carried out once.

2. Deep neural networks

Machine learning is a branch of computer science in which models are directly extracted from data and not imposed by the researcher. In essence, the majority of machine learning techniques can be considered to be nonparametric regression techniques that automatically adapt to the existing data and also adapt when new data is added. If these models are sufficiently general, one can apply them to solve complicated inference problems that cannot be easily solved otherwise. One of the first milestones of machine learning was the conception of the perceptron (Rosenblatt 1957), which is a very simple artificial neural network (ANN). Afterwards, ANNs have served many purposes in machine learning. Specially during the 80's and 90's and thanks to several theoretical developments, ANNs were able to solve problems of increasingly difficulty in supervised and unsupervised regression and classification. The discovery that ANNs with a single hidden layer are a universal approximant to any nonlinear function (Jones 1990; Blum & Li 1991) allowed them to be used as a fast substitute on complex inference problems. This was, in large part, facilitated by the development of the backpropagation algorithm (Rumelhart et al. 1986), which enabled the training of neural networks using training examples and computing the effect of the difference between the prediction of the ANNs and the training set on the parameters of the network.

The concept of the ANN had a difficult time during the start of the 21st century for several reasons. First, other techniques with stronger theoretical grounds (for instance, support vector machines and Gaussian processes) allowed the researchers to understand how the methods were fitting the data and how they can be generalized. Second, shallow ANNs only allowed researchers to solve relatively simple problems, and once the networks were made very deep, backpropagation was not able to correctly train them. The reason was that the gradients with respect to the neural network parameters vanish in deep topologies, so that training using conjugate gradient methods stalls. Fortunately, this has radically changed in the last five years thanks to some breakthroughs. First, it was realized that one of the causes for the failure of backpropagation in deep architectures was the usage of activation functions, such as the usual hyperbolic tangent, which produced vanishing gradients during backpropagation. This was solved by using activation functions, such as the rectified linear unit (ReLU; Nair & Hinton 2010) that we use in this work, which do not produce such stalls. Second, fully connected layers were substituted by convolutional layers that apply a set of small kernels to the input and give as output the convolution of the input and the kernels. This induced a reduction in the number of free parameters of the networks without sacrificing any predictive power. Finally, the appearance of graphical processing units (GPU) on the scene allowed researchers to train neural

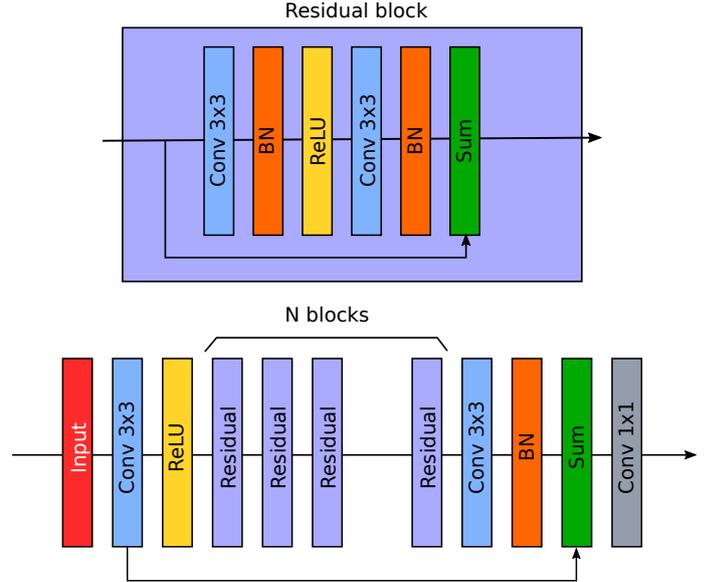


Fig. 1. Representation of residual block and the neural network. *Upper panel:* residual block is shown. *Lower panel:* full architecture of the neural network made of the concatenation of many residual blocks and a skip connection from the input to the output is shown. The selected number of residual blocks for DeepVel is $N = 20$.

networks much faster than was possible before. This also opened the possibility to train the networks via huge training sets. This last point can arguably be considered the main reason for the success of deep learning. Conceptually, deep learning is a set of machine learning techniques based on learning multiple levels of abstraction of the data. If these multiple levels are learned well, deep learning is supposed to generalize well.

In this paper we consider the problem of inferring the horizontal velocity field in the solar surface from two consecutive continuum images. The case of only two images represents the worst case scenario and we could have used more frames for the prediction. However, according to the results we present in the following, we consider that two frames give surprisingly good results. The end-to-end solution given in this paper, which we term DeepVel, is a deep neural network whose topology is described in the following and is trained with velocities extracted from magneto-hydrodynamical (MHD) simulations. The network assumes as input two continuum images of $N_x \times N_y$ in size that are separated by 30 s. The outputs are maps of v_x and v_y at all locations and at three heights in the atmosphere, corresponding to $\tau_{500} = 1, 0.1, 0.01$, where τ_{500} is the optical depth at 500 nm. Only the results at $\tau_{500} = 1$ can be compared with other algorithms such as LCT.

2.1. Deep neural network topology

The deep network that we use has a fully convolutional architecture, which applies a series of convolutions with several small kernels (to be inferred during the training) to the input of every layer. The architecture is graphically represented in Fig. 1. Each colored rectangle in the figure represents a different layer, which we describe in the following:

- *Input* (red): this layer represents the two input images of $N_x \times N_y$ in size. Consequently, this layer represents tensors of $2 \times N_x \times N_y$ in size.
- *Conv 3 × 3* (blue): these layers represent three-dimensional convolutions with a set of 64 kernels (channels) of

$N_{\text{input}} \times 3 \times 3$ in size. We keep the number of kernels and their size fixed because they give very good results with the advantage that convolutions with 3×3 kernels can be made very fast in GPUs. The output tensors of these layers are $64 \times N_x \times N_y$ in size.

- *ReLU* (yellow): these layers represent rectified linear units, which apply the following operation to every pixel and channel of the input: $\text{ReLU}(x) = x$ if $x \geq 0$ and zero elsewhere.
- *BN* (orange): this layer represents batch normalization (Ioffe & Szegedy 2015), which is a trick used to increase the convergence speed of the training. It is based on normalizing the input so that it has zero mean and unit variance, which has been verified to greatly accelerate the training.
- *Sum* (green): this layer describes pixelwise addition between the two inputs.
- *Conv* 1×1 (gray): this layer defines three-dimensional convolution with six $64 \times 1 \times 1$ kernels, which is just a very convenient way to collapse the 64 channels of the last Sum layer of the neural network into the six velocities that we want to predict. The output tensor of this layer is $6 \times N_x \times N_y$ in size.

As seen from the lower panel of Fig. 1, the network is made of the concatenation of N so-called *residual* blocks (He et al. 2016). We choose $N = 20$ for our implementation and did not carry out any hyperparameter optimization, which we leave for the future with the aim of optimizing the network. The internal description of each residual block is shown in the upper panel of Fig. 1. Each block is essentially made of two convolution layers; each convolution layer is followed by batch normalization layers and only the first convolution contains a ReLU activation. Finally, the input of the block is added pixelwise at the end. The main advantage of the residual blocks is that it accelerates the training thanks to the skip connection between the input and output. The output of the set of residual blocks is then transferred through an additional convolutional layer with 64 kernels of 3×3 and a batch normalization layer. The output is then obtained after convolution with 6 kernels 1×1 in size. The total number of free parameters of the network is $\sim 1.6 \times 10^6$.

2.2. Training data and training process

The network is trained via synthetic continuum images from the magneto-convection simulations described by Stein & Nordlund (2012) and Stein (2012). This simulation box is ~ 48 Mm wide in both directions and 20.5 Mm deep, extending from the temperature minimum down to 20 Mm below the visible surface. The simulated solar time spans more than an hour in steps of 30 s. The horizontal resolution turns out to be 48 km with a total size of 1008×1008 pixels. This simulation displays an appropriate balance between the amount of solar surface simulated and the horizontal resolution. The mhd48-1 snapshots that we used are obtained by advecting a uniform field at the bottom boundary. The field is increased until it reached 1 kG at the bottom boundary and then it is kept constant.

We then treated the synthetic images to simulate a real observation. We chose the Imaging Magnetograph eXperiment (IMaX; Martínez Pillet et al. 2011) on board the SUNRISE balloon borne observatory (Solanki et al. 2010; Barthol et al. 2011) as a target. The SUNRISE has a telescope of 1 m diameter and the images that IMaX provides have a spatial sampling of 39.9 km. The spatial sampling of the simulated images used in the training and those of IMaX do not exactly coincide (48 km vs. 39.9 km). However, we demonstrate later than an appropriately trained

network generalizes correctly independent of the size of the structures.

Given that SUNRISE was a balloon mission that observed at a height of ~ 40 km above the Earth surface, the observations are barely affected by the atmosphere. One of the reasons we chose IMaX in our tests is that the instrument has provided long time series of very high-quality diffraction-limited images in both flights (e.g., Lagg et al. 2010; Martínez González et al. 2011; Solanki et al. 2017), which have been used often for LCT studies (e.g., Bonet et al. 2010; Yelles Chaouche et al. 2011; Requerey et al. 2014, 2017). We simulated the effect of IMaX following the approach of Asensio Ramos et al. (2012), which is based on the detailed analysis of Vargas Domínguez (2009). We considered telescope aberrations up to 45 Zernike modes. We considered the amplitudes to be normally distributed with diagonal covariance and a total rms wavefront error (WFE) amounting to $\lambda/9$. These telescope aberrations can be considered to be constant during an observation, so we keep them fixed. The remaining atmosphere was accounted for by considering a wavefront with turbulent Kolmogorov statistics (Noll 1976) with a rms WFE of $\lambda/9$. Although these perturbations are very specific for SUNRISE/IMaX, we think that the neural network trained with these images can be safely applied (or perhaps easily re-trained for different instrumental configurations).

From the available three-dimensional volume of 1008×1008 continuum images for all time steps, we randomly extracted two patches of 50×50 pixels in the same spatial position and separated by 30 s in time. A total of 30 000 such pairs were randomly selected as input for the training. The outputs are the six 50×50 images containing v_x and v_y at $\tau_{500} = 1, 0.1$ and 0.01 , respectively, for each time step. We also generated another set of 1000 samples with the same strategy, which is used as a validation set to avoid overfitting. These are used during training to check that the deep network generalizes correctly and is not memorizing the training set.

The neural network has been developed using the Keras Python library with the Tensorflow backend for the computations. All the training is compiled by Tensorflow to be run in the NVIDIA Tesla K40 and Titan X GPUs¹. The training is carried out by minimizing the squared difference between the output of the network and the velocities in the training set. It is known that optimizing the ℓ_2 norm of the difference might lead to predictions that are too smooth, which is not appropriate for typical uses of deep networks for machine vision in natural images. In the last few years, improvements in this direction have been performed with a second deep network that is used to measure the quality of the prediction. Both networks are trained as a generative adversarial network (GAN; Goodfellow et al. 2014), which results in impressive results (Ledig et al. 2016). We know from the simulations that the horizontal velocity fields are relatively smooth, so we stick with the simpler ℓ_2 norm for our case. We leave the analysis of using GANs for the training for the future.

All inputs are normalized to the median intensity of the quiet Sun (which also needs to be carried out once the network is applied to real observations) and velocities are normalized to the interval $[0, 1]$ using the minimum and maximum velocities in the training set. The optimization is carried out with the Adam stochastic first-order gradient-based optimization algorithm (Kingma & Ba 2014) with a learning rate $\epsilon = 10^{-4}$. As in any stochastic optimization method, the gradient

¹ The trained neural network ready to be applied to solar images, together with the infrastructure to train it with different simulations, can be found in <https://github.com/aasensio/deepvel/>.

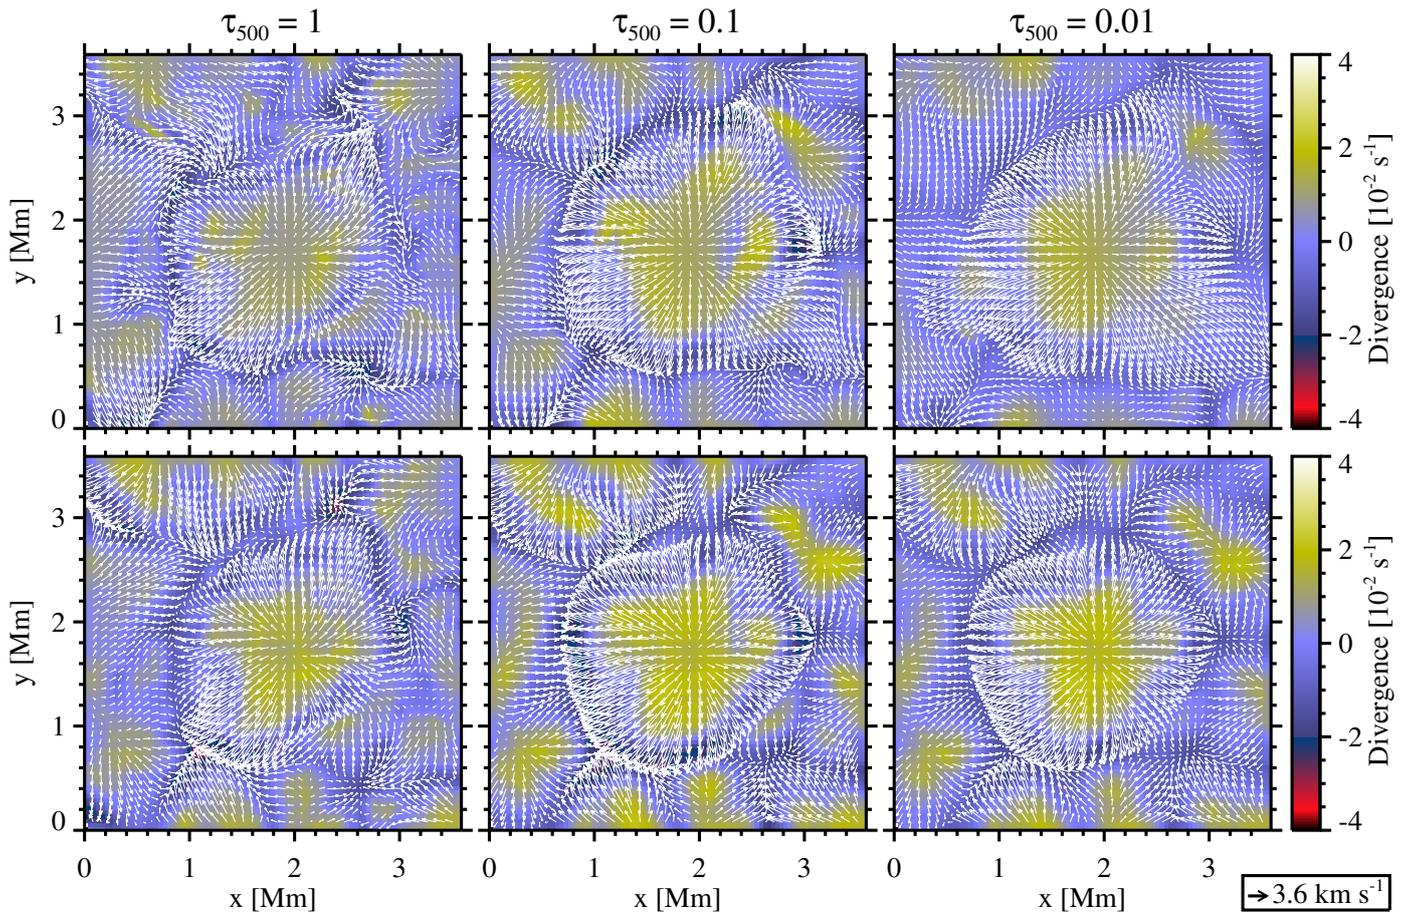


Fig. 2. Instantaneous horizontal velocity field (white arrows) and divergence maps (background images) at three heights in the atmosphere, corresponding to $\tau_{500} = 1, 0.1, 0.01$, for MANCHA (upper row) and DeepVe1 (lower row) velocities.

is estimated from subsets of the input samples, also known as batches. We used batches of 32 samples and trained the network for 30 epochs, where an epoch is finished once all training samples have been used. Therefore, the number of iterations is then 900 000.

2.3. Validation

In absence of a technique similar to DeepVe1 that can be applied to observations, we validated the method using three-dimensional magneto-convection simulations carried out with the MANCHA code (Felipe et al. 2010; Khomenko et al. 2017). The extent of the simulation domain is $24 \text{ Mm} \times 24 \text{ Mm}$ in the horizontal plane and 1.4 Mm vertically with 1152 grid cells in each horizontal direction and 102 uniformly spaced grid points in the vertical direction. The domain is open for the mass flows at the bottom boundary and closed at the top. The radiative transfer losses were computed assuming local thermodynamical equilibrium with precomputed opacities. The magnetic field is initiated through the Biermann battery and amplified by the local dynamo (similar to Vögler & Schüssler 2007). The snapshots used in this study were taken when the total magnetic field reached 10 G at the unit optical depth. The synthetic continuum maps were degraded so that the pixel size is equivalent to those of the training set. We used two consecutive continuum maps to infer the horizontal velocity field and compare it with those extracted from the simulations. The upper panels of Fig. 2 show the velocity

fields of the simulations at three different optical heights in the atmosphere. The underlying map is the divergence of the horizontal field, which is computed as $\nabla \cdot \mathbf{v} = \partial v_x / \partial x + \partial v_y / \partial y$ for $\mathbf{v} = (v_x, v_y)$; positive values indicate diverging flows, while negative values point to converging flows. The lower panels show the results of DeepVe1, which very nicely reproduce the results from the simulation, especially for $\tau_{500} = 1, 0.1$. The results for $\tau_{500} = 0.01$ are slightly less similar although the general appearance is still valuable. Figure 2 only shows a small field of view (FOV), but similar results are found for the rest of the simulated field. Specifically, the velocity field vectors of the whole simulated FOV have a Pearson linear correlation coefficient of 0.82, 0.85, and 0.76 for $\tau_{500} = 1, 0.1$, and 0.01, respectively. Additionally, v_x displays correlation coefficients of 0.82, 0.84, and 0.75 for the same values of τ_{500} , while the figures turn out to be 0.83, 0.86, and 0.78 for v_y for the same τ_{500} heights. We consider that this experiment validates DeepVe1.

3. Results

Once the network is trained, we apply it to real IMAx observations from the first SUNRISE flight. The observational data were obtained on 2009 June 9 from 01:30:54 to 02:02:29 UT, in a quiet-Sun region close to disk center. The 31.6 min length dataset has a temporal cadence of 33.25 s. We point out that the temporal cadence is larger than that used in the training, so that

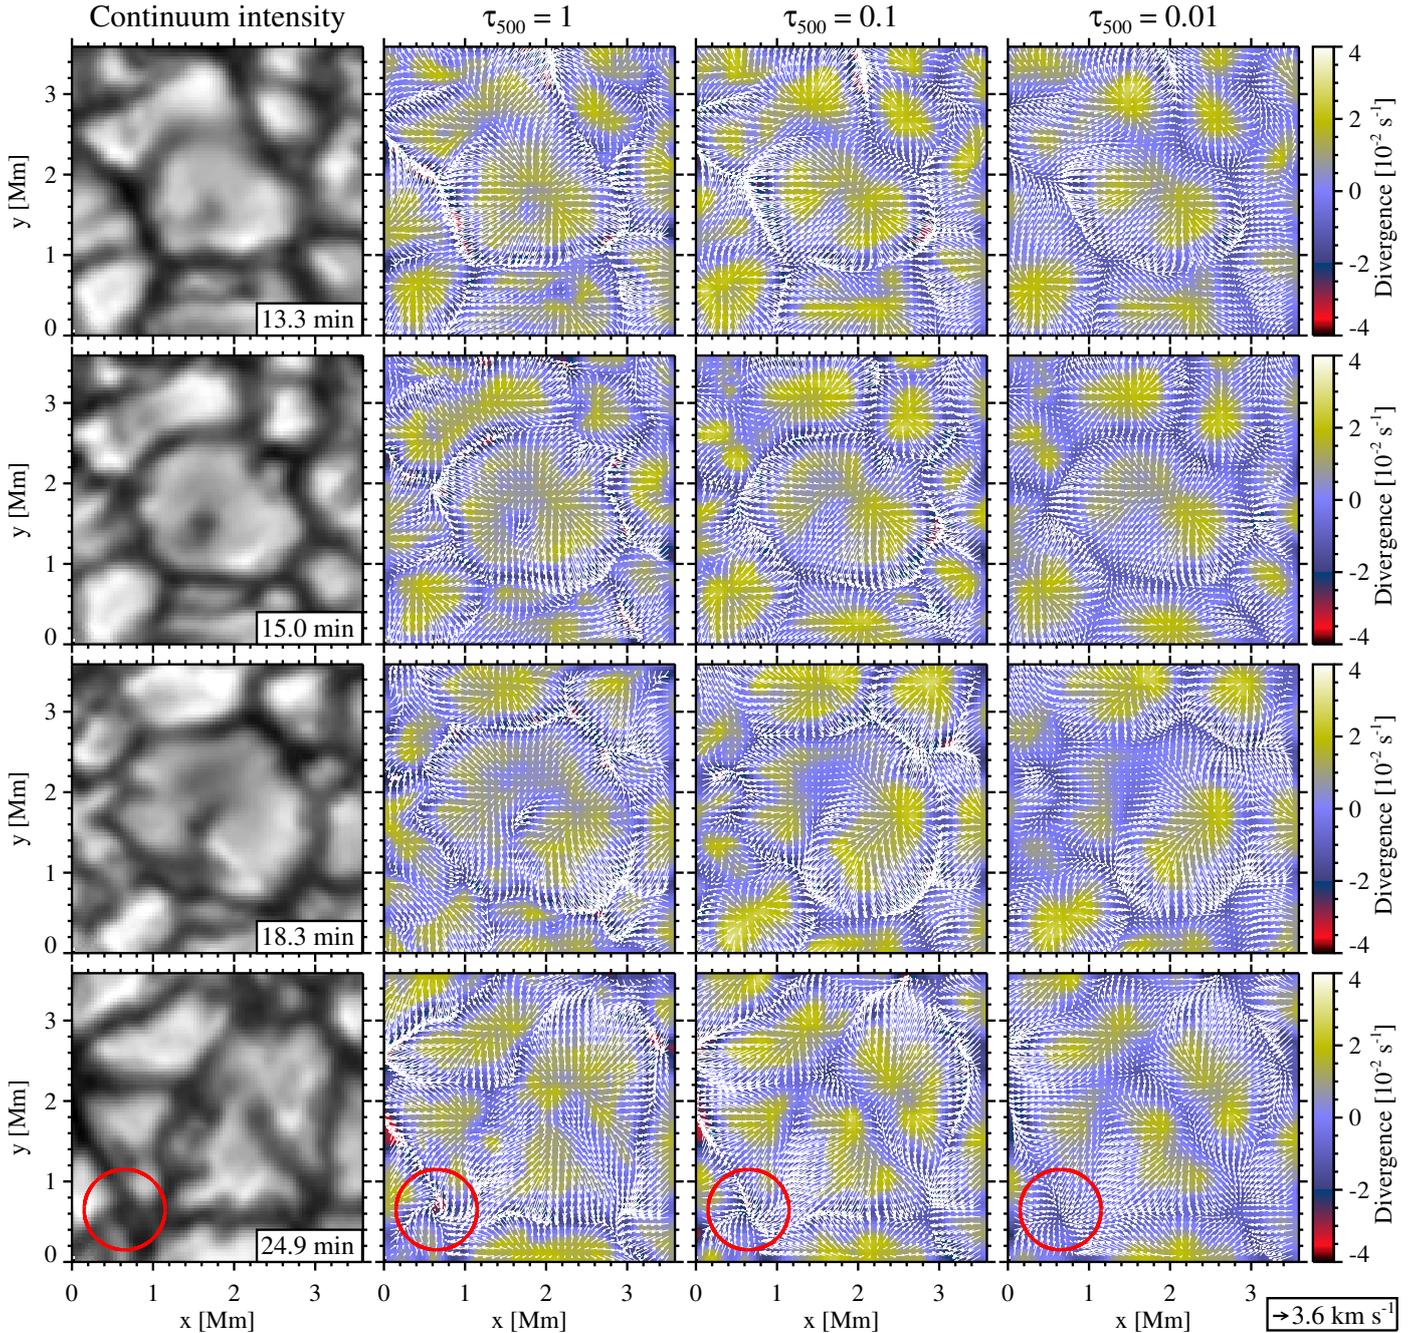


Fig. 3. Evolution of granules as seen in close-ups of continuum intensity maps (*left column*) and the instantaneous horizontal velocity field (white arrows) and divergence maps (background images) at three heights in the atmosphere (*right columns*), corresponding to $\tau_{500} = 1, 0.1, 0.01$. The area is located within the black solid rectangle in Fig. 4a. Red circles (with a radius of 0.5 Mm) in the *bottom row* indicate the location of a small-scale vortex flow. Elapsed time is given in the lower right corner of each continuum image. The arrow at the lower right corner can be used to estimate visually the amplitude of the velocity field. The whole evolution is presented in the [online movie](#).

the network might slightly overestimate the velocities². In addition, the spatial resolution of the observations is slightly better than those of the simulations, but we do not expect large effects. Even though the training was performed with images of 50×50 in size, given the fully convolutional character of the network, we applied it seamlessly to the full FOV of the instrument, which amounts to 736×736 pixels ($29.3 \times 29.3 \text{ Mm}^2$). The computing

² This could have been alleviated if the simulated snapshots had been obtained with the IMAx temporal cadence. We did not follow this path because we wanted to work with public simulations that are available for anyone willing to retrain DeepVel.

time is ~ 2 s per image using a Titan X GPU, and an order of magnitude larger if the computation is performed in a CPU.

3.1. Inferred velocity fields

In Fig. 3 and the online movie³, we show the inferred horizontal velocity field for a small portion of the FOV at four different time steps. The first column shows the continuum image, while

³ The movie can also be obtained from the code repository, located at <https://github.com/aasensio/deepvel>

the rest of columns represent the divergence of the horizontal velocity field at the three different heights in the atmosphere, together with the instantaneous vector field. The results shown in Fig. 3 are impossible to obtain using LCT because of the somehow large spatial and temporal smearing windows that need to be used to increase the correlation and produce robust results.

In general, velocities in lower layers tend to be larger in absolute value and also have larger spatial complexity. Although horizontal diverging flows are similar at the three heights considered, stronger converging flows are seen in intergranular lanes in deeper layers. Additionally, the horizontal size of these zones of converging flows is much smaller in deep layers.

We find the behavior observed in some granules at $t = 15$ min, such as the one at positions $(x, y) = (0.8, 2.6)$ Mm, to be very interesting. This granule appear to be formed by the aggregation of two or more smaller portions. The continuum image of the granule shows a slightly dark lane separating bright regions. This structure is also clearly seen in the velocity field at $\tau_{500} = 0.1$ and $\tau_{500} = 0.01$, while it disappears for deeper layers. It looks like these dark lanes in the middle of granules are a consequence of a converging flow in upper layers that does not reach very deep. Interestingly, the velocity fields before ($t = 13.3$ min) and after ($t = 18.3$ min) help us understand what is happening. We are witnessing a fragmenting granule that is being divided into two by a converging velocity field taking place at higher layers, which later propagates toward lower layers. At $t = 13.3$ min, the division is barely visible in the intensity image, but it is already present at $\tau_{500} = 0.01$. Finally, at $t = 18.3$ min, the granule is divided into two parts with a clear intergranule between them. Our observations seem to be compatible with the buoyancy-braking mechanism (Massager & Zahn 1980; Ploner et al. 1999). In this model, the gas above large granules reduces their upward velocity because of the increase in the mass in upper layers. It loses buoyancy and catastrophically collapses and forms a dark lane if energy losses cannot be compensated. According to our observations, it also develops strong converging velocity flows when eventually forming a new intergranular lane. The fact that DeepVe1 shows this behavior means that this mechanism has to be predominant in the simulations and that the behavior in the upper layers is connected with what is going on at lower layers.

Other instances of fragmenting granules and the appearance of dark structures inside bright granules can be seen in the observations. They all appear to share a similar behavior, except for the specific details. For instance, the dark spot at $(x, y) = (1.6, 1.4)$ Mm at $t = 15$ min seems to reach lower layers slightly faster than the previous example.

3.2. Comparison with LCT

In order to test the performance of DeepVe1, we compare it with the well-known LCT algorithm. The LCT technique recovers horizontal proper motions by tracking intensity features in continuum images. We use a Gaussian tracking window with an $FWHM = 600$ km and then average the cross-correlation function over 30 min. We confront this velocity field with that obtained by DeepVe1 at $\tau_{500} = 1$, temporally averaged over 30 min. In addition, the vector field is spatially averaged with the LCT tracking window size. The upper panels of Fig. 4 show the horizontal velocity arrows and the corresponding divergence maps obtained by DeepVe1 (Fig. 4a) and LCT (Fig. 4b) for the whole FOV of IMAx. We find that the DeepVe1 velocities are 1.15 times larger in magnitude than the LCT ones, which almost perfectly coincides with the overestimation expected because

measurements in IMAx are taken every 33.25 s, instead of the 30 s used in the training (the correction factor would be $\times 1.11$). The velocity fields have Pearson linear correlation coefficients of 0.81 and 0.84 for v_x and v_y , respectively, while it goes down to 0.8 for the whole velocity vector. Additionally, the correlation coefficient for the divergence is 0.71.

This correlation is already evident from the visual inspection of the flow and divergence maps. The lower panels of Fig. 4 show an enlarged view of a smaller region, indicated by a black dashed rectangle in the upper panels. Even though there exist some morphological differences, both maps show the same flow features. In particular, they display an equivalent meso-granular pattern, which is revealed through positive divergence structures on scales between granulation and supergranulation. Such a cellular pattern is commonly found in both observations and hydrodynamical simulations when the LCT technique is applied to intensity images (Simon et al. 1988; Muller et al. 1992; Roudier et al. 1998; Matloch et al. 2010; Yelles Chaouche et al. 2011; Requerey et al. 2017). At the junction of meso-granular cells, smaller structures such as converging flows (Bonet et al. 2010; Vargas Domínguez et al. 2011; Requerey et al. 2017) are also observed. The locations of the convergence centers are indicated by red circles in Figs. 4c and d. Despite their small size, they are equally retrieved in both velocity fields.

The time-averaging increases the correlation between the simulated plasma velocities and the LCT velocities (Rieutord et al. 2001; Matloch et al. 2010). The same applies for the correlation between the smoothed DeepVe1 velocities and the LCT flows. Specifically, we get correlation coefficients of 0.71, 0.75, 0.79, and 0.80 for averaging times of 5, 10, 20, and 30 min, respectively. This results from the fact that the DeepVe1 velocities are a reliable representation of the instantaneous horizontal flow fields, while the LCT velocities are only comparable to plasma velocities typically at timescales longer than the average granule lifetime.

3.3. Small-scale vortex flow

Small-scale vortex flows have been first detected as swirling motions of bright points (Bonet et al. 2008), and later through LCT (Bonet et al. 2010; Vargas Domínguez et al. 2011; Requerey et al. 2017). They have a diameter of ~ 1 Mm (Bonet et al. 2008, 2010; Vargas Domínguez et al. 2011), their lifetime varies from 5 to 20 min (Bonet et al. 2008, 2010), and they appear to be located at meso-granular junctions (Requerey et al. 2017). The spatial distribution and size of such vortices is shown by red circles in Fig. 4d. The temporal and spatial scales are larger than those expected from simulations, where the vortices have lifetimes of only a few minutes (~ 3.5 min) and diameters of ~ 100 km (Moll et al. 2011). This differences are likely due to the smoothing produced by the time and spatial average of the LCT technique.

Figure 5 shows the velocity fields for a close-up of the vortex flow denoted in Fig. 3 by red circles. The underlying maps in the upper panels show the divergence, while the lower panels show the vertical vorticity of the horizontal velocity field, which is defined as $(\nabla \times \mathbf{v})_z = \partial v_y / \partial x - \partial v_x / \partial y$. The vortex has a very small size, surely smaller than 300 km in diameter (see red circle in Fig. 5), and lasts for a very short time, in the range 30–60 s, because it is only clearly visible in one time step and can be guessed in the previous and next frames.

The vortex flow has a central zone with a very strong negative divergence at $\tau_{500} = 1$, reaching values up to -0.03 s^{-1} , which are more than an order of magnitude larger than the

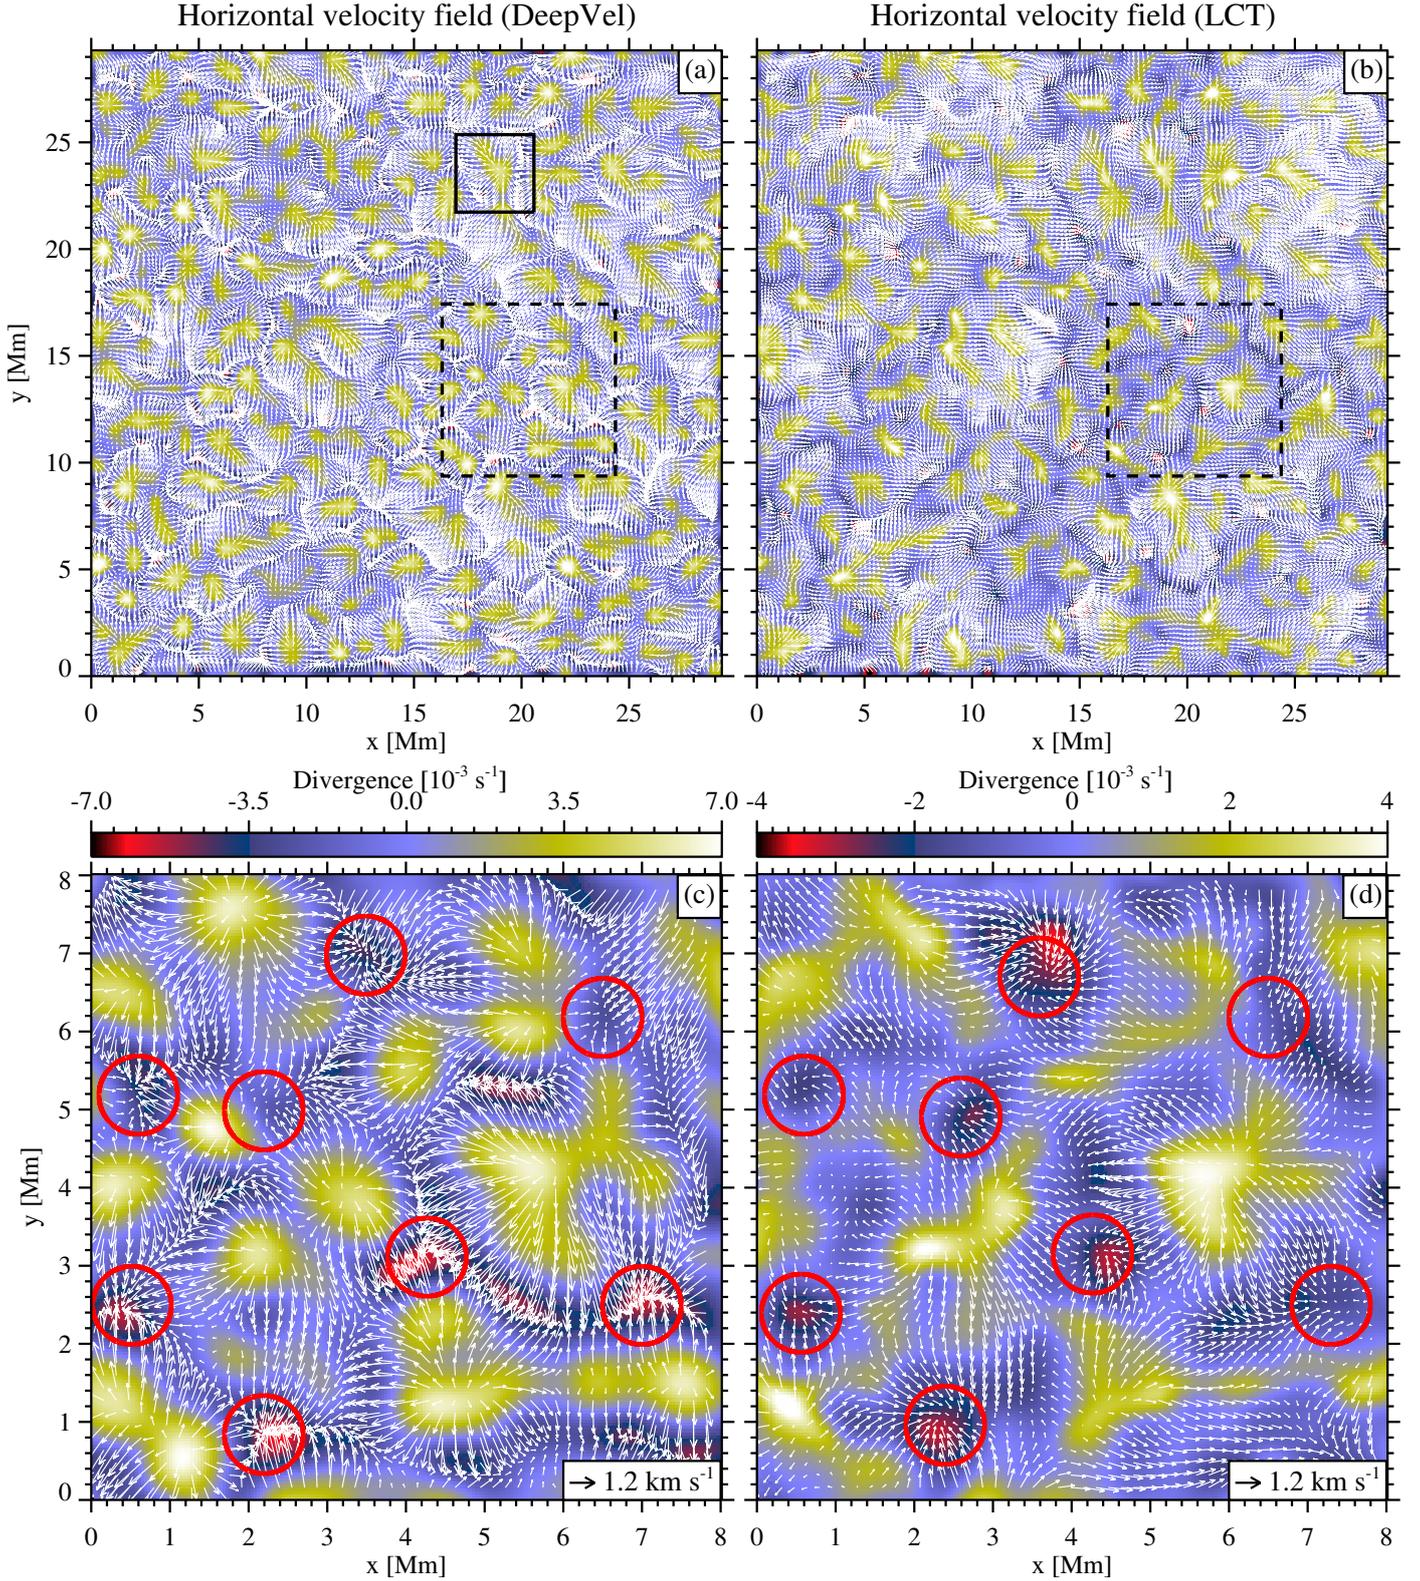


Fig. 4. Average horizontal velocity field (white arrows) and divergence maps (background images) for the FOV of IMAx computed by DeepVel **a)** and LCT **b)**, respectively. The black solid box indicates the region analyzed in Fig. 3. Close-ups of the black dashed rectangles are shown in panels **c)** and **d)** for DeepVel and LCT, respectively. Red circles represent the locations of converging flows.

median value found by Requerey et al. (2017) as a consequence of the much smaller size. The same behavior is seen in the vorticity, with values that also reach -0.03 s^{-1} at $\tau_{500} = 1$, where negative means clockwise rotation. This value of vorticity is an order of magnitude larger than that detected

with LCT (Bonet et al. 2010; Vargas Domínguez et al. 2011; Requerey et al. 2017), but comparable to that found in simulations (e.g., Kitiashvili et al. 2011).

Interestingly, the DeepVel results show a different picture of the vortex flow at higher layers. First, material is still strongly

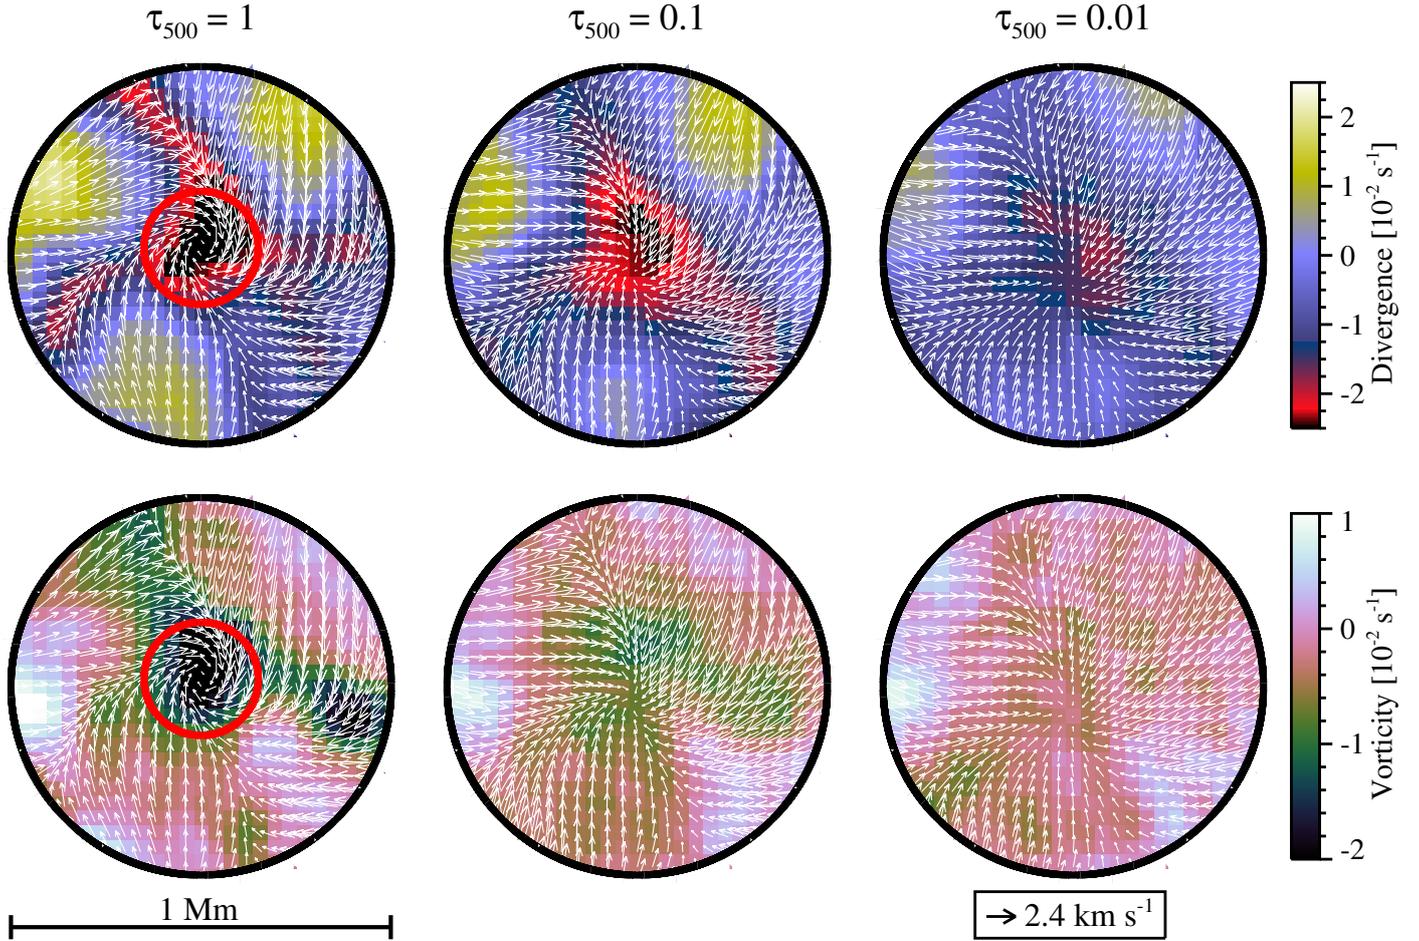


Fig. 5. Close-up of the small-scale vortex flow shown in Fig. 3 with red circles. The *upper panels* show the divergence of the horizontal velocity field at the three selected heights, while the *lower panels* represent the vertical vorticity. Red circles (with a radius of 150 km) show the size of the small-scale vortex flow.

advected toward the center of the vortex with divergences that are still on the order of -0.02 s^{-1} , while the vorticity becomes much smaller in higher layers. This behavior resembles that of the “bathtub effect” (Nordlund 1985), in which the circular velocity is amplified as the plasma contracts with depth.

4. Conclusions

We have developed DeepVe1, which is an end-to-end approach for the estimation of instantaneous and per-pixel horizontal velocity fields based on a deep network. The network is fully convolutional, so that it can be applied to input images of arbitrary size, providing outputs of exactly the same size. In addition, it is very fast, has no parameters to be tuned, and is available to the community as open source. Concerning speed, it can be improved with some architectural changes, but we think that it is already fast enough for our standards.

We checked that the spatially and temporally averaged horizontal velocity field provided by the network is very similar to that obtained with LCT. However, the power of DeepVe1 is that this same information can also be obtained instantaneously, contrary to LCT. Additionally, we provide the velocity field at three different heights in the atmosphere, something that might look counterintuitive at a first look. It is clear from the results presented here (both from simulations and observations) that the network is able to correctly generalize and is not overfitting.

Similar to this application, we expect deep learning to be increasingly applied to solar physics as more high-quality data is obtained and needs to be analyzed. An incomplete list of possible potential applications, which we are already working on, would include fast image deconvolution, fast two-dimensional and three-dimensional spectropolarimetric inversions, and fast control of adaptive optics.

Acknowledgements. We thank B. Ruiz Cobo & F. J. de Cos Juez for very useful comments on an early version of the paper. We also thank R. Abreu for initial discussions on the subject of deep learning. Financial support by the Spanish Ministry of Economy and Competitiveness through projects AYA2014-60476-P Consolider-Ingenio 2010 CSD2009-00038 and ESP2014-56169-C6 are gratefully acknowledged. A.A.R. also acknowledges financial support through the Ramón y Cajal fellowships. We also thank the NVIDIA Corporation for the donation of the Titan X GPU used in this research. We acknowledge PRACE for awarding us access to resource MareNostrum based in Barcelona/Spain. This research has made use of NASA’s Astrophysics Data System Bibliographic Services. We acknowledge the community effort devoted to the development of the following open-source packages that were used in this work: numpy (numpy.org), matplotlib (matplotlib.org), Keras (<https://keras.io>), and TensorFlow (<http://www.tensorflow.org>).

References

- Asensio Ramos, A., Martínez González, M. J., Khomenko, E., & Martínez Pillet, V. 2012, *A&A*, **539**, A42
 Barthol, P., Gandorfer, A., Solanki, S. K., et al. 2011, *Sol. Phys.*, **268**, 1

- Blum, E. K., & Li, L. K. 1991, *Neural Networks*, **4**, 511
- Bonet, J. A., Márquez, I., Sánchez Almeida, J., Cabello, I., & Domingo, V. 2008, *ApJ*, **687**, L131
- Bonet, J. A., Márquez, I., Sánchez Almeida, J., et al. 2010, *ApJ*, **723**, L139
- Brandt, P. N., Scharmer, G. B., Ferguson, S., Shine, R. A., & Tarbell, T. D. 1988, *Nature*, **335**, 238
- De Rosa, M. L., & Toomre, J. 2004, *ApJ*, **616**, 1242
- Felipe, T., Khomenko, E., & Collados, M. 2010, *ApJ*, **719**, 357
- Georgoulis, M. K., & LaBonte, B. J. 2006, *ApJ*, **636**, 475
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., et al. 2014, ArXiv e-prints [arXiv:1406.2661]
- He, K., Zhang, X., Ren, S., & Sun, J. 2016, in 2016 IEEE Conf. Computer Vision and Pattern Recognition, Las Vegas, USA, June 27–30, 770
- Ioffe, S., & Szegedy, C. 2015, in Proc. 32nd Int. Conf. Machine Learning (ICML-15), eds. D. Blei, & F. Bach (JMLR Workshop and Conference Proceedings), 448
- Jones, L. K. 1990, *Proc. IEEE*, **78**, 1585
- Khomenko, E. V., Vitas, N., Collados, M., & de Vicente, A. 2017, A&A, submitted
- Kingma, D. P., & Ba, J. 2014, CoRR, abs/1412.6980
- Kitiashvili, I. N., Kosovichev, A. G., Mansour, N. N., & Wray, A. A. 2011, *ApJ*, **727**, L50
- Kusano, K., Maeshiro, T., Yokoyama, T., & Sakurai, T. 2002, *ApJ*, **577**, 501
- Lagg, A., Solanki, S. K., Riethmüller, T. L., et al. 2010, *ApJ*, **723**, L164
- Langfellner, J., Gizon, L., & Birch, A. C. 2015, *A&A*, **581**, A67
- Ledig, C., Theis, L., Huszar, F., et al. 2016, CoRR, ArXiv e-prints [arXiv:1609.04802]
- Longcope, D. W. 2004, *ApJ*, **612**, 1181
- Louis, R. E., Ravindra, B., Georgoulis, M. K., & Küker, M. 2015, *Sol. Phys.*, **290**, 1135
- Martínez González, M. J., Asensio Ramos, A., Manso Sainz, R., et al. 2011, *ApJ*, **730**, L37
- Martínez Pillet, V., Del Toro Iniesta, J. C., Álvarez-Herrero, A., et al. 2011, *Sol. Phys.*, **268**, 57
- Massager, J. M., & Zahn, J.-P. 1980, *A&A*, **87**, 315
- Matloch, L., Cameron, R., Shelyag, S., Schmitt, D., & Schüssler, M. 2010, *A&A*, **519**, A52
- Moll, R., Cameron, R. H., & Schüssler, M. 2011, *A&A*, **533**, A126
- Muller, R., Auffret, H., Roudier, T., et al. 1992, *Nature*, **356**, 322
- Nair, V., & Hinton, G. E. 2010, in Proc. 27th Int. Conf. Machine Learning, June 21–24, 2010, Haifa, Israel, 807
- Noll, R. J. 1976, *J. Opt. Soc. Amer.*, **66**, 207
- Nordlund, A. 1985, *Sol. Phys.*, **100**, 209
- November, L. J., & Simon, G. W. 1988, *ApJ*, **333**, 427
- Ploner, S. R. O., Solanki, S. K., & Gadun, A. S. 1999, *A&A*, **352**, 679
- Potts, H. E., Barrett, R. K., & Diver, D. A. 2004, *A&A*, **424**, 253
- Requerey, I. S., Del Toro Iniesta, J. C., Bellot Rubio, L. R., et al. 2014, *ApJ*, **789**, 6
- Requerey, I. S., Del Toro Iniesta, J. C., Bellot Rubio, L. R., et al. 2017, *ApJS*, **229**, 14
- Rieutord, M., Roudier, T., Ludwig, H.-G., Nordlund, Å., & Stein, R. 2001, *A&A*, **377**, L14
- Rosenblatt, F. 1957, *Cornell Aeronautical Laboratory Report*, **85**, 460
- Roudier, T., Malherbe, J. M., Vigneau, J., & Pfeiffer, B. 1998, *A&A*, **330**, 1136
- Roudier, T., Rieutord, M., Malherbe, J. M., & Vigneau, J. 1999, *A&A*, **349**, 301
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. 1986, *Nature*, **323**, 533
- Schuck, P. W. 2005, *ApJ*, **632**, L53
- Schuck, P. W. 2006, *ApJ*, **646**, 1358
- Simon, G. W., Title, A. M., Topka, K. P., et al. 1988, *ApJ*, **327**, 964
- Solanki, S. K., Barthol, P., Danilovic, S., et al. 2010, *ApJ*, **723**, L127
- Solanki, S. K., Riethmüller, T. L., Barthol, P., et al. 2017, *ApJS*, **229**, 2
- Stein, R. F. 2012, *Liv. Rev. Sol. Phys.*, **9**, 4
- Stein, R. F., & Nordlund, Å. 2012, *ApJ*, **753**, L13
- Strous, L. H. 1995, in Helioseismology, *ESA SP*, **376**, 213
- Vargas Domínguez, S. 2009, Ph.D. Thesis, Universidad de La Laguna, La Laguna, Spain
- Vargas Domínguez, S., Palacios, J., Balmaceda, L., Cabello, I., & Domingo, V. 2011, *MNRAS*, **416**, 148
- Verma, M., Steffen, M., & Denker, C. 2013, *A&A*, **555**, A136
- Vögler, A., & Schüssler, M. 2007, *A&A*, **465**, L43
- Welsch, B. T., Fisher, G. H., Abbott, W. P., & Regnier, S. 2004, *ApJ*, **610**, 1148
- Yelles Chaouche, L., Moreno-Insertis, F., Martínez Pillet, V., et al. 2011, *ApJ*, **727**, L30
- Yelles Chaouche, L., Moreno-Insertis, F., & Bonet, J. A. 2014, *A&A*, **563**, A93