Solar Physics
DOI: 10.1007/●●●●●-●●●-●●●-●●●●-●

Observables Processing for the Helioseismic and Magnetic Imager Instrument on the Solar Dynamics Observatory

S. Couvidat, J. Schou, P.H. Scherrer, J.T. Hoeksema, Y. Liu, T.L. Duvall, Jr., A.A. Norton, R.S. Bogart, R.I. Bush, R. Shine

Received: 11 May 2015

Abstract

The Solar Dynamics Observatory (SDO) satellite was launched in February 2010 with three instruments onboard, including the Helioseismic and Magnetic Imager (HMI). After two months and a half of commissioning, the HMI started normal operations on May 1, 2010, and has since then been taking observation sequences of the entire solar surface almost continuously. Two observables pipelines have been implemented at Stanford University, to compute line-of-sight and vector magnetic-field observables from the level 1 HMI data. We detail the functioning of these two pipelines, their known issues affecting the resulting physical quantities, and the regular updates to the instrument calibration impacting them. Initial calibration of the HMI was performed on the ground, using a variety of light sources and calibration sequences, and at different locations. However, during the five years of the SDO prime mission, numerous calibration sequences have been taken on orbit with the goal of improving and regularly updating the instrument calibration, and monitoring any change in the HMI. This resulted in several changes in the observables processing that are detailed here.

W.W. Hansen Experimental Physics Laboratory, Stanford University, Stanford, CA 94305

Keywords:

1. Introduction

The HMI prime mission will be completed on April 30, 2015, after five highly successful years of near continuous observations of the full solar disk in the Fe I line at 6173 Å. The extended mission is expected to provide the same level of high quality data to the scientific community. A successor to the Michelson and Doppler Imager (MDI) on board SOHO, the HMI benefited from a more detailed ground calibration which resulted in a better knowledge of some essential properties of the instrument and a better understanding of the data taken. However, the geosynchronous orbit of SDO, unlike the halo orbit of SOHO around the Lagrange L5 point, produces some large daily variations in many physical quantities that have proved difficult to remove from the data. This results in artifacts in most observables. This article reviews these HMI observables: how they are computed, what on-orbit calibration sequences and instrument monitoring steps are taken to ensure that they are produced with up-to-date information, what are their known issues, and what future plans we have to improve them. The line-of-sight and vector magnetic-field observables pipelines have already been presented in other publications: here they are described in greater detail and with any available update at the time of writing. Similarly this paper draws on already published articles based on ground calibration results, but it updates them with on-orbit results. Data processing from the level 0 to the level 1.5 observables and some of the issues with these observables are described in *e.g.*, Schou *et al.* (2012), Couvidat *et al.* (2012), Liu *et al.* (2012), and Hoeksema *et al.* (2014).

Since May 1, 2010, HMI has taken more than 80 million images (as of the end of January 2015) with a resolution of 4096×4096 pixels, and through its two CCD cameras. Close to 800 articles using HMI data are listed on the NASA Astrophysics Data System website (also as of the end of January 2015): the success, scope and breadth of the HMI make it necessary to provide the solar physics community with up-to-date information regarding their processing and what the issues plaguing these observables are. This is required to ensure a better understanding of what can be accomplished and what are the limitations of our data.

In Section 2, we remind the reader of how the observables are computed, both for the line-of-sight (LOS) and vector-field quantities. We provide more information than was already available regarding this observables processing. We also detail any update in the processing pipelines, using on-orbit calibration results. In Section 3, we detail some of the known errors and uncertainties affecting the observables. In Section 4 we mention some other instrumental issues, and what future improvements we plan to implement in the observables pipelines. Finally, we conclude in Section 5.

2. Observables Computation

The HMI observables are also known as level 1.5 data, in contrast to level 0 data (raw HMI images) and the level 1 filtergrams (level 0 images that have been corrected for various effects). These observables are separated into LOS and vector magnetic field.

LOS observables are obtained from data taken by the front camera (LOS camera). They include Dopplergrams, LOS magnetic-field strength (magnetogram), continuum intensity, Fe I line width, and line depth. They are produced in two modes: definitive and near-real-time (NRT). The definitive LOS observables are stored in the following DRMS (Data Record Management System) series: hmi.V_45s, hmi.M_45s, hmi.Ic_45s, hmi.Lw_45s, and hmi.Ld_45s (for, respectively, the Dopplergrams, magnetograms, continuum, line width, and line depth). The NRT LOS observables are stored in the following DRMS series: hmi.V_45s_nrt, hmi.M_45s_nrt, hmi.Ic_45s_nrt, hmi.Lw_45s_nrt, and hmi.Ld_45s_nrt. All are produced with a 45-second cadence, using only the left circular (LCP, Stokes parameter I+V) and right circular (RCP, Stokes I-V) polarizations. Therefore, they represent the projection of different quantities onto the LOS, hence their name.

Vector-field observables, on the other hand, use the full Stokes vector and are obtained from side-camera data. The basic vector-field observable is the Stokes vector itself, obtained at a 12-minute cadence (from a series of 135-second sequences). The Stokes vector is stored in the following DRMS series: hmi.S_720s (for the definitive mode) and hmi.S_720s_nrt (for the NRT mode). From this Stokes vector, the full-vector magnetic field and other quantities are derived using the VFISV inversion code (described in Borrero *et al.*, 2011). Using the Stokes vector, it is also possible to apply the LOS observables algorithm and to produce LOS quantities (by using only the I+V and I-V polarizations). These LOS observables produced from side-camera data are stored in the following DRMS series: hmi.V_720s, hmi.M_720s, hmi.Ic_720s, hmi.Lw_720s, hmi.Ld_720s, and their NRT equivalent. Their main use is as guess values for the VFISV inversions.

All of the observables are computed from level 1 records. The production of these level 1 data is described in Bush *et al.* (2015). These records are stored in the hmi.lev1 (definitive records) and hmi.lev1_nrt (NRT records) DRMS series. Each record contains two data segments: an image (image_lev1) taken by the instrument, and a list of bad pixels (bad_pixel_list) on this image. When the image was taken at a specific wavelength (*i.e.* with the instrument co-tuned), we refer to it as a filtergram. Although this is not part of the observables processing, it is useful to briefly remind the reader of how level 1 data are obtained. From the raw data level 0 images, a dark frame is subtracted and a flat field is applied. The overscan rows and columns of the CCD are removed. A program detecting cosmic ray hits is run, and the potentially affected pixels are listed in the bad pixel list. This list also includes known permanent bad pixels (pixels of the CCD that are known to be defective). Another program, the limb finder, is then run and finds the solar disk center coordinates on the CCD and the observed solar radius. However, the formation height of the signal changes (as is further detailed in

section 2.12), and the solar limb position measured by the HMI changes with the wavelength away from the Fe I line center. Consequently, the radius measured by the limb finder and corrected for the SDO-Sun distance also varies as a function of the difference between target wavelength and wavelength corresponding to the known SDO-Sun radial velocity, OBS_VR. Since the velocities at the east and west solar limbs also differ by ± 2 km s^{-1} due to the solar rotation, an offset appears in the center position of the image. Therefore, the keywords for the disk center location (X0_LF and Y0_LF) and solar radius at the SDO distance (RSUN_LF) returned by the limb finder are corrected by up to half a pixel for the atmospheric height sampled at each wavelength. Corrected values are stored in the CRPIX1, CRPIX2, and R_SUN keywords. Furthermore, the reported plate scale (CDELT1=CDELT2) is consistent with the corrected values. The production of level 1 images from the level 0 ones may take some time, explaining why the level 1 data are produced in two modes (NRT and definitive). NRT records purport to be used in near-real time for various time-sensitive applications (*e.g.* space-weather forecasting), and consequently the hmi.lev1_nrt data are produced within minutes of receiving the raw HMI images. There are minor differences between the NRT level 1 images and their definitive counterparts. NRT level 1 records usually do not list cosmic-ray hits (only the permanent bad pixels are known) because the code computing them requires a time series. Moreover the flat fields applied to the NRT level 1 images may not be up-to-date. NRT level 1 records are used to produce NRT observables, while definitive level 1 records (usually available with a 3 to 4 days delay) are used for definitive-observables processing.

2.1. LOS observables processing

There are two distinct software pipelines for the production of observables. Therefore, two different observables codes are implemented to produce the LOS or Stokes vector quantities. These programs are written in C: HMI_observables.c for the LOS observables, and HMI_IQUV_averaging.c for the Stokes vector (VFISV is fed these Stokes vectors and is not discussed here). The two codes use the C API to access the DRMS and SUMS (storage unit management system) tables in the JSOC database (a PostgreSQL relational database).

HMI_observables.c is the focus of this section. It calls many subroutines written by various HMI team members. Excluding the DRMS functions called to access the JSOC database, HMI_observables.c has currently more than 12000 lines of code.

Here is a description of how it operates. It starts by initializing various variables and testing that the command line parameters entered by the user are within permissible ranges. It initializes the interpolation (spatial and temporal) routines and the polarization calibration routine. It then reads selected keywords of all the level 1 records in a time interval around the target time T_REC at which an observables record is requested by the user. If one of these keywords is missing, cannot be read, or has an unacceptable value, then the code flags the corresponding level 1 record as unusable. If there are no level 1 records in the time interval for which the user requests the observables, then the code exits with

an error message (and no observables record is created, to avoid populating the database with dubious records, in case a user accidentaly requests observables for a T_REC in the future or for a T_REC at which the level 1 data have not been produced yet). If the user requests rotational flat fields for the level 1 records rather than the standard PZT ones, then the code reads the rotational flat field corresponding to the camera and time requested.

Weekly flat fields for both cameras are obtained by offsetting the HMI field of view (FOV) with the piezo-electric transducers (PZT) of the instrument stabilization system (ISS) (Wachter and Schou, 2009). Periodically, observables calculated for a T_REC close to midnight use level 1 filtergrams corrected with different flat fields. These flat fields are stored in the hmi.flatfield DRMS series. Every three months, offpoint flat fields are obtained by moving the legs of the instrument. These offpoint flat fields are used to correct the PZT ones. There exists yet another way of computing flat fields: using the solar rotation itself to smooth HMI images. Such flat fields are referred to as rotational flat fields, and are stored in the hmi.flatfield_update DRMS series. It is possible to apply a rotational flat field to a level 1 image, rather than a PZT one, in the observables code. In practice this option has not been used so far other than for test purpose.

Then, HMI_observables locates a target filtergram: a level 1 record that is as close as possible to the target time T_REC for which the observables record is to be produced and that was taken with the correct wavelength tuning. This wavelength is provided by the user as a command line parameter. From the launch of SDO onwards we have always been using wavelength=3, which corresponds to the HMI filter closest to the core of the Fe I line (at rest) in the blue wing. Once a target filtergram has been located, its image data segment (the 4096×4096 image) is read and the rotational flat field is applied, if needed. Some keywwords of the target filtergram are labeled as reference values. For instance, the focus block used to take the filtergram is used as the reference focus block for the entire run of the observables code. If another filtergram is taken at a different focus block, the observables code will consider this to be an error (as focus blocks are not expected to change during a normal observables run) and will exit. The code then locates another filtergram taken with the same wavelength filter as the target one, to linearly interpolate values of OBS_VR, OBS_VW, OBS_VN, DSUN_OBS, CRLT_OBS, CROTA2, and CRLN_OBS at the target time T_REC. Although it might seem questionable to apply a linear interpolation to an angle (CROTA2 is the negative of the p-angle), in practice this is not an issue as under normal circumstances CROTA2 does not vary significantly between consecutive images and remains close to -180 degrees. Only during roll maneuvers does CROTA2 change. If such a linear interpolation is not possible, the code issues an error message, then creates an empty observables record (with the QUALITY keyword appropriately set to reflect the issue) and finally moves on to the next T_REC.

A gap filling routine (Schou and Couvidat, 2013) is then called on the target filtergram. This routine aims at spatially interpolating over the bad pixels listed in the bad pixel data segment. In the event that the bad pixel list cannot be read, then the observables code exits with an error message.

Both HMI cameras are affected by a small non-linearity in their response to light exposure (see Figure 19 of Wachter *et al.*, 2012). This non-linearity, on the order of 1% for intensities below 12000 DN/s, is corrected separately for the front and side cameras. This correction is computed from the spatially averaged images, and therefore ignores any CCD quadrant difference. A third-order polynomial, based on data obtained from ground-calibration sequences and averaged over the four quadrants of the CCDs, is applied to level 1 intensities. During the course of the prime mission, we changed the values of this polynomial fit (see Section 2.3). The version of the non-linearity correction used in the observables code is indicated by the CALVER64 keyword.

After the non-linearity correction, the observables code retrieves all of the information needed regarding the observables sequence that was run around the T_REC time: number of wavelengths, polarizations, order of the filtergrams, need or not to combine both cameras, etc... If such information is not available or if the sequence that was run at T_REC is not an observables sequence, but rather a calibration one, then the code exits with an error message after creating an empty observables record with the QUALITY keyword set accordingly.

Once the details of the observables sequence are known, then the code loops over each (wavelength,polarization) pair. It retrieves all of the filtergrams needed to produce an image at time T_REC with the specific wavelength and polarization setting. It locates all of the level 1 filtergrams and makes sure that their tuning positions (characterized by the HWLxPOS keywords where x goes from one to four) and their polarization settings (HPLxPOS where x goes from one to three) are as expected. If their data segments have already been read and are in memory, then they can be used immediately. Otherwise the code reads them, it applies the rotational flat field if needed, it gap fills them, and it applies the non-linearity correction. Typically, for a given (wavelength,polarization) pair, HMI_observables requires six level 1 filtergrams in definitive mode and two in NRT mode to interpolate them at the requested T_REC. The two-point temporal interpolation is a basic linear one, while the six-point one uses a specific weighting scheme described in Martinez Oliveros *et al.* (2011).

When enough filtergrams have been read, HMI_observables calls do_interpolate(), which performs several tasks: first, it corrects each and every image for instrumental distortion. The distortion as a function of field position is reconstructed from Zernike polynomials determined during pre-launch calibration and using a random-dot target mounted in the stimulus telescope (Wachter *et al.*, 2012). Section 2.4 describes the instrumental distortion. The routine also corrects the Sun-center coordinates and solar radius keywords as their values have been modified by the distortion correction. Then, it corrects each image for solar rotation. Indeed, each level 1 record used to compute an observables is taken at a slightly different time T_OBS. In between two images, solar rotation moves the solar disk a tiny fraction of a CCD pixel, so that a given pixel on two images does not map the same location on the solar surface. This rotation is corrected at subpixel accuracy using a Wiener spatial-interpolation scheme. The time difference used to calculate the pixel shift is the precise observation time of the filtergram, T_OBS. The routine also re-centers and re-sizes (if needed) all of the images to a common set of values. These values are obtained by averaging those of all

of the level 1 images expected to be used to produce the observables. Finally do_interpolate() performs a temporal interpolation at target time T_REC. When the loop over all wavelengths and polarizations is finished, what is left is a set of filtergrams with the same solar radius, Sun center position, and interpolated in time at the same T_REC.

HMI_observables then calls the polarization calibration routine, polcal(), to return a set of filtergrams which represent true I+V and I-V polarizations, with no cross-contamination: those are called level 1p data. These level 1p records can be saved: this feature was initially thought of as a debugging tool, but proved quite useful for certain scientific purposes. By default, the observables code does not save the level 1p records, as they take a lot of space, but this can be done on an on-demand basis. The polcal() routine requires the temperatures of some HMI components impacting the polarization calibration. Such temperatures are read from the hmi.temperature_summary_300s DRMS series (populated once a day for the entire day through a cron job). The TSEL and TFRONT keywords in the observables series record the temperatures used. In NRT mode, default values are used instead.

After the filtergrams are properly calibrated in polarization, HMI_observables retrieves the look-up table required by the MDI-like algorithm to produce the observables, as well as the polynomial coefficients used to correct the Doppler velocities. These quantites are detailed in later sections. The LUTQUERY keyword in the observables series records the look-up table that was used. The code retrieves two sets of polynomial coefficients from the hmi.coefficients series, usually separated by 12 hours, and linearly interpolates them in time at T_REC before passing them to the MDI-like algorithm.

Finally, HMI_observables calls the routine performing the MDI-like algorithm, Dopplergram() or Dopplergram_largercrop(), to compute Dopplergram, LOS magnetic-field strength, continuum intensity, line width, and line depth from the six wavelengths and two circular polarizations. Dopplergram() was used until January 15, 2014, and Dopplergram_largercrop() has been used ever since. The only difference between these two versions is that Dopplergram_largercrop() computes the observables for a larger crop radius: 90 pixels off the solar limb, rather than the 50 pixels of Dopplergram(). Finally, a statistical routine fstats() is run on the output to compute the statistics keywords: mean, median, rms variation, kurtosis, etc... Appropriate keywords and data segments are written in the observables series, memory is freed, and the code exits.

2.2. MDI-like algorithm

In this section we describe the algorithm implemented in the Dopplergram() and Dopplergram_largercrop() routines. Front camera observables are computed by an MDI-like algorithm, so called because it is based on the one employed to produce the SOHO/MDI observables.

This algorithm has already been described in several publications (*e.g.* Couvidat *et al.*, 2012). Here we provide further details about its implementation in the LOS HMI-observables pipeline. For each of the 16 million — or so — pixels on an HMI image, the MDI-like algorithm starts by estimating the first and second

Fourier coefficients a_n and b_n (with $n = 1$ or $n = 2$) of the Fe I line profile $I(\lambda)$, where λ is the wavelength:

$$a_1 = \frac{2}{T}\int_{-\frac{T}{2}}^{+\frac{T}{2}} I(\lambda)\cos\left(2\pi\frac{\lambda}{T}\right)\mathrm{d}\lambda \; ; \; b_1 = \frac{2}{T}\int_{-\frac{T}{2}}^{+\frac{T}{2}} I(\lambda)\sin\left(2\pi\frac{\lambda}{T}\right)\mathrm{d}\lambda \qquad (1)$$

and

$$a_2 = \frac{2}{T}\int_{-\frac{T}{2}}^{+\frac{T}{2}} I(\lambda)\cos\left(4\pi\frac{\lambda}{T}\right)\mathrm{d}\lambda \; ; \; b_2 = \frac{2}{T}\int_{-\frac{T}{2}}^{+\frac{T}{2}} I(\lambda)\sin\left(4\pi\frac{\lambda}{T}\right)\mathrm{d}\lambda \qquad (2)$$

where T is the period of the observation wavelength span, nominally $T = 6 \times 68.8 = 412.8$ mÅ (*i.e.* six times the nominal separation between two HMI filter transmission profiles). Consequently, the algorithm assumes that the Fe I line profile is periodical with a period T.

MDI was designed so that the FWHM of its filter transmission profiles matches the FWHM of the Ni I line, and the four — equally spaced — wavelength samples cover a period equal to twice this FWHM. Consequently with MDI nearly all of the spectral power in the solar line is captured by the first Fourier coefficients. Unlike MDI, HMI was not designed this way because the dynamic range corresponding to twice the FWHM of the Fe I line is too small to accommodate the large velocity variations resulting from the SDO orbit.

We assume that the Fe I line has the following Gaussian profile:

$$I(\lambda) = I_c - I_d \exp\left[-\frac{(\lambda - \lambda_0)^2}{\sigma^2}\right] \qquad (3)$$

where I_c is the continuum intensity, I_d is the line depth, λ_0 is the Doppler shift, and σ is a measure of the line width (FWHM$= 2\sqrt{\log(2)}\,\sigma$).

The Doppler velocity $v = dv/d\lambda \times \lambda_0$ can be expressed as:

$$v = \frac{\mathrm{d}v}{\mathrm{d}\lambda}\frac{T}{2\pi}\,\mathrm{atan}\left(\frac{b_1}{a_1}\right) \qquad (4)$$

where $dv/d\lambda = 299792458.0$ m s^{-1} /6173.3433 Å$= 48562.4$ m s^{-1} Å^{-1}. The second Fourier coefficients could also be used:

$$v_2 = \frac{\mathrm{d}v}{\mathrm{d}\lambda}\frac{T}{4\pi}\,\mathrm{atan}\left(\frac{b_2}{a_2}\right) \qquad (5)$$

The line depth $[I_d]$ is equal to:

$$I_d = \frac{T}{2\sigma\sqrt{\pi}}\sqrt{a_1^2 + b_1^2}\exp\left(\frac{\pi^2\sigma^2}{T^2}\right) \qquad (6)$$

while σ is equal to:

$$\sigma = \frac{T}{\pi\sqrt{6}}\sqrt{\log\left(\frac{a_1^2 + b_1^2}{a_2^2 + b_2^2}\right)} \qquad (7)$$

However, HMI samples the iron line at only six points and therefore what we compute is a discrete approximation to the Fourier coefficients, rather than the actual coefficients. For instance:

$$a_1 \approx \frac{2}{6}\sum_{j=0}^{5} I_j \cos\left(2\pi\frac{2.5-j}{6}\right) \tag{8}$$

The b_n are determined by a similar formula with cosine replaced by sine. In the LOS observables code, these a_n and b_n are calculated separately for the LCP (I+V) and RCP (I-V) polarizations. Applying Equation (4) returns two velocities: $v_{\rm LCP}$ and $v_{\rm RCP}$.

Departing from the assumptions made, the actual Fe I line profile is not Gaussian (*e.g.* see Figure 19). Moreover, the discrete approximations to a_n and b_n are not accurate due to a reduced number of sampling points and because the HMI filter transmission profiles are not δ-functions (consequently the observables calculated are relative to the Fe I line convolved by the filters). Therefore, $v_{\rm LCP}$ and $v_{\rm RCP}$ need to be corrected.

This is the role of look-up tables. They are obtained from a realistic model of the Fe I line at rest (and in the quiet Sun) and from calibrated HMI filter transmission profiles. Look-up tables are described in Section 2.13, but here is a brief overview: the model line profile is shifted in wavelength to simulate a Doppler velocity. At each shift, the line profile is convolved by the filter transmittances. The MDI-like algorithm is then applied to the simulated HMI intensities. The velocity returned by the algorithm is a function of the actual (input) Doppler velocity. The inverse function is called a look-up table (a misnomer that is a legacy of the MDI implementation). The tables vary across the HMI CCDs. In Dopplergram() and Dopplergram_largercrop() they are linearly interpolated at $v_{\rm LCP}$ and $v_{\rm RCP}$ to derive corrected Doppler velocities $V_{\rm LCP}$ and $V_{\rm RCP}$. We tested a quadratic interpolation, but this did not improve the performances. The standard look-up tables have the following dimensions: $1642 \times 128 \times 128$ where the first dimension is the number of test velocities (times two because the look-up tables store both the first and second Fourier velocities), and the second and third dimensions are the x and y CCD locations. They are stored in several different DRMS series: hmi.lookup, hmi.lookup_corrected, hmi.lookup_expanded, and hmi.lookup_corrected_expanded where "corrected" refers to the fact that the phase maps have been corrected for an interference fringe pattern and "expanded" refers to look-up tables computed for a larger off-limb radius.

Calibration of the HMI filters shows some residual errors (at the percent level) on their transmittances, resulting in imperfect look-up tables. Conversely, the SDO orbital velocity is known very accurately and can be used to, partly, improve these tables. In the HMI pipeline this additional step is referred to as polynomial correction. Twice a day, 24-hour worth of Sun-SDO radial velocity OBS_VR minus the median velocity RAWMEDN accross the solar disk and returned by the MDI-like algorithm (and after applying the look-up tables) are fitted by a third-order polynomial as a function of RAWMEDN. The coefficients of this polynomial are stored in the hmi.coefficients DRMS series as a function of time. Prior to calling Dopplergram() or Dopplergram_largercrop(), HMI_observables

reads two sets of these polynomial coefficients for times before and after the target time T_REC, and performs a linear interpolation of their values at T_REC. The result is sent to Dopplergram() or Dopplergram_largercrop() and is used to further correct $V_{\rm LCP}$ and $V_{\rm RCP}$. The polynomial correction is imperfect, as OBS_VR only covers a relatively small velocity range. In NRT mode, the observables code uses the most recent polynomial coefficients even if they were computed for days old T_REC.

Finally, the resulting $V_{\rm LCP}$ and $V_{\rm RCP}$ velocities are combined to produce a Doppler-velocity estimate $[V]$:

$$V = \frac{V_{\rm LCP} + V_{\rm RCP}}{2} \tag{9}$$

while the LOS magnetic-field strength $[B]$ is estimated as:

$$B = (V_{\rm LCP} - V_{\rm RCP}) K_m \tag{10}$$

where $K_m = 1.0/(2.0 \times 4.67\ 10^{-5}\, \lambda_0\, g_L\, c) = 0.231$ G m^{-1} s, $g_L = 2.5$ (Norton *et al.*, 2006) is the Landé g-factor, and c is the speed of light.

An estimate of the continuum intensity $[I_c]$ is obtained by reconstructing the solar line from the Doppler-shift, line-width, and line-depth estimates:

$$I_c \approx \frac{1}{6} \sum_{j=0}^{5} \left[I_j + I_d \exp\left(-\frac{(\lambda_j - \lambda_0)^2}{\sigma^2} \right) \right] \tag{11}$$

where λ_0, I_d, and σ are values retrieved by Equations (4), (6), and (7), and λ_j are the nominal wavelengths corresponding to each filter profile.

The algorithm is actually implemented slightly differently in the HMI pipeline. Tests on artificial Gaussian lines and using proper HMI filter transmittances showed that the theoretical algorithm overestimates the line width of Gaussian lines by $\approx 20\%$ for a line with $I_d = 0.62$ and $\sigma = 0.0613$ Å (values in Norton *et al.*, 2006). Conversely, the line depth is underestimated by $\approx 33\%$. The continuum intensity seems only slightly underestimated (by $\approx 1\%$). We surmise that these errors on the parameters of artificial Gaussian lines arise because the number of wavelength samples is too small and the filters are not δ-functions. Unlike velocity shifts (and therefore magnetic-field strength) the line width and depth are not corrected by look-up tables.

Due to an initial error in implementation, in Dopplergram() and Dopplergram_largercrop() a factor $K_1 = 5/6$ multiplies σ returned by Equation (7), while a factor $K_2 = 6/5$ multiplies I_d returned by Equation (6). The integral of a Gaussian is proportional to $I_d\, \sigma$: therefore, multiplying these two quantities by K_1 and $K_2 = 1/K_1$ keeps the integral constant and produces the same continuum intensity with Equation (11). In the current implementation, the line depths and line widths returned by Equations (6) and (7) are still intentionally multiplied by K_2 and K_1, respectively, so that both values are expected to be closer to the actual ones, while the continuum intensity remains unchanged.

Finally, when computing line depth and continuum intensity with Equations (6) and (11), the σ used in these equations is not the one derived from Equation

(7) and that is saved in the observables series. A fifth-order polynomial was fitted to an azimuthal average (about the solar disk center) of an HMI line-width map obtained from Equation (7) (and corrected by K_1) during a period of low solar activity, as a function of center-to-limb distance. Line widths calculated from this polynomial are used instead of those from Equation (7) because the latter sometimes returns spurious values in presence of locally strong magnetic fields, especially for pixels away from solar disk center.

2.3. Non-linearity

Wachter *et al.* (2012) derived the non-linearity of the HMI CCDs from ground-calibration data. The intensity (in DN) on a given pixel does not vary linearly with the number of photons. Because, amongst others, of the daily change in the Sun-SDO distance, the number of photons received by each CCD pixel is not constant for a given exposure time and therefore the lack of linear response must be corrected. The non-linearity correction is implemented in the observables codes (both HMI_observables and HMI_IQUV_averaging), on each level 1 image used. The initial correction was based on the results of Wachter *et al.* (2012). The actual intensity minus the linear one was fitted as a function of the linear intensity by a 3rd-order polynomial. The coefficients of this polynomial were: -8.28, 0.0177, -3.716^{-6}, and 9.014^{-11} for the side camera; -11.08, 0.0174, -2.716^{-6}, and 6.923^{-11} for the front camera. After January 15, 2014, other coefficients have been used. The main reason is that a negative value for the zeroth-order coefficient (intercept term) in the original values means that some pixels ended up with a negative (albeit small) intensity, which does not make physical sense. The coefficients used after January 15, 2014 are: 0, 0.0254, -4.009^{-6}, and 1.061^{-10} for the side camera; 0, 0.0207, -3.187^{-6}, and 8.754^{-11} for the front camera. Each CCD quadrant has a slightly different non-linear response, but we measured the spatial average over the entire CCD. Calibration sequences are taken regularly on-orbit to check that the non-linearity of the CCDs does not vary with time. So far, this non-linearity has proven constant. Figure 1 shows the result of the analysis of a typical non-linearity sequence dated from October 16, 2013.

2.4. Distortion correction

As previously mentioned, the instrumental distortion is corrected in do_interpolate(). This correction is based on Zernike polynomial coefficients measured from ground data prior to the SDO launch (Wachter *et al.*, 2012). The Venus transit of June 5-6, 2012 gave us a unique opportunity to test the accuracy of this instrumental distortion model. Using almost seven hours of side-camera images taken in linear polarization and in true continuum (with FIDs of 10004, 10005, 10006, and 10007), and after the instrumental distortion was removed using the same procedure as in HMI_observables, the trajectory of the center of Venus across the solar disk was computed during its transit. A PSF estimate (see Section 4.4) of HMI was deconvolved from the images with a Richardson-Lucy algorithm, in order to improve the estimates of the Venus-center locations. By comparing the

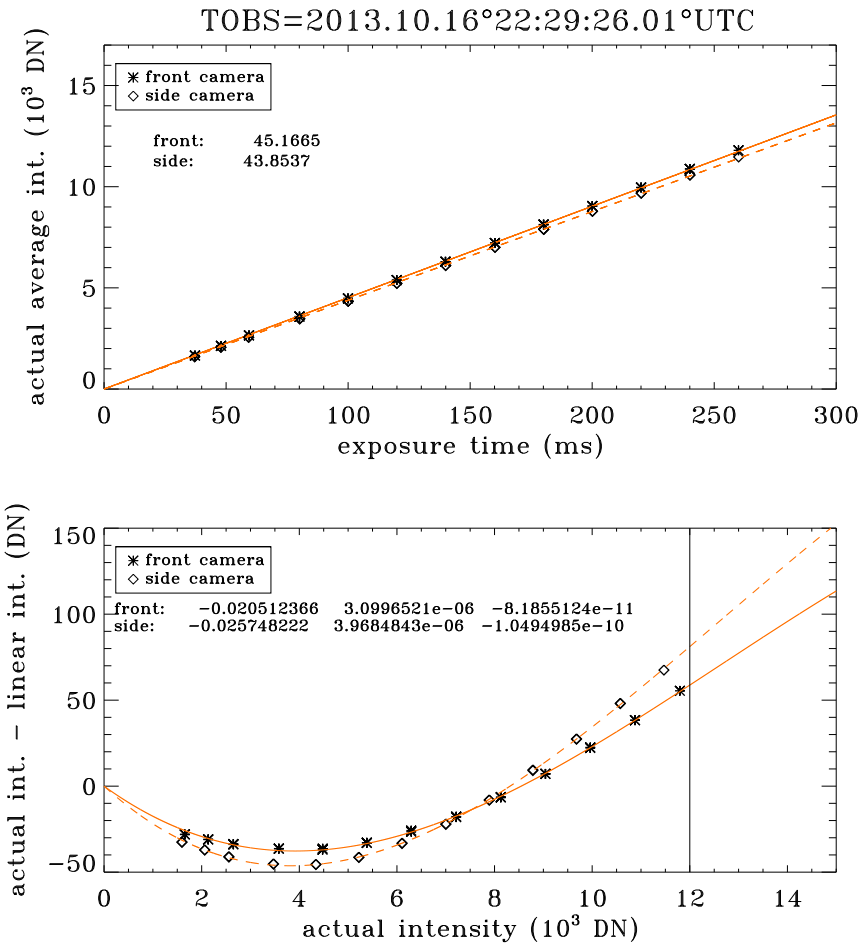

Figure 1. Results of the non-linearity analysis for the front and side cameras, based on data taken on October 16, 2013.

measured Venus trajectory with the theoretical one returned by an ephemeris provided by the Goddard Space Flight Center, it was possible to estimate the residual distortion along the path of Venus, as well as the actual plate scale (CDELT1) and p-angle (CROTA2). CROTA2 and CDELT1 are fitted with a non-linear least-squares fit algorithm (Levenberg-Marquardt) to minimize the Venus-center differences between ephemeris and HMI data.

From Figure 10 the residual distortion across the solar disk along the path of Venus, and for the side camera, is less than 0.1 pixels in the X direction, and less than 0.15 in the Y direction. This confirms that the instrumental distortion model is quite accurate, at least for the pixels lying on the Venus path.

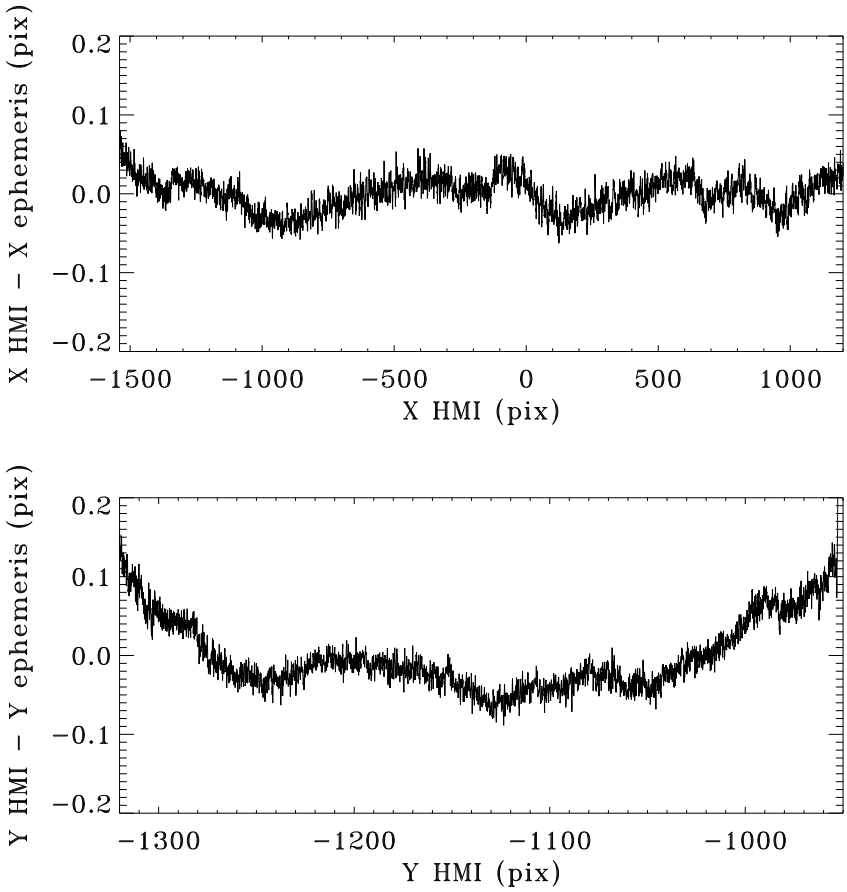

Figure 2. Residual instrumental distortion along the path of Venus during its solar transit of June 2012.

2.5. Fringe removal

Another correction, implemented outside the observables codes but with effect in the call to Dopplergram() or Dopplergram largercrop(), is the removal of interference fringes on the computed level 1.5 images. The front window of HMI acts as a weak Fabry-Perot interferometer due to its multiple layers made of different glass and glue (and thus different refractive indices). This results in interference fringes on Calmode images, as the front window is in focus in this mode. Phase maps of the tunable elements are produced in Calmode and therefore also exhibit interference fringes. As the phase maps are used to derive the look-up tables for the MDI-like algorithm (see Section 2.13), the interference fringes bleed onto line-of-sight observables at the time of their computation in Dopplergram() or Dopplergram largercrop(). For helioseismic purposes this is mainly a cosmetic issue, that should not have a scientific impact. On the other hand it does substantially impact secondary objectives, such as attempts at

determining the surface flows directly from the Doppler shift. We did manage to mostly correct the phase maps, as described below. The original (uncorrected) phase maps are stored in the DRMS series hmi.phasemaps, while the corrected ones are stored in hmi.phasemaps corrected. Each (roughly bi-weekly) phasemap record contains five images: the phase maps of the three tunable elements, and the linewidth and linedepth maps of the Fe i line fitted together with the phases.

To model the fringes we start by assuming that they can, in each of the five variables, be written as

$$A(x, y, t) = A_0(x, y) + A_L(x, y)t + A_C(x, y)cos(phi(t)) + A_S(x, y)sin(phi(t))$$

A_0 represents a constant term, A_L a term describing the overall drift, A_C and A_S describe the fringes, all arbitrary functions of space, and ϕ describes the phase of the fringes (effectively the glass thickness) as a function of time. As it turns out the fringe is most cleanly visible in the line depth, so we start by fitting Eq. 2.5 to that. The fit is performed by using the terms of an SVD as the initial guess and alternately fitting ϕ and the spatial terms.

Having determined ϕ, Eq. 2.5 is now fitted to each of the variables and the corrected variables are determined by subtracting the fringe term $A_C(x, y)cos(phi(t))+$ $A_S(x, y)sin(phi(t))$ from the original values.

Clearly, modeling the fringes like this is far from a perfect, but the result is nonetheless that the amplitude of the large scale fringes is dramatically reduced.

Unfortunately this leaves behind a number of smaller scale fringes. Repeating the above procedure (replacing the A terms with equivalent B terms and ϕ with ϕ_1) to remove these is not nearly as efficient as for the large fringe, but does nonetheless improve the results significantly, so these phase maps are the ones finally applied.

Figure 18 shows the A_C, A_S, B_C and B_S terms.

Figure 4 shows the impact of the correction applied: the left panels shows the raw tunable-element phase maps, while the right panels show the same phase maps after the interference fringes have been removed. It appears that large-scale fringes have been almost completely removed by the correction but small-scale fringes are still present, albeit with a smaller peak- to-peak amplitude.

It is interesting to note that $\phi_1 \approx 2.25\phi$. Given the thicknesses of the different glass elements, one might naively have expected a factor of 2.00 or 2.50, the deviation presumably being due to the different thermal expansion coefficients and/or different changes in the refractive index with temperature.

More elaborate correction schemes were also attempted, but without much success.

We usually compute a new correction each time we re-tune the HMI. The reason is that a re-tuning requires computing new look-up tables for the MDI-like algorithm. Therefore this is a convenient time to update the interference-fringe correction.

2.6. I-ripple correction

An I-ripple is an intensity variation in the HMI output depending on the instrument tuning, especially visible when using a uniform and constant light

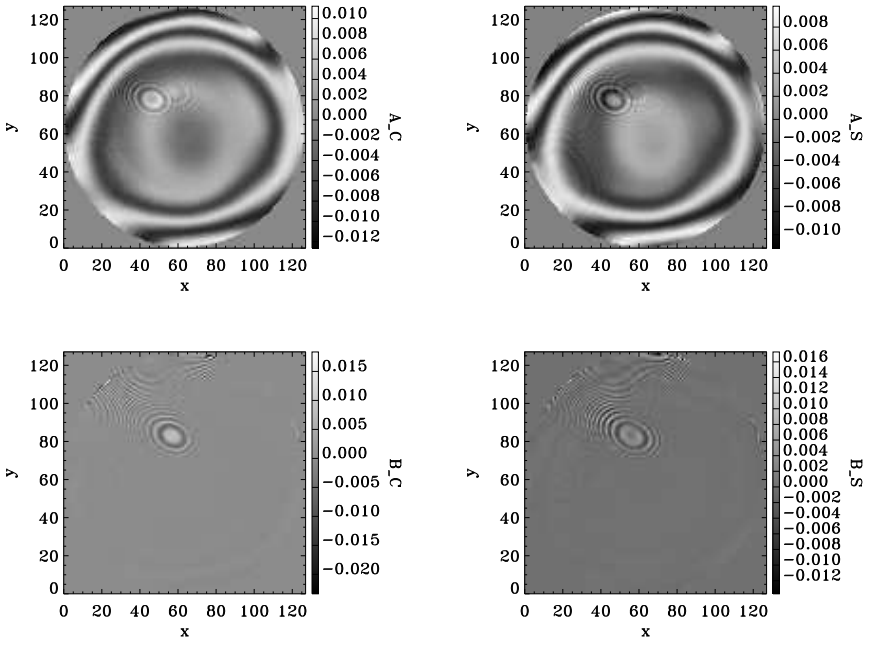

Figure 3. First four singular vectors

source. It results from defects in the tunable filter elements, like misalignements in wave-plates. When deriving the filter transmission profiles, I-ripples are not accounted for. This introduces a small error in these profiles and therefore in the look-up tables used by the MDI-like algorithm in Dopplergram() or Dopplergram_largercrop().

For a specific tuning phase ϕ of the tunable Lyot element E1, its transmitted intensity can be modeled as:

$$\frac{I(\lambda)}{\bar{I}(\lambda)} = K_0 + [K_1 \cos(\phi/2) + K_2 \sin(\phi/2)]^2 \tag{12}$$

where $\bar{I}$ is the average intensity over all tuning phases ϕ possible, and K_0, K_1, and K_2 characterize the I-ripple. This specific equation was derived for an I-ripple resulting from an issue in the Lyot half-wave plate, like a misalignement, but it proved to be equally good for other I-ripple inducing issues, like a combination of a tilt in the entrance polarizer and quarter-wave plate of a Michelson interferometer. Therefore this same equation is also used to model the I-ripples of the two Michelson interferometers.

Figure 5 shows the temporal evolution of the I-ripples (in term of peak-to-peak amplitudes) for the three tunable elements since the beginning of the mission. It is obtained by fitting the I-ripples of each tunable elements on the intensities of a detune sequence and using Equation 12. The peak-to-peak variation in the transmitted intensity I is computed as $K_1^2 + K_2^2$.

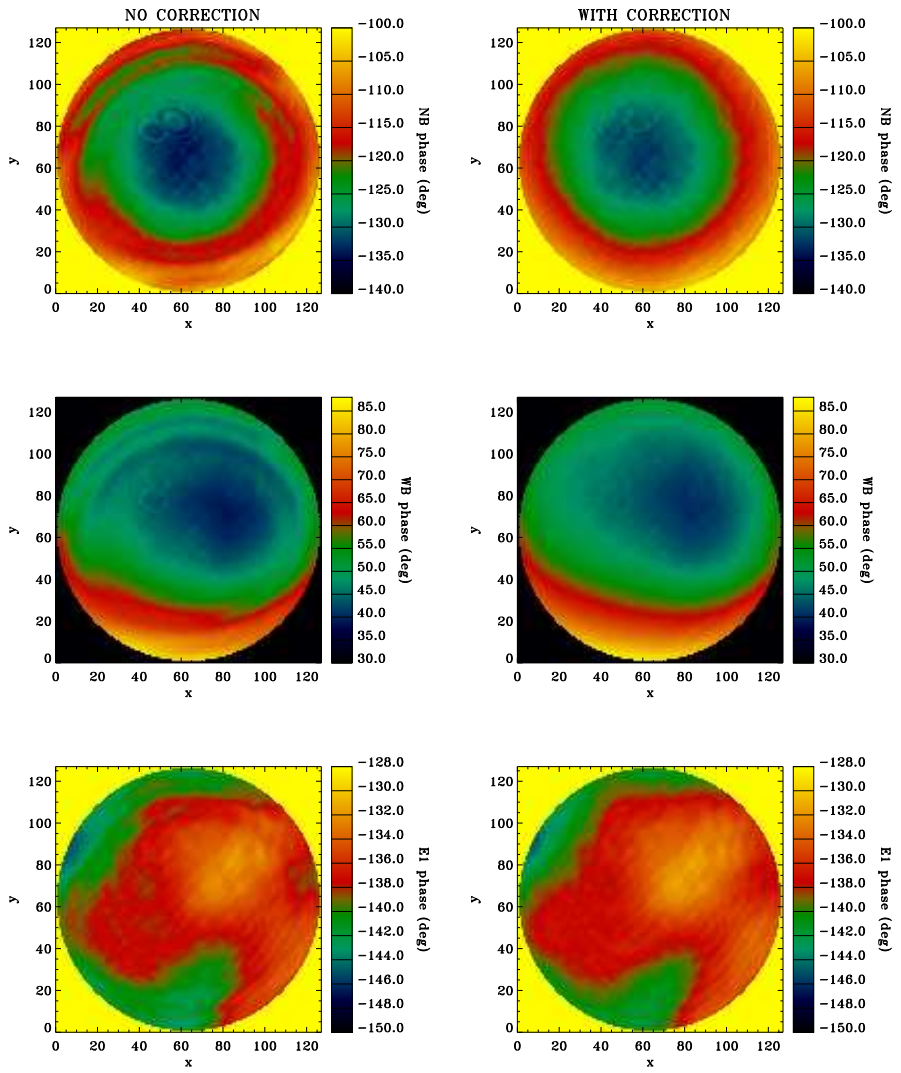

Figure 4. Impact of the interference fringe removal on the phase maps of the tunable elements.

Unfortunately, it is not clear whether the result of each fit only includes I-ripple effects or if other time-dependent defects in the tunable elements — and that are not included in our transmittance model — also bleed onto these results. Regardless, Figure 5 shows that the intensity transmitted by the tunable elements varies with time. Currently, the I-ripple is not taken into account when computing the observables. Including this effect when fitting the detune sequences does improve the goodness of fit (as is expected from the addition of extra parameters), but including the I-ripple in the filter transmission profile does not positively impact some of the issues we encountered with the observ-

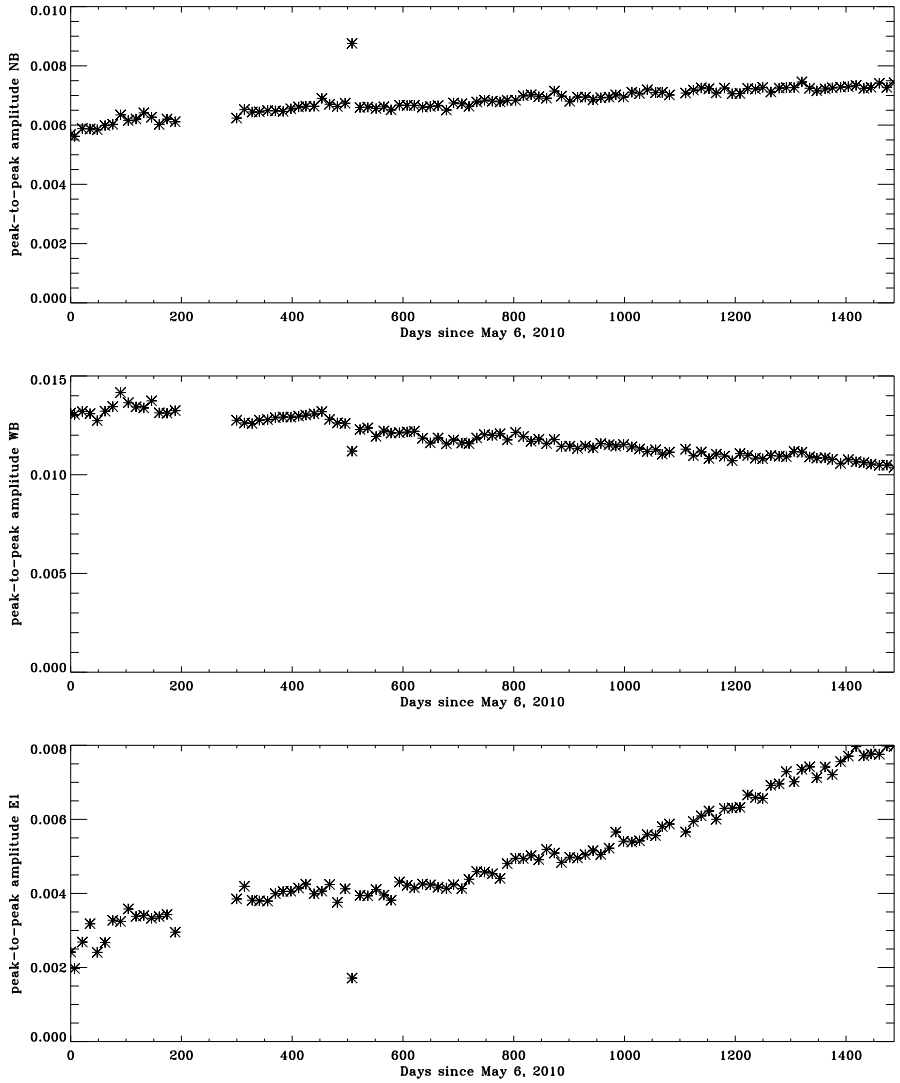

Figure 5. Peak-to-peak amplitudes of the tunable-element I-ripples.

ables (especially the 24-hour oscillations detailed in Section 3.1). The C code lookup_Iripple.c can compute look-up tables with the I-ripple taken into account.

2.7. Image alignment in the observables code

The production of definitive LOS observables typically involves 72 level 1 filtergrams, while the production of a definitive 12-minute averaged Stokes vector requires 360 level 1 filtergrams. All of these filtergrams have slightly different Sun-center location coodinates (X0_LF,Y0_LF) and (CRPIX1,CRPIX2). They

may also have slightly different p-angles (the negative of CROTA2). Therefore, before performing a temporal interpolation at T_REC, the do_interpolate() routine of the observables codes re-register (re-align) each level 1 filtergram. First, it rotates each image around its center to a common CROTA2 value (p-angle) and a common CRLT_OBS (b-angle). Then it takes into account the effect of the solar differential rotation to spatially interpolate the Sun as it should be at a given T_REC. The spatial interpolation is based on Wiener interpolation, which requires knowledge of the correlation function of the data. Finally, the de-rotated images are re-centered to a common pair (X0_LF,Y0_LF) that is obtained by averaging all of the level 1 filtergrams used to produce the observables.

2.8. Limb finding and Sun center error

When processing level 1 images to obtain the level 1.5 observables, the instrumental distortion is removed on each and every image. Rather than re-running the limb finder, once the distortion has been corrected for, an analytical correction is applied to the Sun-center position and the solar radius measured by the limb finder on the level 1 records. Running the limb finder after distortion is removed was deemed too cumbersome as it would slow down the observables processing. Unfortunately, the analytical correction applied to X0_LF, Y0_LF, and RSUN_LF is not perfect. For Y0_LF the difference between the value inferred from the analytical correction and the value returned by the limb finder applied to the un-distorted image reaches about 0.13 to 0.17 pixels (depending on the FID of the level 1 record). For X0_LF, this difference is roughly ten times smaller, and for RSUN_LF, it is even less (below a hundredth of a pixel). Consequently, the CRPIX2 keyword in the observables records is off by roughly 0.15 pixels. Figure 6 shows the difference in Y0_LF between the analytically corrected value and the one returned by the limb finder after distortion correction of the level 1 records. Figure 7 shows the difference in X0_LF, while Figure 8 shows the difference in RSUN_LF.

2.9. IQUV generation

The Stokes-vector observables code HMI_IQUV_averaging.c produces averaged I, Q, U, and V images at six wavelengths on a regular 12-minute cadence centered at the time given in the data series keyword T_REC.

The vector-field observing sequence, run on the HMI side camera, captures six polarizations at each wavelength according to a repeating 135-second framelist.

HMI_IQUV_averaging shares many steps with HMI_observables: *e.g.*, the gap filling, the de-rotation, the re-centering of the images, and the polarization calibration. Rather than performing a temporal interpolation at T_REC though, it performs a temporal averaging. Conceptually, this averaging is executed in two steps. First a temporal Wiener interpolation of the observed filtergrams onto a regular temporal grid with a cadence of 45 s is performed; this results in a set of 25 frames for each wavelength/polarization state constructed using the ten original 135 s framelists. The full time window over which the interpolation is performed is 1350 s, which is wider than the averaging window; a wider window is

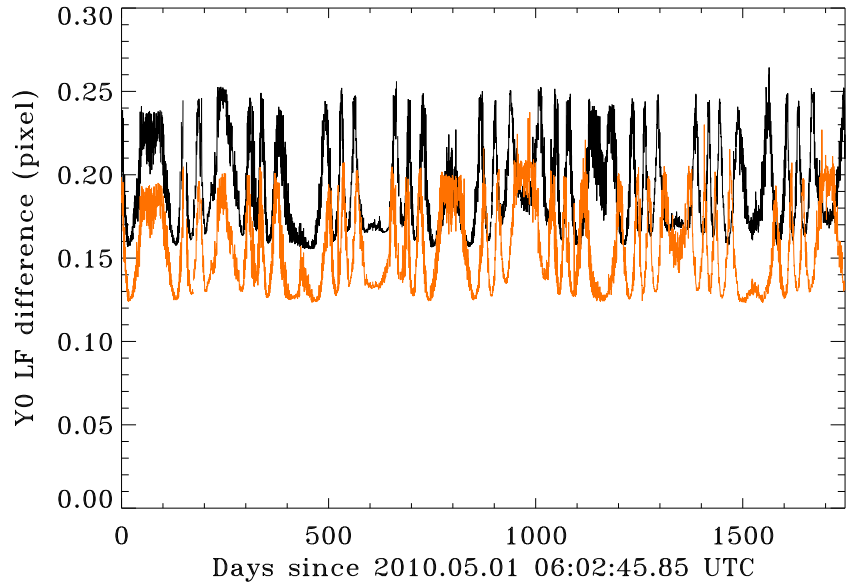

Figure 6. Difference in the Sun center vertical position (Y0_LF) between the value analytically corrected from the level 1 records and the actual value measured by the limbfinder after the level 1 record is corrected for instrumental distortion.

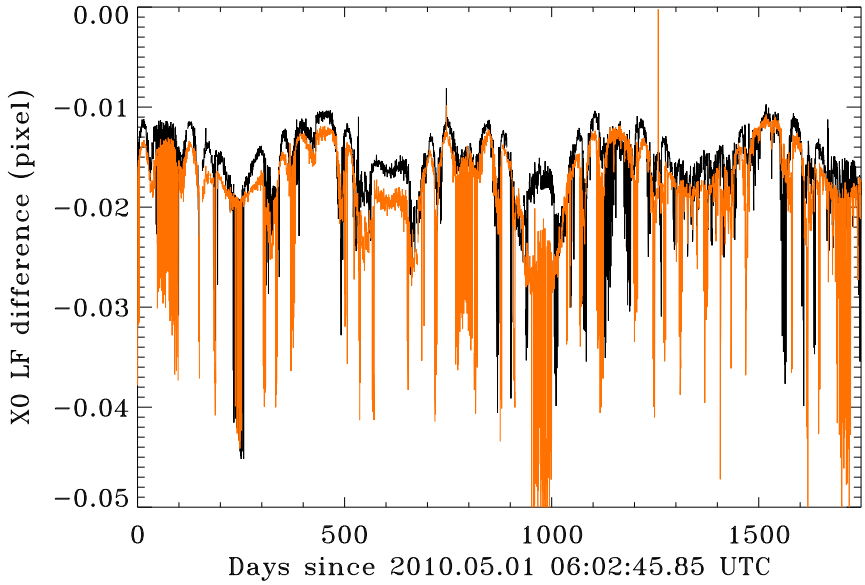

Figure 7. Difference in the Sun center horizontal position (X0_LF) between the value analytically corrected from the level 1 records and the actual value measured by the limbfinder after the level 1 record is corrected for instrumental distortion.

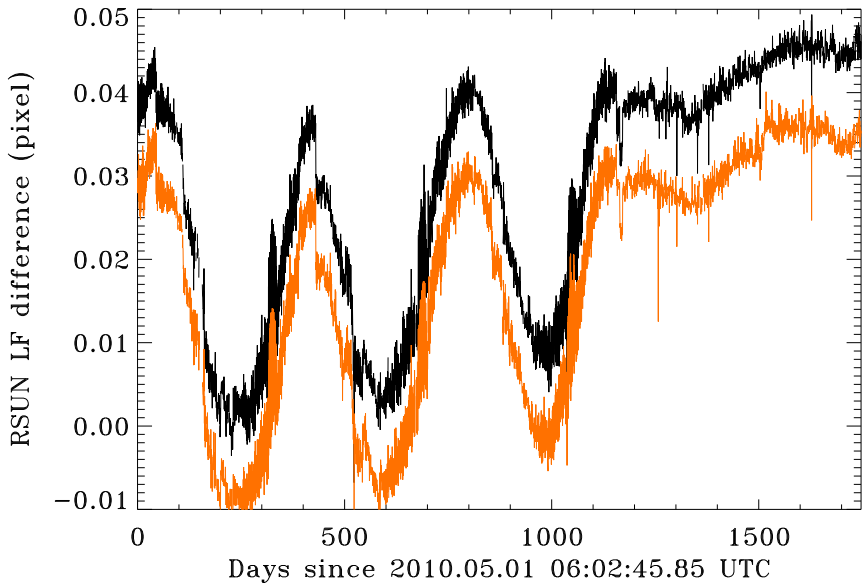

Figure 8. Difference in the Sun radius (RSUN_LF) between the value analytically corrected from the level 1 records and the actual value measured by the limbfinder after the level 1 record is corrected for instrumental distortion.

required as the interpolation needs filtergrams before and after the interpolated times. That is followed by the averaging of the frames using an apodized window with a FWHM of 720 s; the window is a boxcar with $\cos^2$ apodized edges that nominally has 23 nonzero weights, of which the central nine have weight 1.0. Temporal gap filling is also performed if needed.

Following the temporal averaging, HMI_IQUV_averaging calls the polarization calibration routine polcal() that converts the six polarizations taken by the observables sequence into a Stokes [I,Q,U,V] vector. To perform this conversion the polarimetric model described in Schou et al. (2012b) is used, including the corrections that depend on the front window temperature and the polarization selector temperature.

At each point in the image a least-squares fit is then performed to derive I, Q, U, and V from the six observed polarization states. Two additional corrections based on post-launch analysis are applied to the model described in Schou et al. (2012a). The first compensates for what looks like telescope polarization, a spatially dependent term proportional to I that appears in the demodulated Q and U at the level of about a part in 104. The dependence of Q/I, U/I, and V/I on distance from disk center was determined using the good-quality images from 3 May to 3 September 2010. The effect on V is negligible, so no correction is performed on V. The coefficients of proportionality for Q and U are given as fourth-order polynomials in the square of the distance from the center of the image, and this allows for a correction to better than a few parts in 105. It may be noted that this is not strictly a telescope polarization term, because it depends

on the polarization-selector setting. While the need for the first correction was anticipated, the second was not. A perceptible (relative to the photon noise) granulation-like pattern appears in Q and U (again, V is largely unaffected). This signal appears to be caused by a PSF that differs with polarization state. The consequence is that a contamination from I convolved with a different PSF is added to the two linear polarization signals. This is corrected by convolving I with a five by five kernel and subtracting the result. At present a spatially independent kernel is used. The final result of the Stokes-vector observables code are the four Stokes parameters at the six wavelengths.

2.10. Line profile tweaking and calibration changes

To produce the look-up tables for the MDI-like algorithm as described in Section 2.2, we need a model of the Fe I line profile (at rest and in the quiet Sun). Figure 9 shows, in black, two observed Fe I line profiles: one from the Kitt-Peak atlas, and one provided by Roger Ulrich. Both line profiles are rather different in terms of depth, width, and asymmetry. In general, it is difficult to know which profile to use as different instruments return different estimates, because of, amongst others, different PSF and wavelength resolutions. Moreover, the detune sequences regularly taken by the HMI and used to estimate the filter and Fe I profiles are only sensitive to a combination of different quantities. The fitting code used to estimate the best parameters for the filters and the solar lines also has to deal with partly degenerate quantities (for instance, the filter contrast and the Fe I linedepth). Finally, some of the quantities characterizing the HMI filters can only be measured from the ground (because they require access to a large wavelength range), and are not know with a high precision. For instance, the filter-element free spectral ranges (FSR) have been measured several times prior to the SDO launch and yielded conflicting results. For all of these reasons, it is difficult to precisely measure the Fe I line profile using detune sequences. Following the wavelength drift in the Michelson interferometers (and other long-term changes in HMI), we have had to re-tune the instrument several times and to change our estimate of the Fe I line profile to produce better look-up tables. Three different calibrations have been used since the SDO launch: the solar line profile used to produce the look-up tables has been slightly modified. Figure 9 shows, in red, the three line profiles we have been using so far.

In the code producing the look-up tables for the MDI-like algorithm, the Fe I line is approximated by a Voigt profile and two Gaussians (to simulate the line asymmetry):

$$\begin{aligned} I &= Ig - dg\exp(-l^2)\left(1 - \frac{a}{\sqrt{\pi}\,l^2}\left[(4l^2+3)(l^2+1)\exp(-l^2) - \frac{2l^2+3}{l^2}\sinh(l^2)\right]\right) \\ &- A\exp(-(\lambda+B)^2/C^2) \\ &+ D\exp(-(\lambda-E)^2/F^2) \end{aligned} \tag{13}$$

where $l = \lambda/wg$ and for $|l| \leq 26.5$. With the initial calibration: $Ig = 1.0$, $dg = 0.5625$, $wg = 0.06415$, $A = 0.015$, $B = 0.225$, $C = 0.2$, $D = 0.004$, $E = 0.150$, $F = 0.22$, and $a = 0.03$. With the second calibration used: $Ig = 1.0$,

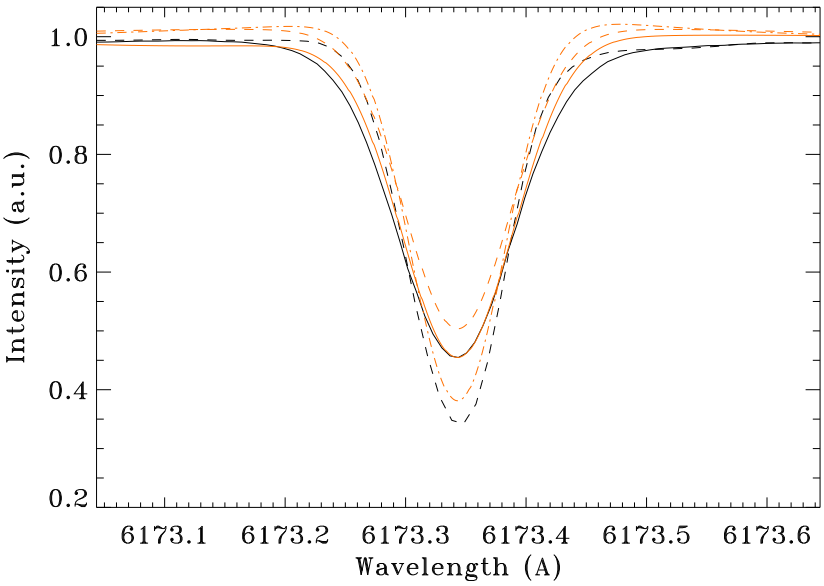

Figure 9. Different Fe I line profiles used for the calibration of HMI (in red) and two observed line profiles (in black).

$dg = 0.53$, $wg = 0.0615$, $A = -0.01$, $B = 0.225$, $C = 0.2$, $D = 0.015$, $E = 0.1$, $F = 0.25$, and $a = 0.03$. Finally, with the third calibration: $Ig = 1.0$, $dg = 0.58$, $wg = 0.058$, $A = -0.0074$, $B = 0.2$, $C = 0.13$, $D = 0.021$, $E = 0.05$, $F = 0.18$, and $a = -0.09$.

The first calibration was used from May 2010 until January 18, 2012 (18:15 UT), the second calibration from January 18, 2012 (18:15 UT) to January 15, 2014 (19:18 UT), and the third calibration ever since. In each case, we used the following FSRs for the filter elements: 168.9 mÅ for the NB Michelson, 336.85 mÅ for the WB Michelson, 695 mÅ for the Lyot element E1, 1417 mÅ for E2, 2779 mÅ for E3, 5682 mÅ for E4, and 11354 mÅ for E5.

2.11. Phase maps

The tunable elements of HMI, the two Michelsons and the Lyot element E1, have a transmittance $T(\lambda)$ that is modeled as:

$$T(\lambda) = \frac{1 + B\cos(2\pi\lambda/\mathrm{FSR} + \Phi + 4\phi)}{2} \tag{14}$$

where ϕ is the tuning phase (the phase resulting from the tuning motor positions), while Φ is the relative phase. Detune sequences are taken every other week to measure maps (over the CCDs) of Φ of all three tunable elements. It was noticed early on that the average phases of the tunable elements differ slightly between front and side cameras, as can be seen on Figure 10. The maximum

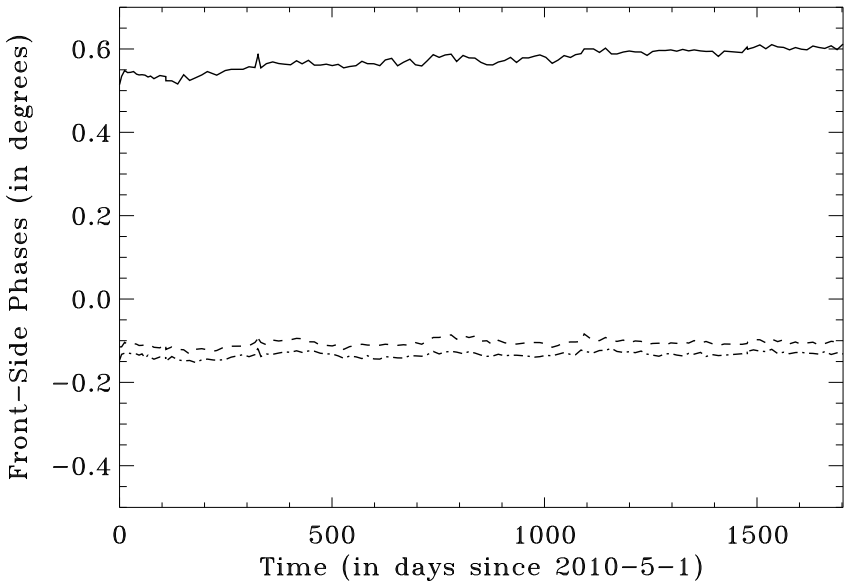

Figure 10. Spatially averaged phase difference for the three tunable elements, between the front and side cameras.

difference is for the narrow-band Michelson, and it exhibits a slow increase with time (by about 4.3×10^{-5} degrees per day). The broad-band and Lyot E1 elements exhibit a smaller phase difference, which seems to be stable with time. The origin of these front/side camera phase differences may reside in the polarizer of the NB Michelson: it may leak the orthogonal polarization that is then picked up by the polarizing beam-splitter used to separate the front and side camera paths.

2.12. Height of formation correction

The limb finder applied to level 1 images in the HMI processing pipeline returns a solar radius that depends on the wavelength at which the image was taken (and characterized by its FID keyword). The standard observables sequence samples the solar Fe I line at 6 wavelengths, each separated by about 68.8 mÅ. Because the heights of formation of various parts of the Fe I line are different, it is normal that the radius returned by the limb finder also exhibits a dependence on the FID.

The values returned by the limb finder (X0_LF, Y0_LF, and RSUN_LF) are corrected for the wavelength dependence in the keywords CRPIX1, CRPIX2, and R_SUN.

The correction routine first determines the wavelength index wl of the filtergram, by dividing its FID by 10. It then computes corr1 and corr2:

$$\mathrm{corr1} = 0.445 \exp(-(wl - 10 - \mathrm{OBS\_VR}/(F\,3 \times 10^8/20) - 0.25)^2/7.1) \quad (15)$$

and:

$$\mathrm{corr2} = 0.39\,(-2(wl - 10 - \mathrm{OBS\_VR}/(F\,3\times10^8/20) - 0.35)/6.15) \quad (16)$$

$$\exp(-(wl - 10 - \mathrm{OBS\_VR}/(F\,3\times10^8/20) - 0.35)^2/6.15) \quad (17)$$

where $F = 0.690/6173$. Then the corrections applied are:

$$\mathrm{CDELT1} = \mathrm{CDELT1}\,\mathrm{RSUN}/(\mathrm{RSUN} - \mathrm{corr1}) \quad (18)$$

$$\mathrm{RSUN} = \mathrm{RSUN} - \mathrm{corr1} \quad (19)$$

$$\mathrm{CRPIX1} = \mathrm{CRPIX1} - \cos(\pi - \mathrm{CROTA2}\,\pi/180)\,\mathrm{corr2} \quad (20)$$

$$\mathrm{CRPIX2} = \mathrm{CRPIX2} - \sin(\pi - \mathrm{CROTA2}\,\pi/180)\,\mathrm{corr2} \quad (21)$$

The stability of this formation height correction has been monitored since the SDO launch, as it is possible for the details of the correction to change after a HMI retuning or/and when we change the calibration used to determine the observables.

Figure 11 shows an example of the impact of the height of formation on the solar radius returned by the limb finder. For a full day on May 17, 2010, we plotted the solar radius (corrected by the Sun-SDO distance) for all of the front camera images. Each image is given an effective wavelength, which is its filter index (given by the FID) corrected by the Doppler shift in the Fe I line profile resulting from OBS_VR. The resulting plot is well fitted by a Gaussian profile:

$$R = A\exp(-(wl - wl_0)^2/wlw^2) + \mathrm{offset} \quad (22)$$

Figure 12 shows how the Gaussian parameters A, wl_0, wlw, and the offset vary with time.

There is clearly a time dependence of these parameters, but in the observables pipeline we adopted the time-constant correction described at the beginning of this section.

2.13. Observables look-up tables

The MDI-like algorithm computes a discrete estimate of the first and second Fourier coefficients of the Fe I line profile, using the six wavelengths of an observables sequence. The phase of the first Fourier coefficients is directly proportional to the Doppler velocity. However, this is true only for a Gaussian line profile, for filter profiles that are delta functions, and for an accurate enough discrete estimate of the Fourier coefficients. None of that is true with HMI data: the Fe I profile is asymmetric, the filter transmission profiles are wide and have significant side-lobes, and the discrete estimate is based on only six points. Therefore, it is necessary to correct the Doppler velocities returned by the MDI-like algorithm. To do so, we compute look-up tables: at each point on the HMI CCDs, we compute the correct velocity as a function of the velocity obtained from the Fourier phase. The correct velocity is determined by using a better model of

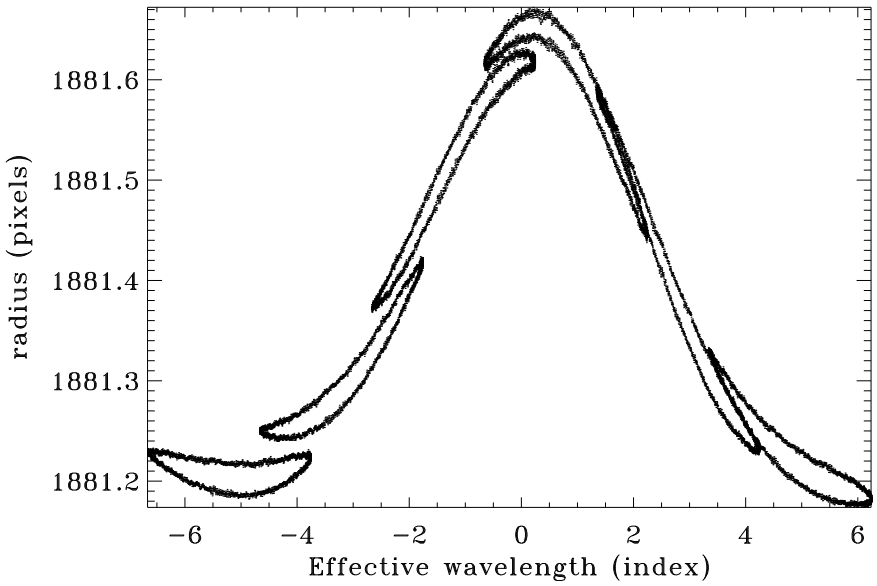

Figure 11. Solar radius measured on the front camera in pixels, as a function of the index of the effective wavelength. An effective wavelength of 1 means that when the filtergram was taken, the Fe I line center was 68.8 mÅ away from its position at rest. Data taken on May 17, 2010.

the Fe I line profile than a Gaussian one (see Section on tweaking the line profiles), and calibrated filter transmission profiles. These look-up tables are imperfect, due to error in the filter profiles and in the model of the Fe I profile. Figure 13 shows an example of look-up table for the pixel at CCD center for the front camera, and for three different calibrations: each time we retune HMI (to compensate for the wavelength drift in the Michelson interferometers), we need to produce new look-up tables.

2.14. Phase maps

The phase maps of the tunable elements exhibit large-scale and small-scale interference fringes, most likely created by the front window of HMI. Initially, we used phase maps uncorrected for these fringes to produce the look-up tables for the MDI-like algorithm. This resulted in a fringe pattern visible in the HMI observables. The large-scale fringe pattern has been successfully corrected based on an algorithm fitting the phase and amplitude of the fringes. All of the observables with a T_REC after 2012.10.1 were produced with look-up tables based on corrected phase maps. The small scale fringe pattern proved more difficult to remove and is still partly present in the observables

The following DRMS series are used to store the tunable-element phase maps: hmi.phasemaps, hmi.phasemaps_corrected (with fringe correction), hmi.phasemaps_cal11 (phase maps computed with the original calibration of HMI, and from May 1,

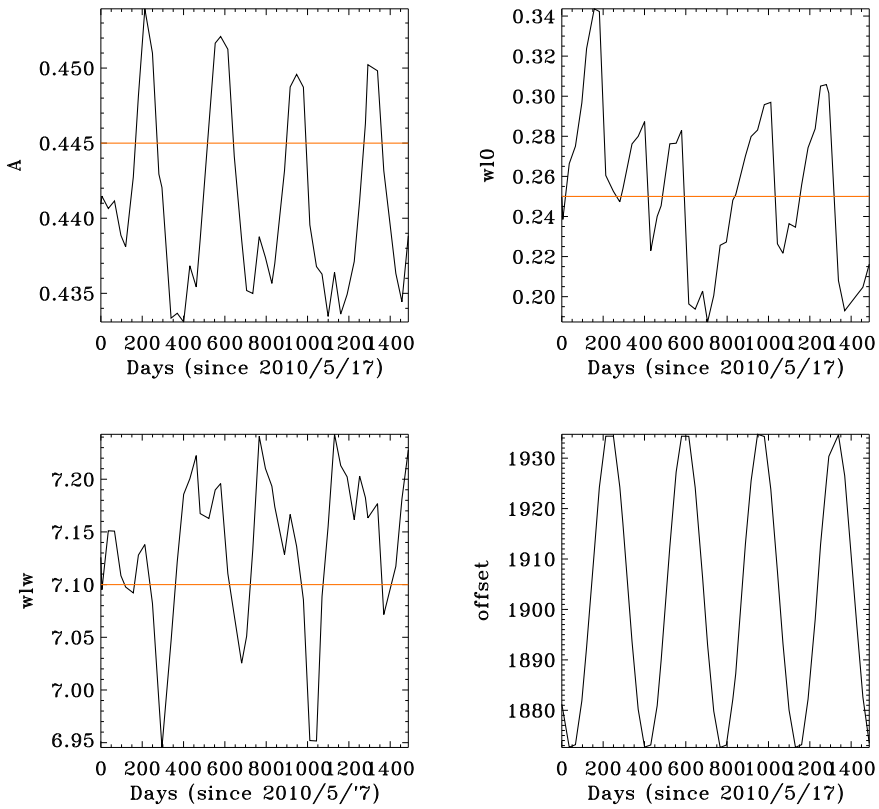

Figure 12. Changes in the parameters of the formation height correction since May 2010.

2010 to December 5, 2010l; used only to compute the interference fringe correction), hmi.phasemaps_cal12 (phase maps computed with the second calibration of HMI, and from May 1, 2010 to December 21, 2011; used only to compute the interference fringe correction), and hmi.phasemaps_cal13 (phase maps computed with the third calibration of HMI, and from May 1, 2010 to December 18, 2013; used only to compute the interference fringe correction).

2.15. Polarization calibration

2.16. Polynomial correction

We correct the Doppler velocity computed by the MDI-like algorithm using a polynomial regression: the median Doppler velocity RAWMEDN returned by the MDI-like algorithm minus OBS_VR is fitted as a function of RAWMEDN over a 24 hour period, by a third-order polynomial. The Doppler velocity returned by the MDI-like algorithm is then corrected at each pixel on the CCD using this polynomial. Figure 14 shows the time variation of the four polynomial coefficients since May 1, 2010. COEFF0 represents the offset between RAWMEDN and

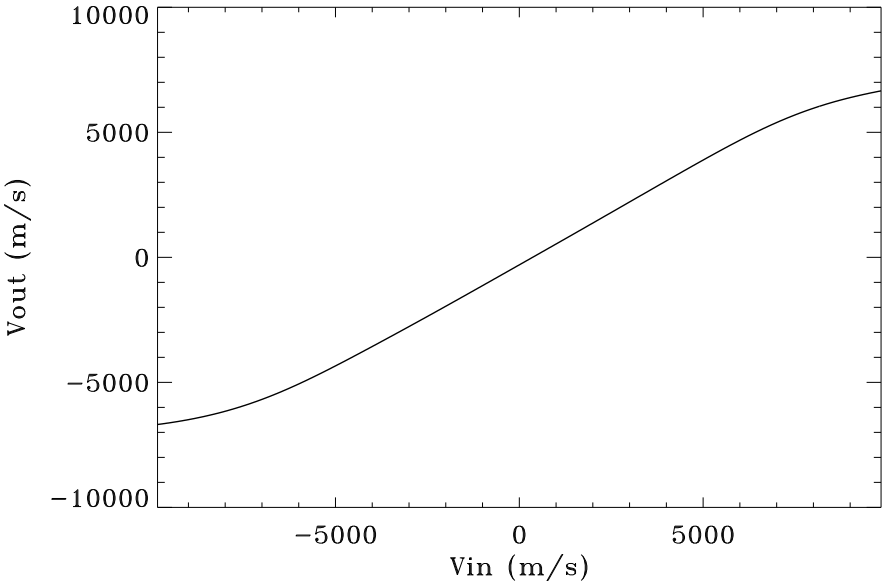

Figure 13. Example of a look-up table for the first Fourier coefficients, and for the central CCD pixel on the front camera.

OBS_VR, and it increases with time because of the wavelength drift in the two Michelson interferometers. The regular jumps in COEFF0 are due to the retuning of the instrument.

The DRMS series where the polynomial coefficients are recorded is name hmi.coefficients. The correction_velocities.c code is run automatically once a day and computes coefficients for T_REC separated by 12 hours. It is run on hmi.V_45s_nrt records, and expects 1920 Dopplergrams per day. For days when calibration sequences were run and the number of Dopplergram is significantly lower than 1920, it is possible to force the code to still produce polynomial coefficients.

The polynomial correction suffers from two main shortcomings: first, it is computed on the velocity range covered by OBS_VR, which is quite limited compared to the velocity range of the solar signal: for velocities outside the OBS_VR range, the polynomial correction amounts to an extrapolation; second, the correction is computed for a spatial average over then CCD, and therefore local differences are not taken into account: in particular the correction seems to be less accurate toward the solar limb.

3. Error Estimates and Impact on Observables

In the following sections we list the problems in the HMI observables that we are aware of.

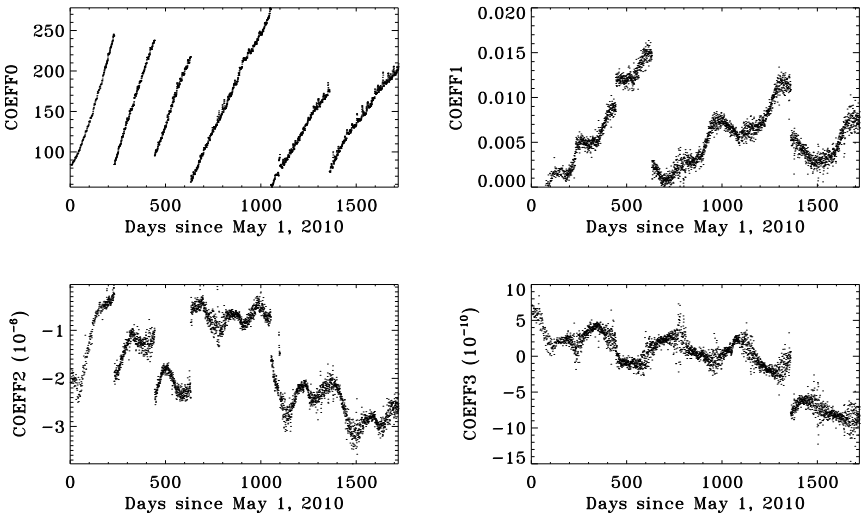

Figure 14. Polynomial coefficients used to correct the Doppler velocities, as a function of time.

3.1. 24h oscillations in observables

Temporal variations with a 24-hour period are found in each and every HMI observables, especially inverted vector magnetograms and line-of-sight magnetograms (*e.g.* Liu *et al.*, 2012; Hoeksema *et al.*, 2014).

The major contribution to this variation is the Doppler shift of the spectral line. The geosynchronous orbit of SDO has a daily change of velocity relative to the Sun that ranges from -3.5 to $+3.5$ km s^{-1} and varies during the year. Solar rotation, together with other additional motions, complicates this issue. The SDO orbit also indirectly impacts the observables. For instance, the relative position of the HMI side camera with respect to the Earth during the course of a day produces a change in the amount of heat received by this camera. This change in thermal environment impacts the level 1 filtergrams, and consequently the observables, in subtle ways.

As an example of the 24-hour oscillations, a vector magnetogram for the simple and stable active region AR 11084 is studied. This active region was tracked for four days, from the first to the fourth of July 2010. For each magnetogram, we selected two groups of pixels that have a relative intensity in the following ranges: 0.0 to 0.35, and 0.65 to 0.75. We also selected a quiet-Sun area of 40×30 pixels. The three groups include pixels in the sunspot umbra, penumbra, and in the quiet Sun region. The temporal profiles of mean magnetic-field strength of the selected pixels are shown in the left columns of Figure 15 for the umbra (top), the penumbra (second), and the quiet-Sun region (third); for reference the mean LOS velocity in the quiet-Sun region appears at the bottom of both columns. The red curves are third-order polynomial fits. Differences between the mean field strengths and their polynomial fits, the residuals, are plotted in the right columns. Variation of field strength in the sunspot is less than 5%.

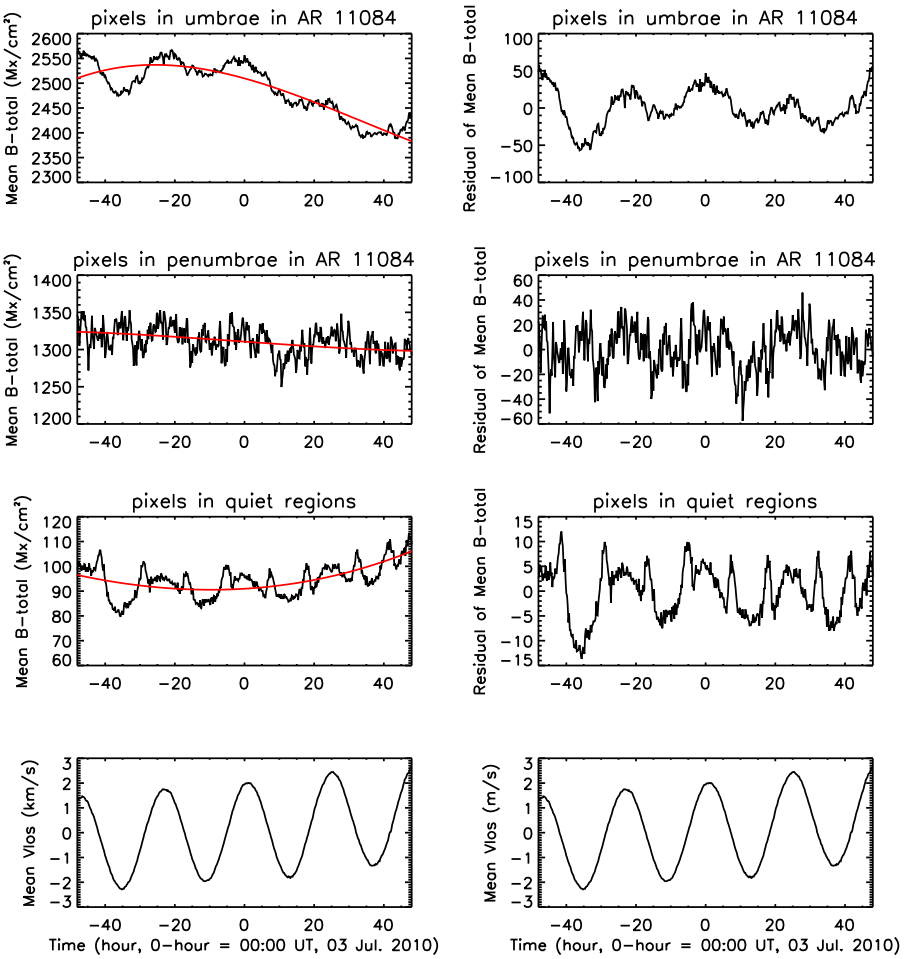

Figure 15. The black curves in the left column in the top three panels are four-day temporal profiles for AR 11084 of the mean magnetic-field strength determined for the umbra, penumbra, and quiet Sun, respectively. The red curves are third-order polynomial fits. Residuals are plotted in the top three panels on the right. The residual is the difference between the mean field strength and the polynomial fit. For reference, the mean line-of-sight velocity observed in the quiet-Sun region is plotted in the two bottom panels. For this analysis the umbra includes pixels where, compared with the quiet Sun, the continuum intensity, $Ic < 0.35$ when corrected for limb darkening. In the penumbra $0.65 < Ic < 0.75$. The areas selected at 06:00 UT on 2 July 2010 are enclosed by the contours shown in Fig.**??**. (From **?**.)

We then decompose field strength into LOS (ME-B$_{los}$ hereafter) and transverse fields (ME-B$_{tran}$ hereafter) to explore how the velocity correlates with these two components separately. Shown in left and middle columns of Figure 3.1 are residuals of LOS field (left column) and transverse field (middle column) of inverted magnetic field for umbra (top), penumbra (middle), and quiet Sun region (bottom). As a comparison, we plot in the left column the residuals of the LOS field, derived by Doppler velocities measured from the LCP and RCP

polarizations (B_{los} hereafter). The solid straight lines are linear fits, while the curve in the top-right panel is a 2nd-order polynomial fit. The velocity-dependent variation is clearly seen in the LOS field of both ME-B_{los} and B_{los} in the sunspot; whereas the transverse field ME-B_{tran} in the sunspot appears to not contain this variation. Strong fields may shift either LCP or RCP polarizations away from one of the HMI sampling positions but do not shift the linear polarizations (LP) away. LCP and RCP make up the Stokes V parameter used to measure the LOS field, while LP makes up the Stokes Q and U used for the transverse field. This may be one of the reasons behind the velocity-dependent variation in strong field that is significant only in the component of LOS field. The quadratic function in the residuals of B_{los} indicates that, in this determination of the LOS magnetic field, the variation is not sensitive to the direction of the velocity.

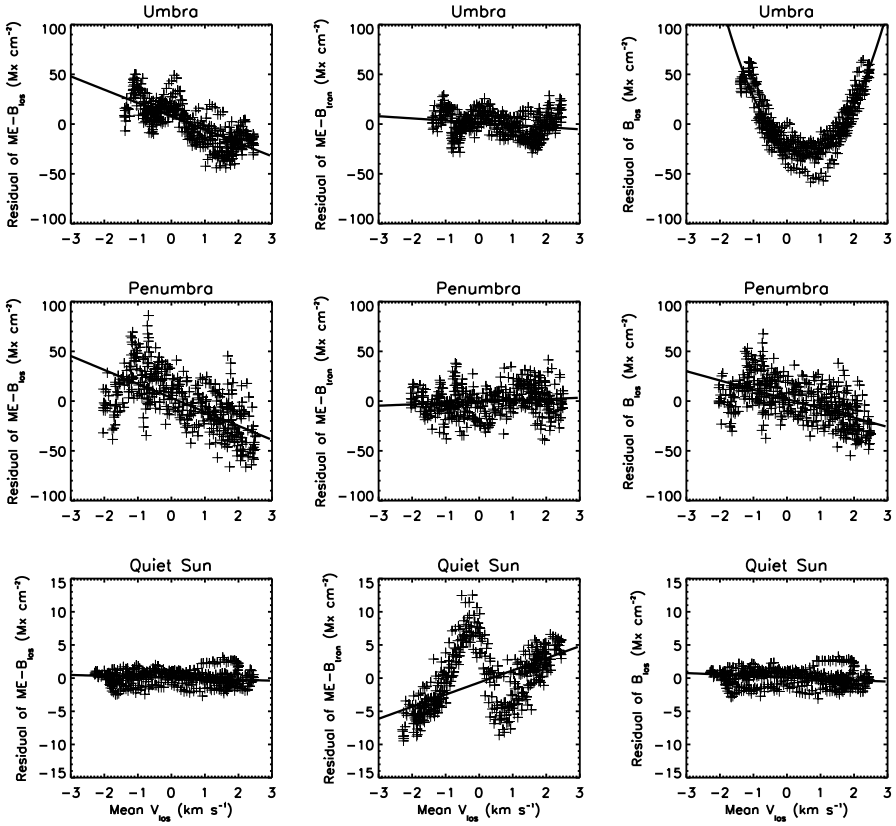

Figure 16. Scatter plots show relationship between residuals of mean field and velocity. Residual is the difference between mean field and polynomials fit as shown in Fig.15. From left to right are for line-of-sight field, transverse field of inverted vector field, and line-of-sight field derived from Doppler velocities of LCP and RCP. Solid straight lines refer to linear fits; the curve in top right panel refers to a 2-order polynomials fit. From top to bottom are for umbra, penumbra in AR 11084, and quiet Sun regions, respectively.

To further study how the 24-hour oscillations vary with field strength and depend on the LOS algorithm used, we use spectropolarimetric observations

from the IBIS instrument at the Dunn Solar Telescope (Sunspot, New Mexico). The sunspot NOAA 10960 was observed for 7 hours on 8 June 2007, and the full Stokes profiles of the Fe I line were measured at 23 wavelengths. These Stokes profiles are averaged over the first observation hour and interpolated on a fine wavelength grid. A Stokes I profile from a quiet Sun region is used as the reference line profile when computing look-up tables for the MDI-like algorithm applied to IBIS data. The LCP and RCP profiles in a pixel inside the sunspot umbra are taken and shifted in wavelength to simulate various Doppler velocities and various magnetic field strengths. The Doppler velocities tested are the OBS_VR values of December 27, 2011 (ranging from -2016 m/s to +1860 m/s). At each velocity, and for a given field strength, the MDI-like algorithm is applied to simulated HMI intensities (obtained with a sample of HMI filter transmission profiles).

We tested how the peak-to-peak variation in field strength returned by the MDI-like algorithm varies with field strength (always for the same input velocity range). Figure 18 shows that the oscillation amplitude strongly depends on strength. The specific pattern depends on the OBS_VR range: for a given field strength at a given location on the solar surface, the daily peak-to-peak variation will vary during the year with the orbital velocity of SDO.

Conversely, we selected all of the pixels with a given field inclination and calculated the peak-to-peak difference in the field strength returned by the MDI-like algorithm for various field strengths. The result, displayed on Figure 17, shows that there is a strong dependence of the peak-to-peak difference on the field strength. This dependence is probably not linear.

We tested whether computing the LOS observables with a least-squares regression on a Voigt profile reduces the sunspot oscillation, compared to the MDI-like algorithm. Indeed, other studies have shown that a least-squares fit gives more reliable results than the MDI-like algorithm for the LOS observables. Figure 18 shows (in red) the peak-to-peak difference in the field strength returned by the least-squares algorithm for the OBS_VR range of December 27, 2011. Overall, a least-squares fit does not reduce the peak-to-peak difference (even though the average value of the field strength retrieved by the algorithm is closer to the actual value than what the MDI-like algorithm returns). The main advantage of least-squares fits is that, unlike the MDI-like algorithm, they adapt to a change in the Fe I linewidth and linedepth in the presence of a strong field. However, this does not seem to be especially beneficial here. The main issue with the Fe I line profile in strong fields is not that the linewidth or linedepth changes, but it seems to be that the shape of the line is very different from a Voigt or Gaussian profile. Stronger fields alter this shape more than weaker fields, as can be seen on Figure 19.

Several attempts are reducing the amplitude of the 24-hour oscillations have been performed, but to no avail. Current implementation of the MDI-like algorithm ignores the I-ripple of the tunable filter elements, as mentioned in section 2.6. Since we are able to estimate the phase and amplitude of these I-ripples, we tested better models of the filter transmission profiles with the I-ripple taken into account. Similarly, we tested more realistic Fe I line profiles (observed ones, rather than Gaussian or Voigt models). Both cases (better models of the

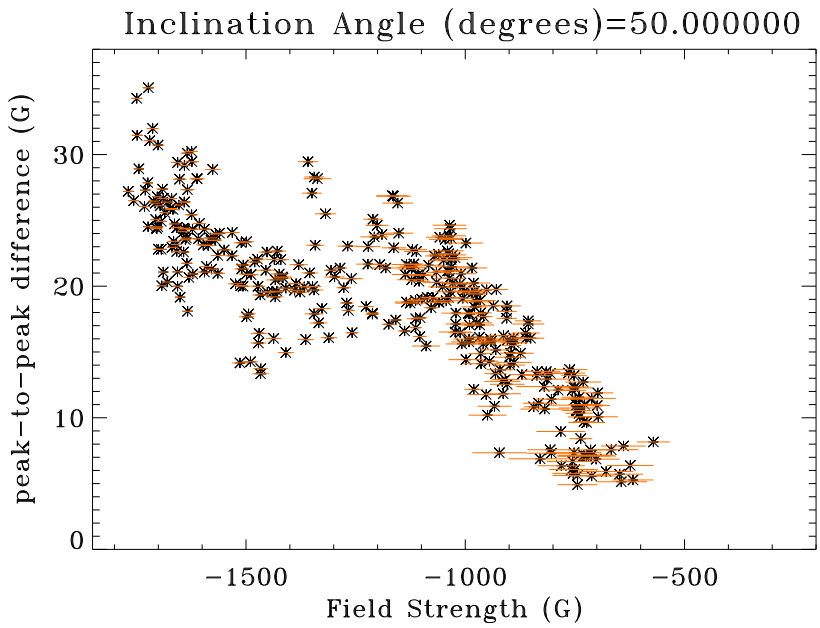

Figure 17. Peak-to-peak difference in the LOS magnetic field strength returned by the MDI–like algorithm as a function of field strength, for the Doppler velocity range of December 27, 2011 (OBS_VR in the range -2016 m/s to +1860 m/s), and for an inclination angle of 50 degrees.

transmission profiles and of the Fe I line) did not result in a lower amplitude for the 24-hour oscillations. We also tried a better polynomial correction. Currently, as mentioned in section 2.16, this correction is applied by computing the median Doppler velocity across the solar disk, and is limited to the daily OBS_VR range. To increase this range, we tested a polynomial correction where average Doppler velocities are computed on small patches located at the center and limb of the Sun. Unfortunately, the strong spatial dependence of the correction makes it difficult to properly reconcile different parts of the OBS_VR range and different locations on the solar disk. Finally, several attempts at improving our knowledge of some filter quantities (like the FSR — full spectral range — of the tunable elements) were performed by revisiting ground calibration results in the light of calibration sequences taken during the commissioning phase. The FSR of the tunable Lyot element is a basic quantity that affects the filter transmission profiles and was measured from the ground at different occasions. Unfortunately, each measurement produced a different outcome and to this day there is still a small uncertainty on the exact value of this FSR.

3.2. Roll calibration

Roll maneuvers are performed in April and October of each year when the SDO spacecraft is rotated 360 degrees around the Sun-spacecraft axis. These rolls are used for measuring the oblateness of the Sun (*e.g.*, Kuhn *et al.*, 2012). They

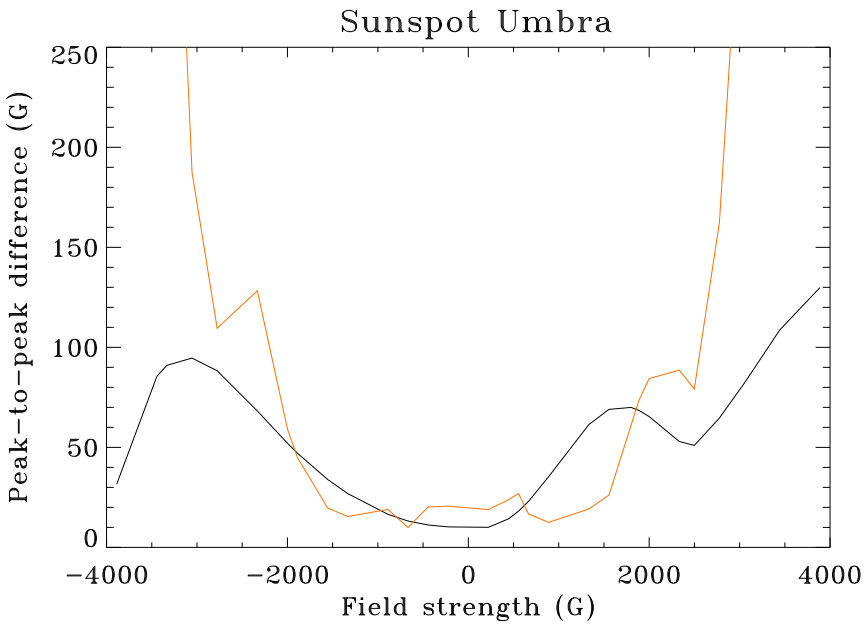

Figure 18. Peak-to-peak difference in the LOS magnetic field strength returned by the observables algorithm as a function of field strength, for the Doppler velocity range of December 27, 2011 (OBS_VR in the range -2016 m/s to +1860 m/s). The black curve is for the MDI-like algorithm, while the red curve is for a least-squares fit with a Voigt profile.

can also be used to test some systematic errors on the Doppler velocity. For instance, by plotting the median Doppler velocity across the solar disk prior to applying the polynomial correction (RAWMEDN) minus the Sun-SDO radial velocity (OBS_VR) as a function of RAWMEDN, the impact of a changing SDO roll angle (CROTA2) informs us on spatial dependence in the line-of-sight observables algorithm: were the calibration of HMI perfect, a same RAWMEDN should correspond to a same RAWMEDN-OBS_VR difference whatever the roll angle is. However, as can be seen on Figure 20, this is not the case. This provides a rough estimate of the systematic error there is on the determination of the Doppler velocity. Different CROTA2 angles mean that a specific location on the Sun is projected onto a different CCD pixel. If the line-of-sight observables algorithm returns different Doppler velocities this means that these velocities have a spatial dependence that should not exist with a perfect calibration. It appears that such a spatial dependence results in errors of less than 15 m/s on the estimated velocities. Figure 21 shows the median velocity across the solar disk (DATAMEDN) as a function of the roll angle CROTA2.

3.3. Solar rotation rate daily variation

The solar rotation rate measured on HMI Dopplergrams exhibits a 24-hour period. An example is shown on Figure 22 for January 5, 2015. That day, the peak-to-peak variation in the equatorial rotation rate was about 65 m/s.

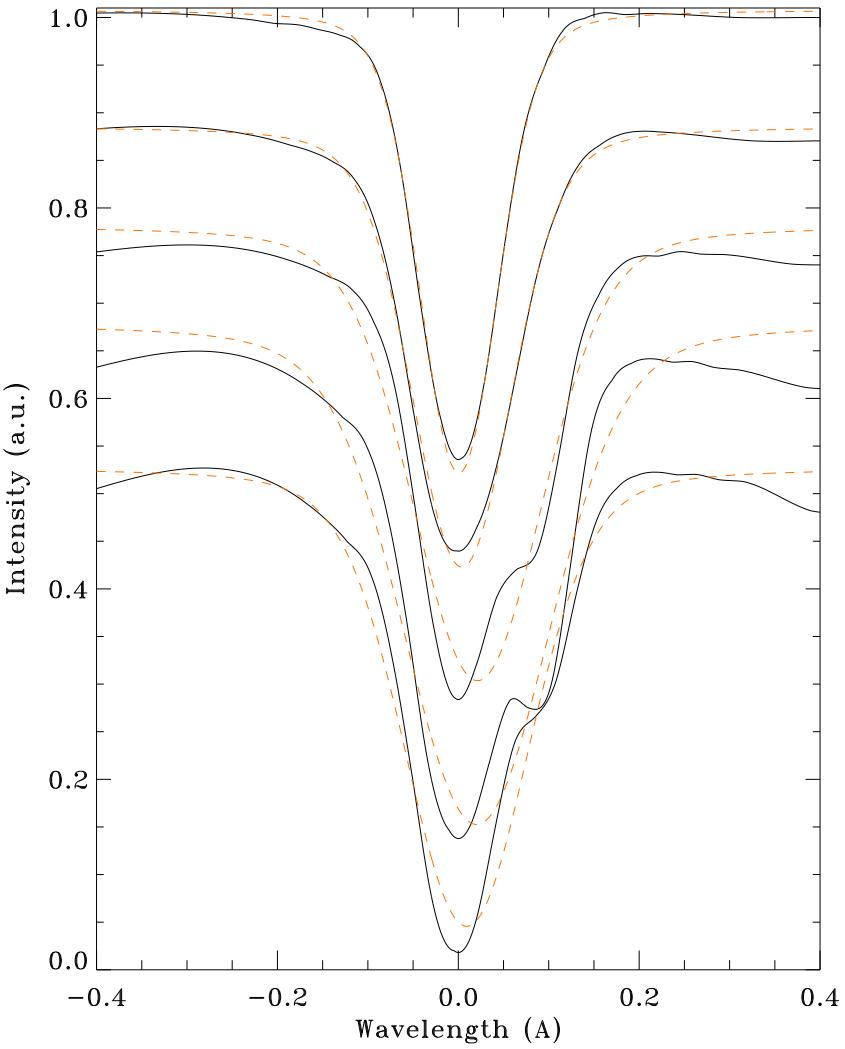

Figure 19. Changes in LCP Fe I line profile from quiet Sun (top curve) to the sunspot umbra (bottom curve). Red lines are the results of a least-squares fit with a Voigt profile.

The determination of the solar rotation rate is very sensitive to Doppler velocity near the limb: the daily variation in the measured rotation rate highlights errors in the computation of the Doppler velocity with the MDIL-like algorithm.

3.4. Errors with the LOS algorithm

As was already mentioned in Section 3.1, in the presence of strong fields the MDI-like algorithm may produce significant errors. The shape of the Fe I line in a strong and inclined field may significantly differ from a Voigt profile and

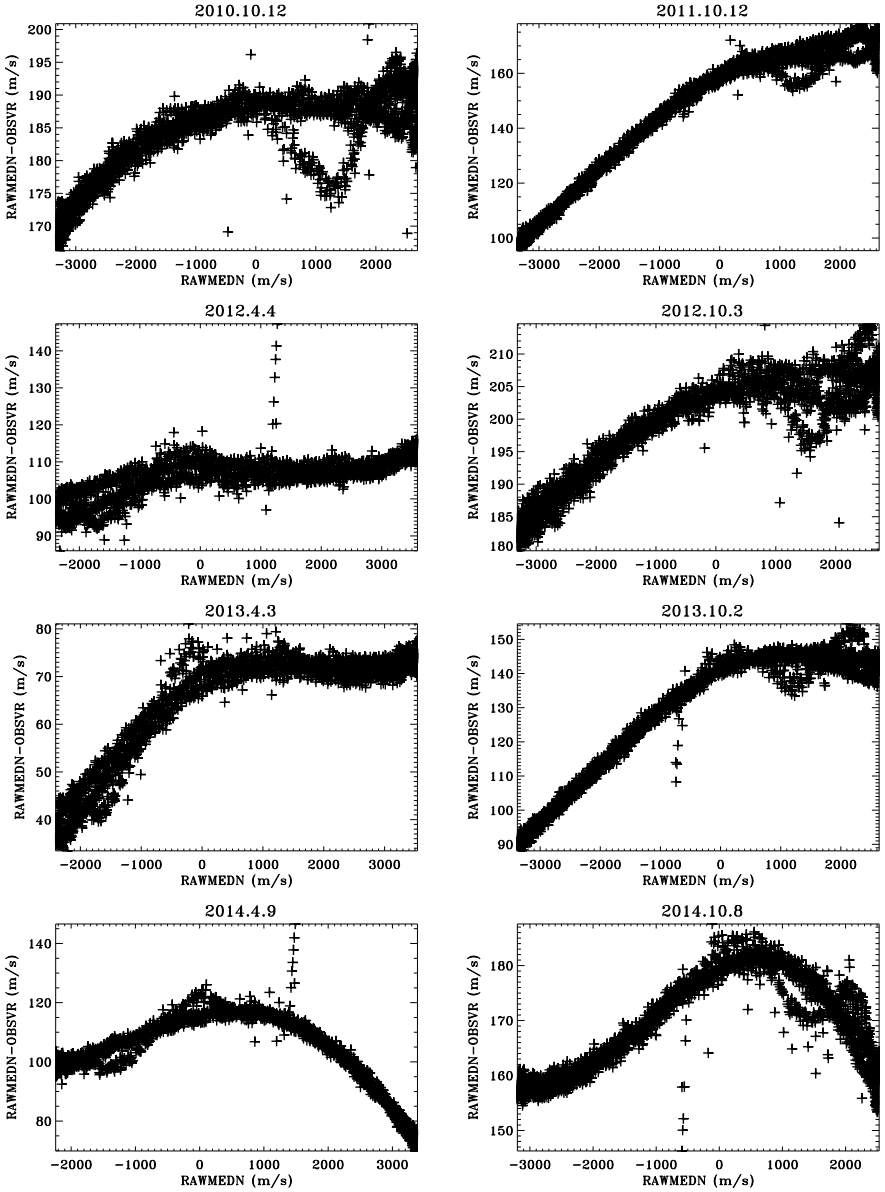

Figure 20. Spatially-averaged Doppler velocity returned by the line-of-sight algorithm, prior to applying polynomial correction, (RAWMEDN) minus the Sun-SDO radial velocity (OBS_VR) as a function of RAWMEDN. We show the results of eight SDO rolls.

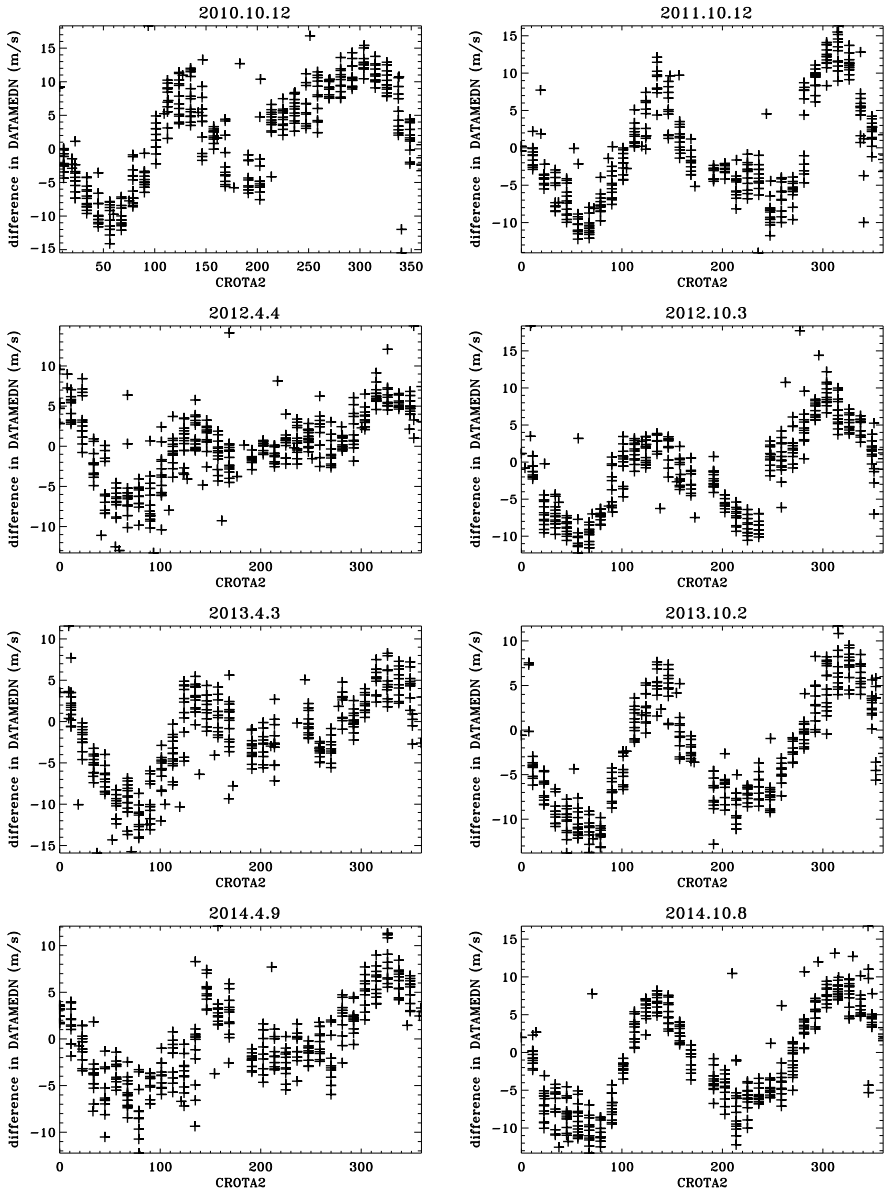

Figure 21. Median velocity across the solar disk (DATAMEDN) as a function of the roll angle CROTA2. We show the results of eight SDO rolls.

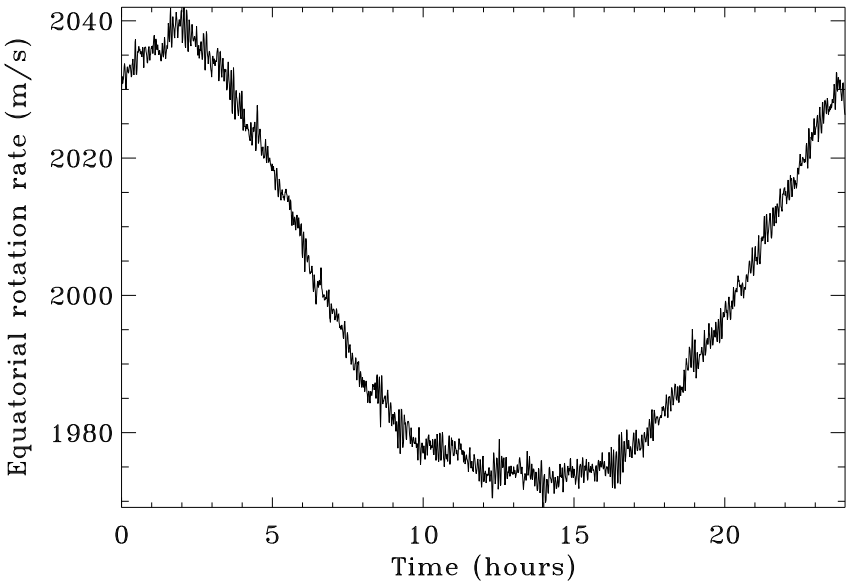

Figure 22. Daily variation in the equatorial rotation velocity, for T_REC=2015.1.5.

from the profile used to produce the look-up tables, while the wavelength shift resulting from the Zeeman effect may push the LCP, RCP, or both components partly or totally outside the dynamic range of HMI.

One way of testing such errors is to take a 10-wavelength observables sequence, rather than the usual six. On 24 October 2014, such a special observables sequence was run for about an hour: the cadence on the front camera was 75 seconds, rather than 45 seconds, and ten equally spaced wavelengths were taken. There was a large sunspot that day, NOAA 12192. Figure 23 shows a comparison of the observables quantities returned by the MDI-like algorithm when using six or ten wavelengths. The six-wavelength plots were obtained by looking at standard observables obtained just prior to the special sequence run. While the continuum intensity appears fairly robust, the linewidth is especially sensitive to a change in the number of wavelengths: with six wavelengths, the computed Fe I line width decreases as we move from quiet Sun toward the sunspot umbra. This is the reverse of what is expected in the presence of a magnetic field, and of what the 10-wavelength sequence returns. The panels showing the individual LCP and RCP velocities highlight the difference between six and ten wavelengths: the LCP velocities derived with six wavelengths show saturation inside the sunspot umbra, due to the wavelength shift produced by the Zeeman effect (the LCP profile is partly outside the dynamic range of HMI with only six wavelengths). Moreover, the RCP velocities with six wavelengths are underestimated compared to 10 wavelengths. This results from the RCP profile partly lying outside the dynamic range of only six wavelengths: the MDI-like algorithm, based on a discrete estimate of the Fourier coefficients, assumes

that the line profile is periodic. If the profile is truncated, then the condition of periodicity produces a line profile whose shape is significantly different from a quiet-Sun profile, resulting in large errors in the estimate of the velocity shift.

The impact of a magnetic field on the line profile can also be seen on Figure 24: it displays the LCP and RCP components obtained during another 10-wavelength sequence taken on September 10, 2014. This figure is for a pixel in a sunspot umbra (solid lines) and in the quiet Sun (dashed lines). The magnetic field in the sunspot shifts the LCP (black lines) and RCP (red lines) components, but they remain within the dynamic range of HMI (even with only six wavelengths). However, the line shapes are significantly distorted compared to the quiet-Sun shape used to produce the look-up tables: both LCP and RCP components are shallower and wider than in a quiet region, and the LCP profile does not have a clearly defined minimum (a fit of the profile by a Gaussian function finds a minimum at a lower wavelength than the location of the actual minimum). This results in an underestimation of the actual LOS magnetic-field strength in the umbra.

Therefore, it appears that in the presence of a strong field the MDI-like algorithm will underestimate (in absolute value) the LOS field strength, and may result in saturation if the Doppler shift resulting from motions (Sun-SDO radial velocity, solar rotation, acoustic waves, convection, etc.) combined with the Zeeman effect send the LCP, RCP, or both profiles outside the range of the instrument. Even in the absence of saturation, the fact that the Fe I line profile is different from the one used to produce the look-up tables will result in errors in the observables determination and contributes to the 24-hour oscillations.

3.5. Magnetic-field error with Stokes vector inversion

Inversion of the full Stokes vector obtained on the side camera of HMI produces better estimates of the magnetic field than the LOS strength returned by the MDI-like algorithm. For one thing, VFISV does fit for the width and depth of the Fe I line profile, thus taking into account the broadening and shallowness of this profile in presence of a magnetic field. However, VFISV assumes that the line profile follows a Voigt profile, which is not necessarily the case in presence of strong fields. For this reason, and others, the full-vector magnetic-field data also have errors. To try and quantify these errors, we use the results of the special ten-wavelength observables sequence of 24 October 2014, described in the previous section. Full Stokes profiles were produced for the T_REC in the time interval covered by this sequence, and VFISV was run to invert the vector magnetic field. The large sunspot NOAA 12192 allowed a comparison of inverted field strengths in presence of strong fields. To make comparison easier with the front-camera observables, we looked at the LOS field strength returned by VFISV (field strength multiplied by the cosine of the inclination angle). Six-wavelength results we obtained by looking at standard VFISV inversions just prior to the start of the special sequence run. As can be seen on Figure 25, the LOS field strength inside the umbra of the sunspot is overestimated (in absolute value) with six wavelength compared to 10 wavelengths. This plot can be used to provide a rough estimate of systematic errors we might have on the inverted field strength from VFISV with the standard observables sequence.

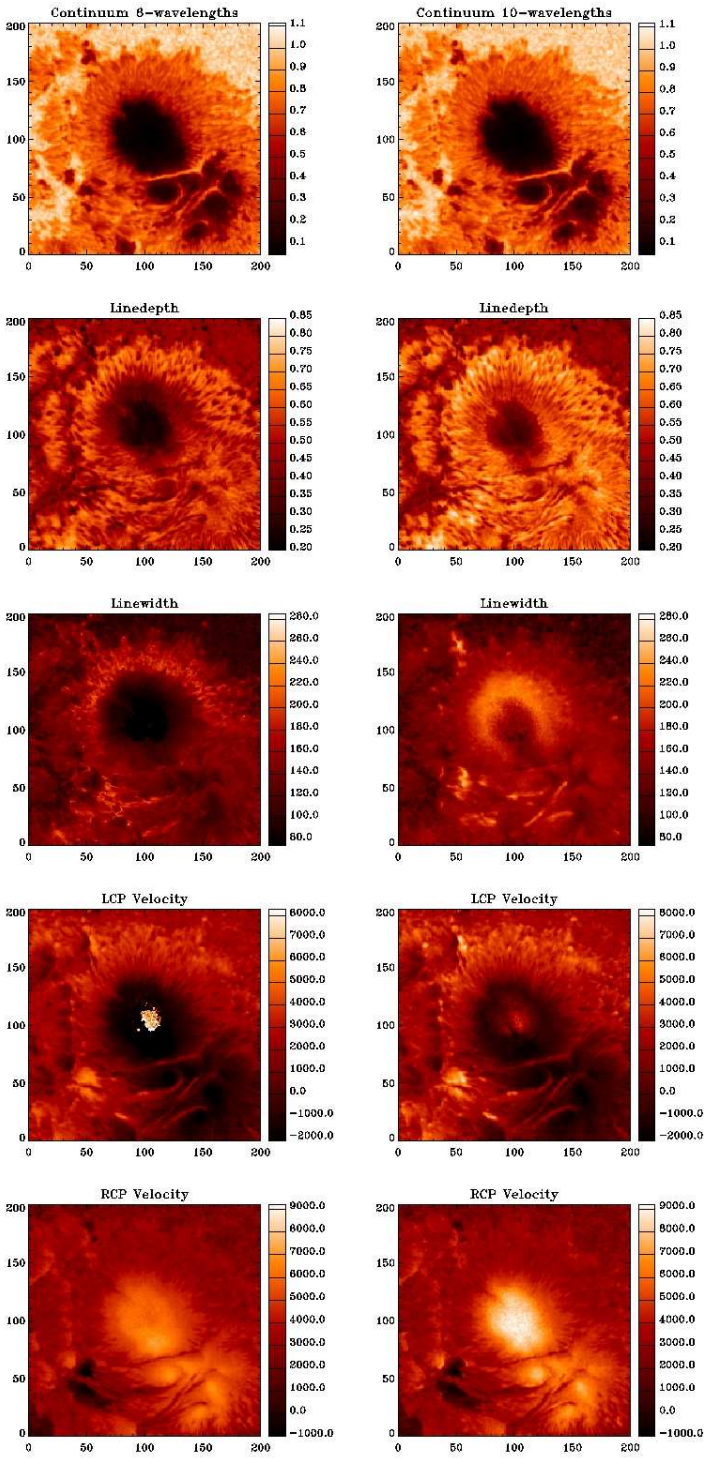

Figure 23. Observables quantities returned by the MDI-like algorithm based on six and ten wavelengths. Results derived from a ten-wavelength sequence taken on 24 October 2014. We show the active region NOAA 12192.

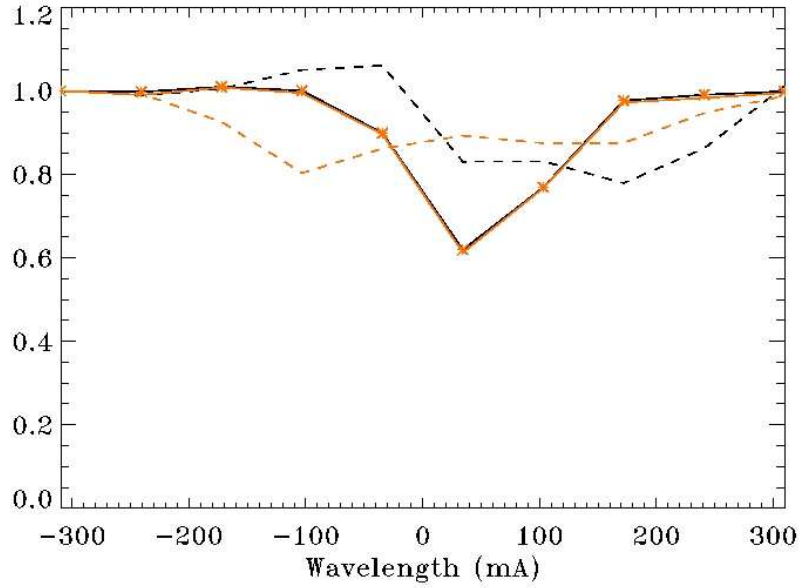

Figure 24. LCP (black lines) and RCP (red lines) components of the Fe I line profile in the pixel of a sunspot umbra (solid lines) and in a quiet-Sun pixel (dashed lines). Results from a ten-wavelength observables sequence taken on 10 September 2014.

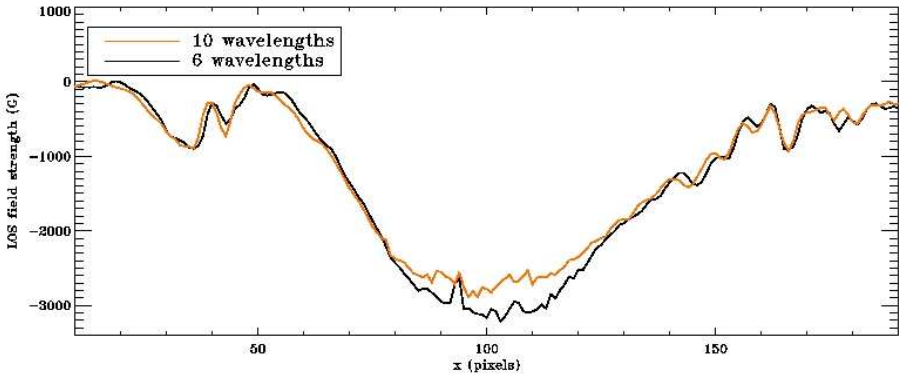

Figure 25. Inversion results of VFISV — cosine of the inclination angle times the field strength — for six and ten wavelengths in the observable sequence. We show a cut in the umbra of sunspot NOAA 12192 on 24 October 2014. A 75-second cadence observables sequence was run with ten wavelength rather than the usual six for an hour, that day.

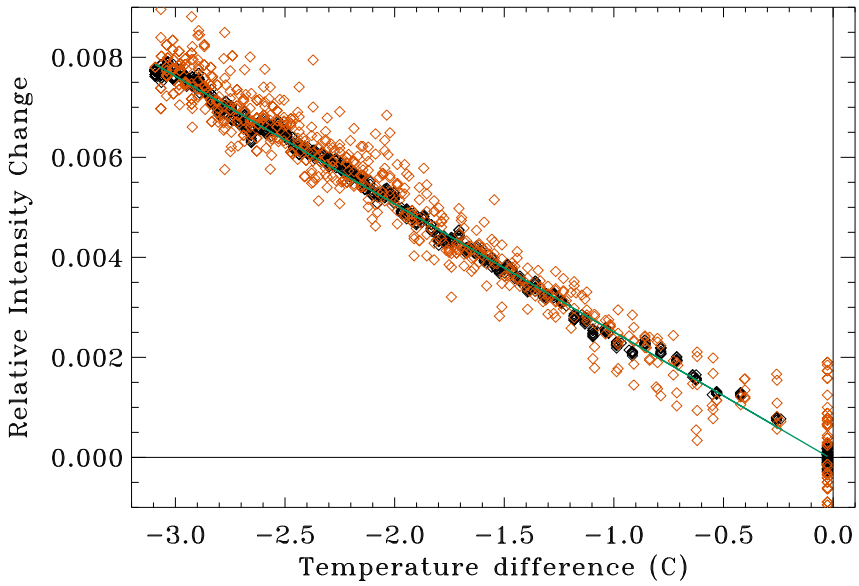

Figure 26. Dependence of the HMI CCD gain on temperature.

4. Other Instrument Issues

In this section, we detail other known issues that have not yet been corrected.

4.1. Temperature dependence of CCD gains

The gain of the HMI CCDs varies with the CCD temperature, as is shown on Figure 26. This figure was obtained from detune sequences taken on September 28, 2011: about 10 hours of detunes were run in Obsmode, and on the front camera of HMI. For a given FID of the level 1 records, the DATAMEDN values were corrected for the non-linearity of the CCDs (this correction is not optimal as it was performed on the median value rather than on each and every individual pixels). The DATAMEDN values were then corrected for changes in Sun-SDO distance. Finally, for a given OBS_VR (Sun-SDO radial velocity), there are two values of DATAMEDN corresponding to two different times of the day, and obtained with two different CCD temperatures. Figure 26 plots the relative change in DATAMEDN as a function of the change in temperature. Low FID values (6000 to 6008) are represented by black diamonds, while higher FID values are shown as red diamonds. The trend line is the result of a linear regression, and the slope of ≈ -0.0025 means that DATAMEDN decreases by about 0.25% when the CCD temperature increases by one degree Celsius.

This sensitivity of the CCD gain to temperature is not negligible, and we plan to include a correction for this effect in the observables codes. Difficulties in dealing with the temperature records have delayed the implementation of this correction so far.

4.2. LCP/RCP cross talk

4.3. Phase diversity

4.4. Stray light and PSF

Stray light is a term used to describe light observed within an image element that does not originate from the corresponding location on the solar disk. Stray light may vary spatially over the field of view and temporally as the optics age. Mathematically, this process can be represented by the point spread function (PSF), its Fourier transform known as the Optical Transfer Function (OTF) and the absolute value of the OTF known as the Modulation Transfer Function (MTF). We report a PSF and deconvolution procedure to remove stray light from HMI data. An MTF was obtained from ground-based calibration data using field stops (Wachter et al., 2012) in order to describe the pre-launch instrumental optics. We first fit the MTF with an exponential convolved with the ideal OTF, often known as a "chat" function. The exponential and "chat" function used to fit the MTF in the frequency domain have the form of a Lorentzian and Airy function in the spatial domain. Post-launch data including solar aureole, lunar eclipse and Venus transit events were used to evaluate how well the PSF was able to reproduce the observed scattering. The PSF reported differs from previous stray light removal efforts since we do not use a single Gaussian or sum of Gaussians as the central mathematical component. Deconvolution was carried out using a Richardson-Lucy algorithm on a graphics processing unit. The recovered images are compared to the originals to determine the increase in the granular intensity contrast and the decrease in minimum umbral intensity. A large, individual sunspot is shown before and after correction for scattered light.

The observed image, O, is modeled by a PSF, Ψ, convolved with the solar image I.

$$O(r,\phi) \;=\; I(r,\phi) * \Psi(r,\phi) \tag{23}$$

A common approach to model Ψ uses a Gaussian for the blurring component (the high amplitude central peak of the function in frequency domain) and a dispersion function or a Lorentzian for the scattering component (the low amplitude tail that accounts for low-level scattering from further away) **??**. The technique of using sums of Gaussians to remove stray light has become popular in the past decade by Mathew et al. (2007) using the sum of three Gaussians and a Lorentzian applied to MDI data, Mathew et al. (2009) using four Gaussians applied to Hinode SOT data, and Yeo et al. (2014) using a sum of five Gaussians applied to HMI data. Another approach is reported by Wedemeyer-Bohm (2008) who derives the PSF for the Hinode Solar Optical Telescope (SOT) by the convolution of ideal diffraction-limited PSFs and Voigt functions. The PSF of optics are best described by an Airy function (Airy, 1938). Gaussians, although versatile and easier to handle mathematically than the Airy function, are not ideal to use in the PSF. Wedemeyer-Bohm (2008) states that a linear

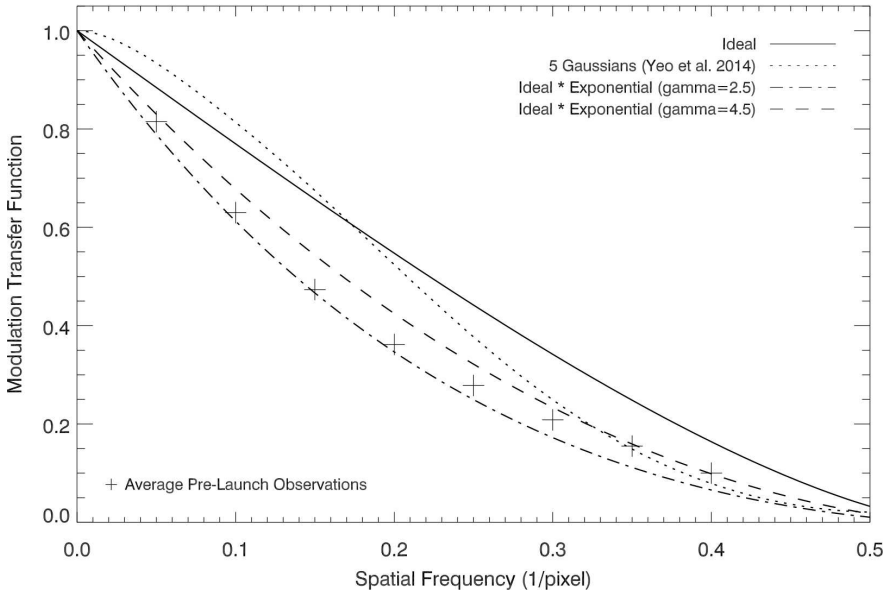

Figure 27. The ideal MTF is shown as a function of spatial frequency. Also plotted are two MTF that are the ideal chat function multiplied by a simple exponential with γ =2.5 and γ4.5. The Yeo et al (2014) MTF is overplotted as is an MTF determined by the ideal multiplied by an exponential function (Equation 9). The symbols represent the average of three of the ground based curves observing during instrument calibration, shown in Fig 2 and reported in Wachter et al. (2012).

combination of functions is, strictly speaking, not correct. Instead, convolution is preferred.

For an incoherent imaging system such as HMI, the ideal PSF is proportional to the Bessel J1 as shown below.

$$PSF_{ideal}(r\prime) = \left(\frac{2J_1(r')}{r'}\right)^2 \tag{24}$$

$$r' = \frac{\pi D r}{f\lambda} \tag{25}$$

where r' is a normalized radius. D is the diameter of the telescope aperture which is 14 cm for HMI. f is the effective focal length which is 4953 mm and λ is 6173 Å. Each HMI pixel is 12 microns so the $r\prime$ value has this as an incremental value. It is worth mentioning that HMI is undersampled, meaning that the size of the pixels are slightly larger than the size that would match the diffraction limit. The OTF is described by Equation 4 as found in **?**.

$$OTF_{ideal}(\rho\prime) = \frac{2}{\pi} acos(\rho') - \rho'\sqrt{(1-\rho'^2)} \tag{26}$$

$$\rho' = \frac{\rho f \lambda}{D} \tag{27}$$

$$MTF_{ideal}(\rho\prime) = |OTF(\rho\prime)| \tag{28}$$

Where ρ is the spatial frequency, $\rho\prime$ is a normalized spatial frequency, and all other symbols are explained above. Because HMI is undersampled, the proper range of values for $\rho\prime$ to use in these equations is $0 - 0.9091$ instead of $0 - 1$. Basically, this indicates that the power does not go to zero at the spatial Nyquist frequency, as is characteristic of an undersampled imaging system.

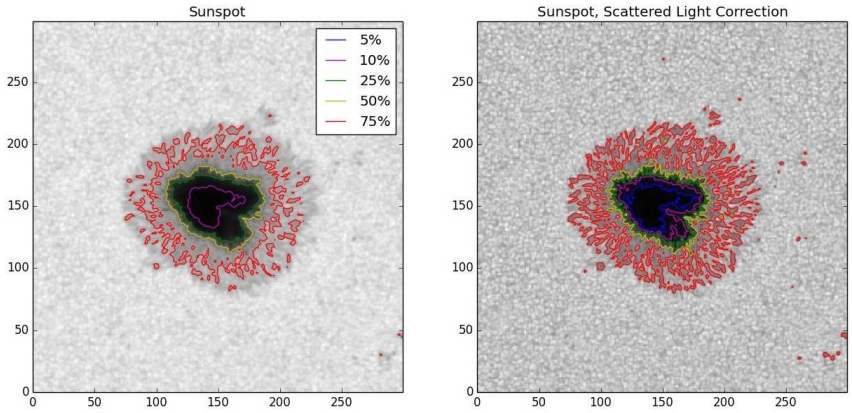

Figure 28. A cropped filtergram image containing a sunspot from the HMI side camera taken on 2013.11.18 is shown at left. The corresponding deconvolved image, with scattered light removed, is shown at right. The dark core of the sunspot changes from being 5.5% of the nearby quiet-Sun continuum intensity in the original image to being 3.3% in the deconvolved image.This change corresponds to temperatures of 3370 in the original to 3140 Kelvin in the deconvolved. The granulation contrast doubles with a standard deviation of the intensity in the quiet-Sun being 3.7% of the average in the original image and 7.2% of the average in the deconvolved image.

The final, non-ideal form of the MTF (Eq 7) and subsequent PSF (Eq 8, where $\mathscr{F}$ denotes the Fourier transform) was determined for the following reasons. First, the ground-based observations showed that the MTF was well-characterized by the ideal MTF (Equation 6) times an exponential function, see Fig 1 and Equation 7. The undersampling of HMI is the reason why the MTF is not zero at a spatial frequency of 0.5 and greater. The final value of gamma in the exponential function of Eq. 7 was determined as $\gamma = 4.5$ through least squares fitting of the transit of Venus data on 2012.06.05 from the side camera data.

$$MTF(\rho') = |OTF(\rho\prime)| \times e^{\frac{-\pi\rho'}{\gamma}} \quad (29)$$

$$PSF(\rho') = \mathscr{F}(MTF) + c \times e^{\frac{-\pi\rho'}{\gamma}} \quad (30)$$

Images of lunar eclipses of the Sun, where a large portion of the solar image is obscured by the lunar disk, enable a measure of the large-scale scattering. One such eclipse occurred on 2010.10.07 and the HMI continuum intensity filtergram shows a light level of 0.34% of the disk-center continuum intensity for a position 200 pixels onto the lunar disk. We found that the light level tended towards a constant far away instead of continually decreasing with increasing

distance from the solar limb. This motivated an additional term, $c \times e^{\frac{-\pi\rho'}{\gamma}}$, to be added to the PSF. This is slightly non-traditional but necessary to fit the tail of the distribution. If we only considered light scattered from 10″ away (such as Wedemeyer-Bohm, 2008 did for SOT), then the additional term would not be necessary. Once we had seen the level off-limb and on the lunar-disk, we wished to include it to emulate the large-scale scattering / dust on the optics. We originally tried to fit the c at the same time as other parameters using the Venus transit data. However, the disk of Venus is too small to effectively measure the large-scale scattering across the CCD.

4.5. Strehl ratio

5. Conclusion

This article describes the computation of line-of-sight and Stokes-vector observables, some of the errors and uncertainties they have, the impact that the HMI calibration on them, and how calibration varied during the life of the mission. The production of HMI observables requires constant monitoring of the instrument for effects that impact, amongst others, filter transmission filter profiles, CCD flat fields, CCD dark currents, etc. Regular calibration sequences are taken on orbit to allow this monitoring. The on-orbit calibration is described in Bush *et al.* (2015), while here we focus on the impact of this calibration on the observables.

References

Bracewell, R. 1995, Two-dimensional Imaging, 191, Prentice Hall, Inc., New Jersey
Couvidat, S., Schou, J., Shine, R.A., *et al.* 2012, Solar Physics, 275, 285
Couvidat, S., Rajaguru, S.P., Wachter, R., *et al.* 2012, Solar Physics, 278, 217
Hoeksema, J.T., Liu, Y., Hayashi, K., *et al.* 2014, Solar Physics, 289, 3483
Jefferies, S.M., and Duvall, T.L. 1991, Solar Physics, 132, 215
Kuhn, J.R., Bush, R.I., Emilio, M., and Scholl, I.F. 2012, Science, 337, 1638
Liu, Y., Hoeksema, J.T., Scherrer, P.H., *et al.* 2012, Solar Physics, 279, 295
Martinez Oliveros, J.C., Couvidat, S., Schou, J., *et al.* 2011, Solar Physics, 269, 269
Mathew, S.K., Martínez Pillet, V., Solanki, S.K. and Krivova, N. A. 2007, Astron. & Astroph., 465, 291
Mathew, S.K., Zakharov, V., and Solanki, S.K. 2009, Astron. & Astroph., 501, L19
Scharmer, G.B., Henriques, J., Kiselman, D., and de la Cruz Rodríguez, J. 2011, Science, 333, 316
Schou, J., Borrero, J.M., Norton, A.A., *et al.* 2012, Solar Physics, 275, 327
Schou, J., Scherrer, P.H., Bush, R.I., *et al.* 2012, Solar Physics, 275, 229
Wachter, R., and Schou, J. 2009, Solar Physics, 258, 331
Wachter, R., Schou, J., Rabello-Soares, M.C., *et al.* 2012, Solar Physics, 275, 285
Yeo, K., Feller, A., Solanki, S., Couvidat, S., Danilovic, S., and Krivova, N. 2014, Astron. & Astroph., 561, 22
Zwaan, C. 1965, Rech. Inst. Astron. Utrecht, 17, 4