

Accuracy analysis of the on-board data reduction pipeline for the Polarimetric and Helioseismic Imager on the Solar Orbiter mission

Kinga Albert^{1,2*}, Johann Hirzberger¹, J. Sebastián Castellanos Durán^{1,3}, David Orozco Suárez⁴, Joachim Woch¹, Harald Michalik⁵ and Sami K. Solanki¹

¹Max Planck Institute for Solar System Research,
Justus-von-Liebig-Weg, Göttingen, 37077, Germany.

²Technical University of Braunschweig, Hans-Sommer-Straße,
Braunschweig, 3329, Germany.

³Georg-August-Universität Göttingen, Friedrich-Hund-Platz,
Göttingen, 37077, State, Country.

⁴Instituto de Astrofísica de Andalucía (IAA - CSIC), Apartado,
Granada, 3004, Spain.

⁵Institute of Computer and Network Engineering at the TU
Braunschweig, Hans-Sommer-Straße, Braunschweig, 3329,
Germany.

*Corresponding author(s). E-mail(s): albert@mps.mpg.de;

Abstract

Context: Scientific data reduction on-board deep space missions is a powerful approach to maximise science return, in the absence of wide telemetry bandwidths. The Polarimetric and Helioseismic Imager (PHI) on-board the Solar Orbiter (SO) is the first solar spectropolarimeter that opted for this solution, and provides the scientific community with science-ready data directly from orbit. This is the first instance of full solar spectropolarimetric data reduction on a spacecraft.

Methods: In this paper, we analyse the accuracy achieved by the on-board data reduction, which is determined by the trade-offs taken to reduce computational demands and to ensure the autonomous

operation of the instrument during the data reduction process. We look at the magnitude and nature of errors introduced in the different pipeline steps of the processing. We use an MHD sunspot simulation to isolate the data processing from other sources of inaccuracy. We process the data set with calibration data obtained from SO/PHI in orbit, and compare results calculated on a representative SO/PHI model on ground with a reference implementation of the same pipeline, without the on-board processing trade-offs.

Results: Our investigation shows that the accuracy in the determination of the Stokes vectors, achieved by the data processing, is at least two orders of magnitude better than what the instrument was designed to achieve as final accuracy. Therefore, the data accuracy and the polarimetric sensitivity is not compromised by the on-board data processing. Furthermore, we also found that the errors in the physical parameters are within the numerical accuracy of typical RTE inversions with a Milne-Eddington approximation of the atmosphere.

Conclusion: This paper demonstrates that the on-board data reduction of the data from SO/PHI does not compromise the accuracy of the processing. This places on-board data processing as a viable alternative for future scientific instruments that would need more telemetry than many missions are able to provide, in particular those in deep space.

Keywords: spectropolarimetry, on-board processing, data pipeline, data reduction accuracy

1 Introduction

The Polarimetric and Helioseismic Imager (PHI; Solanki et al, 2020) is one of the instruments on-board the Solar Orbiter mission (SO; Müller et al, 2020). Solar Orbiter is following heliocentric orbits, that incline relative to the ecliptic plane to access higher solar latitudes. SO/PHI is a spectropolarimeter scanning the photospheric Fe I 617.43 nm absorption line at two different spatial resolutions, through two telescopes: the Full Disc Telescope (FDT) and the High Resolution Telescope (HRT). The HRT is stabilised with an image stabilisation system, that corrects spacecraft jitter and follows the observed features, counteracting solar rotation. SO/PHI samples the spectral line at 6 wavelengths, recording four polarisation states at each wavelength, from which the full Stokes vector (I , Q , U and V), describing the polarisation of the light, can be derived. To obtain a data set with reliable signal-to-noise ratio (S/N), and offer possibilities for trade-offs between S/N and acquisition time, the instrument can parametrise its acquisition scheme. The two most commonly used acquisition schemes are: (1) scanning through the absorption line while recording each of the four polarisation states five times, and accumulating four images in each state, which is completed in less than 100 s or (2) scanning through the spectral line and polarimetric states a single time, accumulating 16 images in each state, completed in less than 60 s (see Solanki et al, 2020). Each data

set results in twenty-four images and provides information about the magnetic field vector and the line of sight velocity at an average formation height of the spectral line. We arrive at these quantities, describing the solar atmosphere, on-board the spacecraft, through a full data reduction pipeline, including the inversion of the radiative transfer equation of polarised light (RTE), assuming a Milne-Eddington approximation of the solar atmosphere. We complement the output of the inversion with the total intensity image from the continuum region next to the absorption line, as well as with metadata about all the details of the data reduction, forming the science-ready data product that is made available to scientists.

In order to facilitate on-board processing, SO/PHI has a custom designed Digital Processing Unit (see [Solanki et al, 2020](#)) on which we implemented a data processing software system (see [Albert et al, 2020](#); [Lange et al, 2017](#)). There were three major drivers in the design of the data processing system: the resource limitations of the hardware, the need for autonomy of the data processing due to the long telecommand to telemetry turnaround times, and the need for the robustness of the system (i.e., to ensure complete and correct data reduction on images from different orbital positions and different solar scenes). To meet the needs with the limited resources, we had to trade off algorithm complexity and computational accuracy.

In this paper, we analyse the effect of these trade-offs on the accuracy of the data reduction pipeline. We compare processing results of a synthetic data set on a representative hardware model of SO/PHI with a reference implementation of the data reduction without trade-offs, which represents the best possible results for the data set. We use synthetic data to exclude errors from sources outside the processing pipeline; these are crucial to analyse, however, they lie outside the scope of this paper. We show that errors accumulated during the processing are negligible, and therefore we achieve the desired quality for the reduced data. We analyse the quality of the Stokes vector achieved by the on-board processing pipeline, the final accuracy of the output data, and the errors introduced during the processing.

2 The on-board data processing

The baseline data processing of SO/PHI consists of the standard spectropolarimetric data reduction steps (see [Fig. 1](#)). The processing pipeline operates on data loaded from mass memory, where we store the acquired raw data. This is also where we store the results of the pipeline, while they wait for data compression and download.

The pipeline starts with dark field and flat field correction. For both of these steps we determine the calibration data (i.e., the dark- and flat fields) on-board, by two separate processes, and store them in the mass memory prior to the initiation of the data reduction. Hence, the only action performed by the pipeline is their loading and their application to the data.

The following step is the prefilter correction. The exact prefilter profiles have been determined on ground at 49 different wavelengths, given in 49 different voltages of the Filtergraph (see [Solanki et al, 2020](#); [Dominguez-Tagle et al, 2014](#)), and uploaded to SO/PHI. The orbits of SO induce a continuous change in the radial velocity of our instrument with respect to the Sun, therefore we determine the voltages for data acquisition as part of instrument calibrations, on-board. This is done such that a reference wavelength λ_0 , falls close to the minimum of the spectral line, and one sample falls into the nearby continuum. At the time of the data processing, we calculate the corresponding values of the prefilter profile for the data set by linear extrapolation of the measured ones, and then apply them to the data.

The polarimetric demodulation is the step that recovers the Stokes vector from the observations. For this step, the pipeline can either use a field dependent demodulation matrix, or one that is uniform across the field of view (FOV; see [Solanki et al, 2020](#)). Preliminary analysis performed on data retrieved up to date from SO/PHI shows that the results are more accurate with a uniform demodulation matrix. Hence, the on-board pipeline currently uses the average of the FOV dependent demodulation matrix, which was measured during the ground testing prior to launch. This is also the demodulation matrix used in this paper. Further improvement to the data demodulation is possible through polarimetric ad-hoc cross-talk correction between the different components of the Stokes vector (see [Sanchez Almeida and Lites, 1992](#); [Schlichenmaier and Collados, 2002](#)). After the demodulation and the cross-talk correction, the pipeline normalises the resulting Stokes vector, using the disc centre continuum intensity. The latter is determined on ground and uploaded to SO/PHI.

To retrieve the physical quantities, we execute the RTE inversion on the normalised Stokes vector. The RTE inversion implemented on-board uses the Milne-Eddington model atmosphere and the Levenberg-Marquardt minimisation method (see [Cobos Carrascosa et al, 2016](#)). The RTE inversion operates on a pixel basis, and needs all values of a pixel from the 24 different images. The pipeline performs this pixel sorting prior to the RTE inversion. The initial conditions for the inversion can be provided in a configuration file, or through numerical calculations, called classical estimates ([Semel, 1967](#); [Rees and Semel, 1979](#); [Landi Degl'Innocenti and Landolfi, 2004](#)). After the inversion we arrive at the physical quantities: the magnetic vector, $\vec{B} = (|\vec{B}|, \gamma, \phi)$ where $|\vec{B}|$ stands for the field strength and γ and ϕ for the inclination and azimuth of the field, respectively, and line of sight velocity, v_{LOS} . They are then sorted back into images. Finally, the pipeline attaches the continuum intensity to them and stores them into the mass memory as the result of the data reduction.

This baseline can be further extended with other modules, such as binning and cropping, used to balance the telemetry volume with the needs of each science case. We also have the capability to further extend the pipeline in the future, e.g., with Fourier filtering to restore the data from optical effects, such as a known point spread function (PSF).

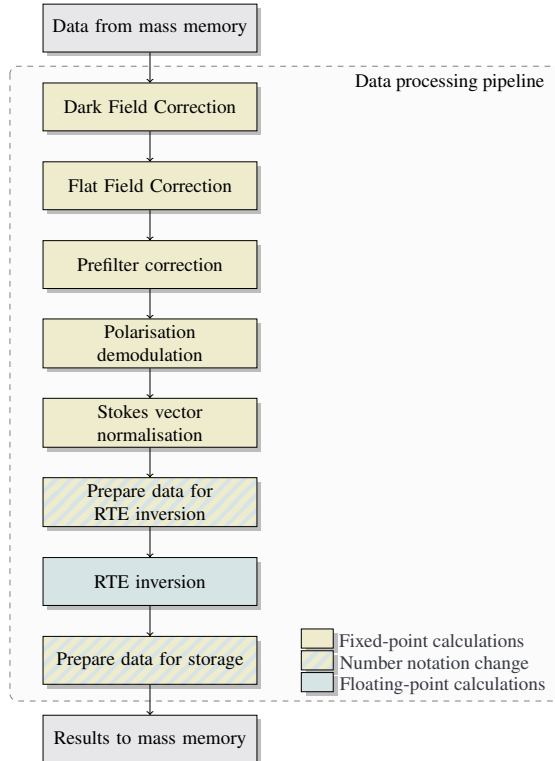


Fig. 1 The current baseline data processing of SO/PHI, executed on-board the spacecraft. To enable the calculations on the limited on-board resources, we combine fixed-point and floating-point number representation in the pipeline.

SO/PHI combines different number representations in the data processing (see Fig. 1). Wherever it was possible, we opted for fixed-point representations as a method to reduce resource usage, trading off accuracy, which is one of the most important sources of accuracy loss in the pipeline. We use floating-point calculations where the accuracy of fixed-point computations did not fulfil the requirements (for instance, the inversion of the RTE). In contrast to floating-point, where number normalisation is inherent in the notation and therefore the precision is better preserved, in fixed-point representation the decimal point is always in the same place, resulting in effectively different number of bits used for the representation of different magnitudes, with 0-padding. For instance, two irrational numbers differing only by a scale factor of 2, both within the range of possible numbers on the allocated bits (i.e., none of them produces overflow), would have an accuracy difference of a bit in fixed point representation, while in floating point they would have the same accuracy. Maximising the accuracy of fixed point representation is possible through scaling up values, to effectively use as many bits as possible. SO/PHI

uses in its data processing 24.8 fixed-point notation, where 24 bits are for the integer part and 8 for the decimal, and single precision, 32 bits floating point.

In order to maximise SO/PHI's processing accuracy, we must make sure that we effectively use all the available bits in the fixed point notation by controlling the magnitude of our data in all steps of the data processing. We always scale the full data set together, to maintain the information in all dimensions of the data: spatial, spectral, and polarimetric. The individual pixel values in the images have no physical meaning throughout the pipeline, it is the normalisation of the Stokes vector by the disc centre continuum quiet Sun intensity (denoted simply as I_c) that creates the suitable input to the RTE inversion module.

Our starting point for scaling the data through the pipeline is in the detector. The exposure time adjusts the brightness of the solar scene such, that it is reliably represented on 12 bits read out from the detector. Then, we accumulate a number of detector readouts to increase the signal-to-noise ratio of the solar data, and pad the result with zeros after the decimal, to reach the 24.8 fixed point notation. Then, we calculate the largest possible integer number that we can obtain through these operations, called from here on *maximum range*, and place it into the metadata of the data set for further reference. For instance, for 20 accumulations, the maximum range would be 20×2^{12} . At the start of the data processing, just after loading our raw data set from the mass memory, we scale up the data to use all the available bits. This is achieved by multiplying with a scale factor calculated as the ratio of 2^{23} (the largest possible integer in two's complement on 24.8 fixed-point notation) and the current maximum range. In each operation that follows, we aim to preserve this largest possible maximum range in the result, by considering the magnitude of the operands and that of the result. For simplicity, we only keep track of the maximum range. However, in few cases the minimum range of the absolute value of the data is also relevant, for instance, when we perform divisions like that necessary to correct the flat field or the prefilter. Here, the smallest value in the divisor determines the maximum range of the result. Since these are not tracked, we make assumptions about the divisor, with the consequence that any pixels with smaller values will create an overflow and the resulting pixel will become not-a-number (NaN).

We scale all calibration data to no higher value than to represent the precision with which they are determined (e.g., in the case of the flat fields we only use 3.8 bits). In those cases where the accuracy of the calibration data is high, we do a trade-off between the accuracy of the data and that of the calibration data. We decide all trade-offs on a case by case basis, based on simulations to verify which scaling gives the best precision results.

3 Test setup

To compare the accuracy of the on-board processing pipeline to what could be achieved on-ground, and isolate it from other sources of errors (e.g., solar




Fig. 2 The data set used in the tests contains an MHD simulated sunspot (Rempel, 2015). We synthesised the spectral line and its nearby continuum from the MHD cube (left) and degraded, as described in the main text, to what we would expect from SO/PHI (right) shown here at the continuum intensity sample wavelength, $\lambda_0 - 300 \text{ m\AA}$. The most obvious result of the degradation is the loss of RMS contrast. The size of the data is 1024×1024 pixels, each pixel corresponding to 48 km on the Sun. The spatial sampling is larger than what SO/PHI achieves at closest approach. However, it is preserved to provide more pixels for statistical analysis. The arrows indicate pixels for which the Stokes profiles are plotted in Figs. 3 and 4.

evolution, and the accuracy of the calibration), we have chosen to process a synthetic data set. We process this data with the pipeline described in Sect. 2 on the Qualification Model (QM) of SO/PHI, which is fully representative in terms of the Data Processing Unit of the Flight Model. We compare the QM results to a reference pipeline run in a PC, in double precision floating point, representing the best possible processing accuracy. The one exception is the RTE inversion, which is performed on the QM in both cases. In the case of the reference pipeline, we upload the RTE input data in floating point and download directly the floating point results that it produces. The data processing pipeline uses the same calibration data that we apply in the data preparation, therefore we have no inaccuracy originating from the determination method and processing of the calibration data. This means that all errors presented in this work are inaccuracies originating in the data reduction pipeline.

The test data set is a magnetohydrodynamic (MHD) simulation of a sunspot (Rempel, 2015). From this MHD cube we synthesized the Fe I 6173 Å spectral line profiles, applying a wavelength sampling of 14 mÅ, with the SPINOR code (Frutiger et al, 2000). SPINOR relies on the STROPO routines to solve the RTE (Solanki, 1987). The simulation is 1024×1024 pixels, with a pixel size of 48 km. The pixel size of the HRT telescope of SO/PHI, at closest approach, corresponds to 101 km, however the dataset is not resampled,

in order to provide more pixels in the umbra and penumbra for statistical analysis.

We degrade the synthesized spectral line profiles in several steps, starting by convolving the wavelength dimension of the synthesised data with the transmission profile of the SO/PHI Filtergraph (see Solanki et al, 2020; Dominguez-Tagle et al, 2014):

$$S_p^{conv}(\lambda, x, y) = S_p^{synth}(\lambda, x, y) * F(\lambda), \quad (1)$$

where $*$ denotes convolution, the index p runs over the four polarimetric modulation states, λ denotes the wavelength, x and y are the spatial image coordinates in pixels. S_p^{synth} are the Stokes profiles synthesized from the MHD cube, F denotes the filter profile, and S_p^{conv} is the Stokes vector from the synthesis, convolved with the spectral profile. We then select the samples for SO/PHI from the resulting spectral profiles. The samples are defined relative to a reference wavelength (λ_0 , which we chose for this test to be 6173.371 Å). λ_0 is placed in the vicinity of the absorption line minima:

$$S_p^{samp}(\lambda, x, y) = S_p^{conv}(\lambda_s, x, y), \quad (2)$$

where λ_s denotes the sample wavelengths in reference to λ_0 :

$$\lambda_s = [\lambda_0 - 300 \text{ mÅ}, \lambda_0 - 140 \text{ mÅ}, \lambda_0 - 70 \text{ mÅ}, \lambda_0, \lambda_0 + 70 \text{ mÅ}, \lambda_0 + 140 \text{ mÅ}],$$

and \vec{S}^{samp} is the Stokes vector sampled in wavelength.

After spectral sampling, we convolve each individual image of the data set (the 24 images, six wavelength samples and four modulation states) with the spatial PSF in the shape of a Lorentzian function, configured with the theoretical parameters of the HRT telescope, adjusted to the plate scale of the simulation:

$$S_p(\lambda, x, y) = S_p^{samp}(\lambda, x, y) * A(x, y), \quad (3)$$

where A represents the PSF of the SO/PHI HRT.

The convolution with the PSF reduces the continuum quiet Sun root-mean-square (RMS) contrast of the synthetic data from 22.83% to 7.9%. The next step is the polarimetric modulation of the synthetic data:

$$I_m^{mod}(\lambda, x, y) = k \sum_{p=1}^4 M_{mp}(\lambda) S_p(\lambda, x, y), \quad (4)$$

where index m denotes modulation states, I_m^{mod} is the modulated data set, M_{mp} is the polarimetric modulation matrix, and S_p is the Stokes vector from Equation (3). The constant k adjusts the quiet Sun continuum mean intensity (I_c) of the data to what is representative of the SO/PHI HRT, converts the number of photons collected by the detector to digital numbers, read out from

the electronics. This constant, furthermore, accounts for the frame accumulations that we perform in order to increase the signal-to-noise ratio of the data, which we adjust to 20 frames for the tests. Due to the higher RMS contrast in the test data, and therefore higher dynamic range, we adjust the I_c slightly below the level observed in the SO/PHI HRT data. The ratio of I_c for the observed to the test data is 1.08. It is important to remark that this difference in I_c creates a slightly worse case from the point of view of accuracy: as the values in the test data are somewhat lower, they are represented on fewer bits, when compared to SO/PHI observations. We remark, that the order in which we apply the instrumental degradation to the synthetic data has been established for the sake of convenience and to be able to correct the data in the same order with the pipeline. For instance, we have applied the polarimetric modulation here, even though the modulation package in the instrument is right after the entrance window following the beam direction. Since only the order of linear operations have been exchanged (only the dark field application is non-linear, which we apply according to the optical path), this does not affect the resulting input data.

Once we have modulated the input Stokes vector and converted the data to digital numbers, we apply the prefilter profiles, dark and flat fields to the synthetic data:

$$I_m^{obs}(\lambda, x, y) = I_m^{mod}(\lambda, x, y)I^p(\lambda, x, y)I_m^f(\lambda, x, y) + I^d(x, y), \quad (5)$$

where I_m^{obs} is the data set produced to match SO/PHI observations, i.e., the input to the processing pipeline, I^p is the prefilter profile which is different for each wavelength and pixel, I_m^f are the flat fields, depending both on wavelength and polarisation states m , and I^d is the dark field of the sensor, the same for all wavelengths and polarisation states.

It is worth noting that the addition of the flat field and dark field to the data further reduces its RMS contrast to 6.4%. During the first months of SO/PHI operations, we have observed the RMS contrast in SO/PHI HRT observations of the quiet Sun to be around 4.5% to 5%, which is expected to change further as SO/PHI changes its distance to the Sun. The smaller the dynamic range in the data (i.e., lower the RMS), the closer we could take all values of it to the maximum possible range, hence achieving a better overall computational accuracy. For this, we could do an adjustment in the exposure time, or in the maximum range assumptions at data acquisition. This is, however, not done in the case of our test data. Therefore, the results are representative of what we would achieve on SO/PHI data. Also note, that the dark- and flat fields that we apply here are calculated on ground from SO/PHI data, and calculated on-board SO/PHI, respectively. These data we downloaded during the commissioning phase of the mission. We produced this early calibration data with methods that we will improve further, however these are representative of the flat- and dark fields that we expect after fine-tuning. Furthermore, it is important, that in order to not introduce further uncertainty into the process, and be able to assess the effects of the numerical

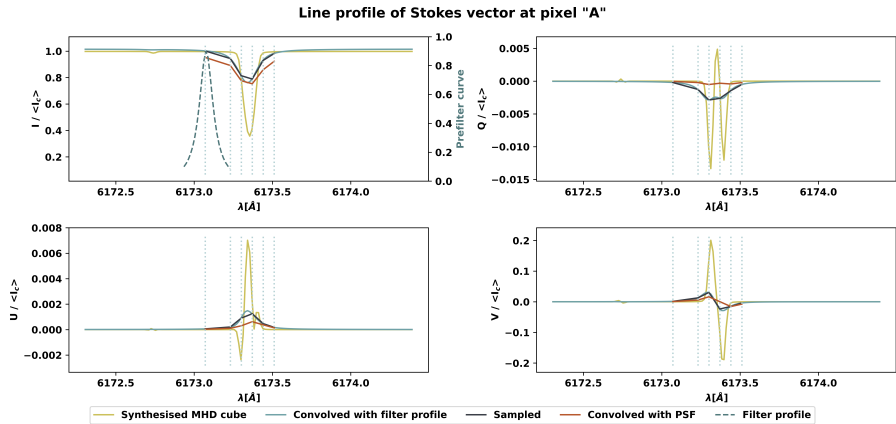


Fig. 3 The spectral line profile of the input data set shows the degradation of the synthesised MHD data. We first convolve the synthesised profile with the filter profiles of SO/PHI (shown in the top left panel), then select the correct wavelength samples, followed by the convolution of the resulting images with the theoretical PSF of the instrument. This is a bright pixel from the quiet Sun. The convolution with the filter profiles significantly reduces the spectral line complexity. We indicate the location of the plotted pixel in Fig. 2.

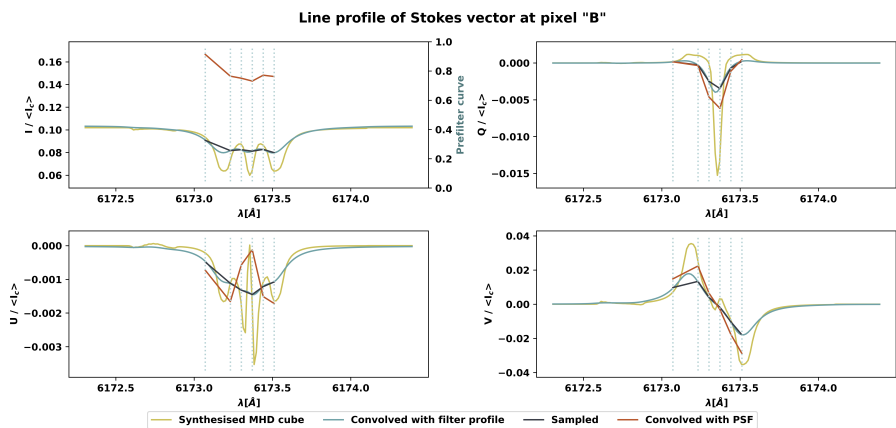


Fig. 4 Same as Fig. 3, for a pixel from the umbra. This dark pixel is surrounded by bright structures, therefore after applying the PSF, there is a large change in the intensity level at continuum from neighbouring pixel contributions. There is also a significant change in the U profile. This is an extreme case, with neighbouring pixels being significantly different from the one selected. Here, the PSF convolution has a much stronger effect on the final profiles, when compared to Fig. 3. We indicate the location of the plotted pixel in Fig. 2.

errors on the RTE inversion, we omit several effects that would appear in real data. We do not introduce noise to the data in the course of its degradation. Likewise, we use the data as instantaneous snapshots of the solar scene without considering the evolution of the solar scene, rotation of the Sun or spacecraft jitter.

We show the effects of the data degradation at the continuum wavelength in Fig. 2. The most obvious effect is the reduction of the image RMS contrast. In the line profiles of the data (see Figs. 3, and 4) we can see that the complex profiles from the MHD simulations smooth out significantly as a result of the convolution with the transmission profiles of the SO/PHI Filtergraph. The same operation also lowers the amplitude of the polarisation signals. The sampling of the data further removes details of the spectral shape by reducing the available information. This effect is especially strong in the sunspot profiles, due to the complex shapes. The spatial PSF convolution strongly changes the Stokes I intensity, especially in the umbral profile, as usually stray light does in real observations. We do not expect the RTE inversion to perfectly reconstruct the resulting profiles, producing especially large differences in the umbra, as it does not account for the stray light. The spatial PSF can lower the amplitude of the polarisation signals further (e.g., in quiet Sun areas) although it is not always the case (as shown here in the umbral profiles) since the final effect on each pixel depends on the surrounding signals. Note, that the final, degraded profiles appear different with respect to the synthesised data. This is mainly an effect of the wavelength sampling. While we sample the absorption line with symmetric offsets, the reference wavelength does not necessarily coincide with the centre of the line, causing a shift in the sampling, and introducing an apparent asymmetry even in the case of symmetric profiles. This, however, does not affect the performance of the RTE inversion. The relatively few samples (only five) also contribute to the strong difference.

We analyse how the data pipeline changes the accuracy of the data throughout each step of the pipeline, grouping them into three categories: the input and the first steps of pre-processing, the polarimetric errors, and the physical parameters. In the first category, we start by analysing the input errors introduced by the data quantisation to fixed point notation. Next, we look at the early pre-processing errors, which are introduced by the first three pipeline steps: dark field, flat field and prefilter correction. The next category comprises the polarimetric sensitivity, where we evaluate the errors in the determination of the Stokes vector during the demodulation process and the cross-talk correction. Achieving a good polarimetric sensitivity of the Stokes vector is the most important requirement of SO/PHI. Finally, in the third category we give a glimpse into the resulting physical parameters, in which we analyse the results of the inversion. These results correspond to the science ready data obtained directly on-board.

4 Analysis

For the analysis, we separate the data into three zones, based on the Stokes vector signal levels: the umbra, the penumbra and the quiet Sun. The regions are defined on the conservative side, with preference on excluding pixels from them, rather than including pixels that do not clearly belong. We consider > 10000 , > 81000 , and > 94000 pixels in the three regions, which corresponds

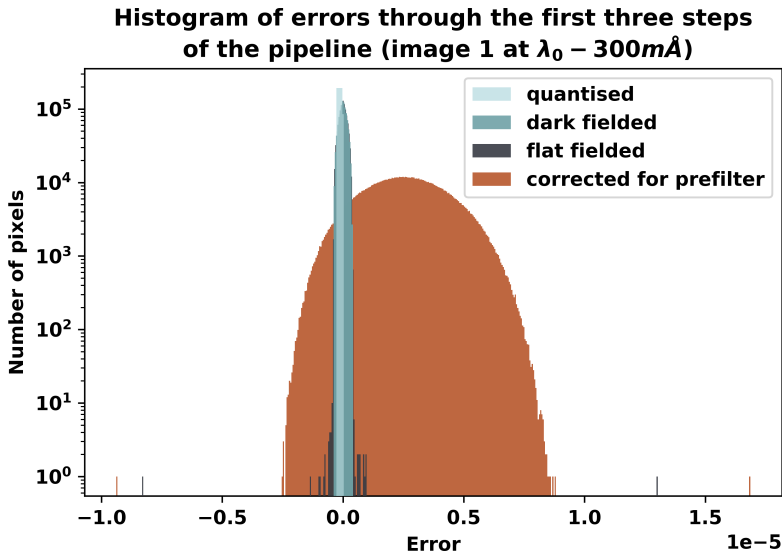


Fig. 5 The histogram of errors after data quantisation and the first three steps of the pipeline shows the accuracy achieved after each step, normalised to the mean of the data. The quantisation and dark field subtraction introduces small errors. The errors introduced by the flat field division are very similar to that of the dark field correction, however a few outliers show up from dust grains in the FOV. The prefilter correction produces a residual of the image with very low intensity due to the inaccuracy at the extrapolation of the prefilter profile to the correct voltage.

to 1%, 7.8% and 90.5% of the full field of view, respectively. We follow this definition in the rest of the paper, and mark these regions on the figures.

4.1 Input and first steps of pre-processing

The first source of error is the quantisation error of the input data to the pipeline. After calculating them in double precision floating-point, according to the description in Sect. 3, we transform the data set to the fixed-point representation as the raw data would be stored: detector read-out on 12 bits, accumulated 20 times, and padded with 0-s for the decimals. This corresponds to a maximum range of 20×2^{12} . The RMS of the error resulting from the quantisation, normalised to the image mean intensity, is between 1.55×10^{-7} and 3.5×10^{-7} across the FOV of the 24 images (the six different wavelengths and four polarisation states). This is consistent with the $1 \div 2^8$ precision of the decimal in the 24.8 fixed point representation. The profile of the histogram is flat, as expected for quantisation errors (see Fig. 5), with a systematic bias towards smaller numbers in the quantised data, due to using bit truncation, rather than rounding.

The first steps of pre-processing are: the subtraction of the dark field from the data, the division of the data by the flat field, and the division of the data by the extrapolated prefilter profiles. After the loading of the data, the pipeline

scales it up by $2^{23} \div (20 \times 2^{12}) = 102$, which is also followed by the scaling of the dark field. The errors after this step originate from data quantisation both in the input (as described before) and the dark field. The subtraction operation itself does not produce any errors, inherently. After this step, we have an error RMS ranging between 1.55×10^{-7} to 7.7×10^{-7} , normalised to the image mean in the 24 images of the data set. This step slightly changes the profile of the error histogram, removing the bias caused previously by bit truncation (see Fig. 5).

The following step, the division by the flat field, divides the data with a max range 2^{23} , i.e., effectively using 24.8 bits, by data with max range 2^3 , i.e., effectively using 3.8 bits. The flat field has been normalised to its mean intensity, scaled by 2^3 , and we assume its minimum range to be 2^2 . Any number below this may cause an overflow, however there is still some room for smaller values, as the data does not fill up the full detector well at acquisition. After the division, we readjust the magnitude of the result to 2^{23} by multiplying it with 2^2 . The errors in this step originate from the errors on the input data (as shown in previous steps), the representation error of the divisor, and the representation errors of the result. In the histogram, a few pixels with larger errors appear due to dust grains in the FOV (which translates to very small numbers in the divisor). However, these errors are only in a handful of pixels, not contributing significantly to the RMS calculated over the full FOV, which ranges between 1.56×10^{-7} and 7.7×10^{-7} , normalised to the image mean intensities.

The next step, the prefilter correction, starts with the interpolation of the prefilter. The pipeline performs this on data scaled to a maximum range of 2^{23} , then scales it down to a maximum range of 2^{10} . Furthermore, we assume a 2^9 minimum range. The interpolation of the data creates an error in the divisor, compared to what we obtain on ground. Then, through the division, the error histogram widens: when subtracting the results, a very small amplitude residual of the data remains. The histogram profile seen in Fig. 5 is the histogram of the test data. After the operation, we scale the result of the division back to 2^{23} by multiplying it with 2^9 . The RMS of the error across the FOV ranges between 3.1×10^{-6} and 3.75×10^{-6} , normalised to the mean intensity of the images.

4.2 Polarimetric errors

The Stokes vector is the output of the polarimetric demodulation of the data. The pipeline is able to do further adjustments with ad-hoc polarimetric cross-talk correction methods. However, for this study, we will limit ourselves to errors due to the demodulation process. We scale the demodulation matrix to the maximum range 2^9 , for which we have to account in the input data. In order to avoid overflow, we divide the input data by $2^9 \times 4$, where the number 4 accounts for the addition of the rows in the 4×4 matrix multiplication. The result of the operation is then on a maximum range of 2^{23} . Since the values of the Stokes vector at this point do not have physical meaning, we normalised

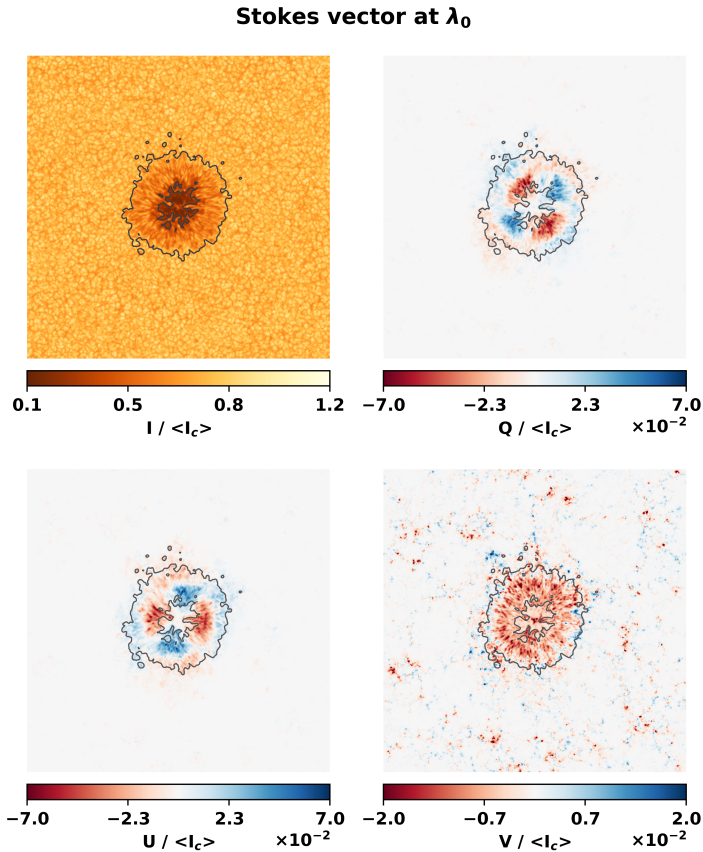


Fig. 6 The normalised Stokes vector, $\vec{S}/I_c = (I, Q, U, V)/I_c$, is shown here at the reference wavelength, λ_0 (6173.371 Å), close to the minimum of the absorption line. As expected, Stokes Q/I_c and U/I_c shows the strongest signals in the penumbra. This is also true for V/I_c , which is due to the large splitting of the line in the umbra, causing a poor sampling by SO/PHI. (see Fig. 3).

them to the mean of the quiet Sun intensity, I_c , before showing them in Fig. 6. The normalisation here is performed in double precision, to show the results without the error introduced by this step, when performed on-board.

The polarimetric sensitivity requirement set for SO/PHI is 10^{-3} , which is met during the processing: the errors accumulated by the end of this step have an RMS across the FOV between 3×10^{-6} and 4.7×10^{-6} (see Fig. 7). This leaves a generous margin to other sources of error, and does not compromise the required polarimetric precision.

The errors across the different Stokes parameters, the different wavelengths and the different regions of the data (umbra, penumbra, and quiet Sun) differ by a maximum of 1.7×10^{-6} , which is considered negligible. The variation of error is visible for I on a linear scale. However, for Q , U and V we need a logarithmic scale to illustrate the differences. The errors in the result depend

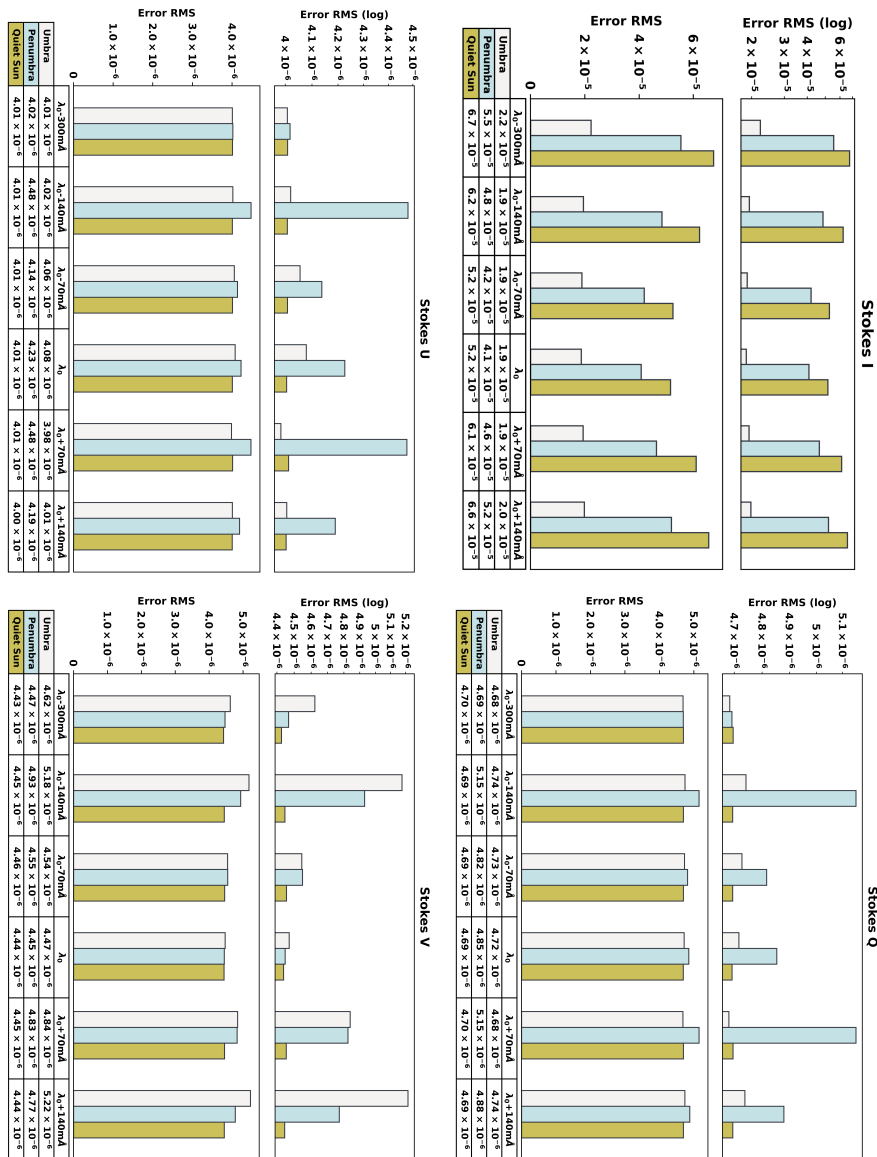


Fig. 7 The RMS of the errors in the Stokes parameters varies between 4.0×10^{-6} and 6.6×10^{-5} , which leaves a good margin to meet the 10^{-3} polarimetric sensitivity requirement of SO/PHI. The errors vary by 6.5×10^{-5} across the Stokes parameters, wavelength and solar regions, as shown in the table on the bottom. Their relation in linear scale (bottom bars) shows that Stokes Q , U , and V are very close to each other, therefore to better see their differences, we use a logarithmic scale (top).

on four factors: the errors that were accumulated prior to this step, the magnitude of the input data, the magnitude of the output data and the terms of the demodulation matrix. These first and second sources oppose each-other: the so far accumulated error correlates linearly with the intensity (due to the image residual after prefilter correction), while the input data is better represented where the values are larger, therefore it has an inverse correlation. The demodulation matrix has only positive terms in the row producing Stokes I , however it has negative terms for Q , U , and V , producing cancellation effects.

Results of the RTE inversion




Fig. 8 The results of the RTE inversion, shown here, together with the continuum intensity image, form the science ready data that is transferred to ground. The results are consistent with what is expected from such a data set: $|\vec{B}|$, γ and v_{LOS} shows smooth transitions between the structures, with magnetic fields up to 4000 G. In v_{LOS} we can see the up and down flows of the solar granulation, as well as of the Evershed flow. ϕ is dominated by noise in the quiet Sun, however it does show the fan-like structure around the penumbra, as expected.

The trend in errors in Stokes I approximately follows the intensity of the output, with only a slight deviation from this trend along the spectrum (see

Fig. 7). This is a case where the intensity in the final results dominates the magnitude of the errors. In Stokes Q , U and V the intensity of the result does not overpower the trend any more, and there is a cancellation effect of previous errors due to negative terms in the demodulation matrix. The result of all these values is a trend in errors that is stochastic.

4.3 Physical parameters

We reach the final physical quantities through the RTE inverter (Cobos Carrascosa et al, 2016). The inversion module can be configured in five different modes, depending on the desired outputs, and on the initial model (apart from special debugging modes). The first three can provide all atmospheric parameters of the Milne-Eddington model (line-to-continuum absorption coefficient ratio, Doppler width, damping coefficient, source function — its slope and its value at the top of the atmosphere, magnetic field vector $\vec{B} = (|\vec{B}|, \gamma, \phi)$, and LOS velocity, v_{LOS}). These modes are: inversion starting with an initial model, called classical estimates, calculated with analytical formulae (centre of gravity technique, see Semel (1967), and the weak-field approximation, see Landi Degl'Innocenti and Landolfi (2004)), inversion with a configurable initial model, and the classical estimates without being followed by an inversion. Aside from these modes, we have two others which only return line of sight (LOS) parameters: the longitudinal mode where we only obtain LOS velocity and LOS magnetic field, and the no polarisation modulation mode where we only obtain LOS velocity. We select the modes based on the required science return and available telemetry.

The output parameters of the inversion are also configurable. Thus, we can request only a subset of the full set of output (physical) parameters for a given mode. For the first modes that calculate all nine Milne-Eddington model parameters, in standard operations we only request the four parameters of interest: the three components of the magnetic field vector and the LOS velocity. For the other modes in standard operations, we select all available outputs.

Before we perform the RTE inversion, we must prepare the data to match the interface of the inverter. After demodulation, the data are normalised to I_c (I_c is calculated on ground). The pipeline performs this operation with I_c represented on 12.8 bits (which corresponds to it being calculated on data with maximum range 2^{12}), and scales the result to a maximum range of 2^{23} . We transform these data from fixed to floating point by assuming a 2.30 fixed point representation. As part of the preparation of the data, the pipeline also rearranges the images, to provide the inverter with a data stream that it can process (i.e. all 24 values that the same pixel in the FOV takes in the data set). A similar step takes place after the inversion, to form images again.

We use the first mode of the inverter in this study, where we do the full inversion of the data, with initial conditions calculated through classical estimates. See Fig. 8 for the results of the RTE inversion. These results show smooth transitions of the values over the FOV in $|\vec{B}|$, γ and v_{LOS} . In $|\vec{B}|$ we




Fig. 9 The RMS of the errors in the RTE inversion results primarily reflects the stability of the inversion in the different regions. It shows how the small errors in the input data, introduced through the processing, affect the final results retrieved with the same method. The error in $|\vec{B}|$ and γ is smallest in the penumbra. The determination of γ and ϕ is challenging in the quiet Sun, due to low signal levels, which is also reflected in the error RMS. v_{LOS} in the umbra has higher error due to a shallower line core and more complex profiles. (The magnitude of the errors is discussed in Sect. 5)

obtain magnetic fields up to 4000 G. The upper limit in the inversion module for $|\vec{B}|$ is 5000 G, which is not reached. In v_{LOS} we can see the up- and downflows of the solar granulation, as well as the Evershed flow in the penumbra (see Evershed, 1909). The upper and lower limits of the inversion module for v_{LOS} is $[-20, 20]$ km/s. In ϕ , the azimuth ambiguity disrupts the smooth transitions, and is dominated by noise in the quiet Sun.

The source of the differences that we obtain between the reference pipeline and the SO/PHI processing pipeline, are the small variations of the input data due to the processing. However, the stability of inversion on such a data set (considering the physics of the MHD simulation, the spectral and spatial convolutions applied, and the spectral sampling points) also determines the amount of error introduced by these small changes. In case of real SO/PHI observations, we would have an additional contribution from the inaccuracy of the calibration data. This is eliminated in this study by using the same data to degrade and calibrate the synthetic data set. In the following paragraphs, we present the results; we discuss their magnitude and significance in Sect. 5.

The RMS of the errors in $|\vec{B}|$ introduced by the SO/PHI pipeline vary between 16.9 G to 5.8 G from the quiet Sun to umbra, see Fig. 9). Values below 1000 G show a larger disagreement as a consequence of lower signal levels, both in the case of the quiet Sun and the penumbra, showing up as a low density scatter in Fig. 10. We note, that pixels with values 1000 G belong to the outer

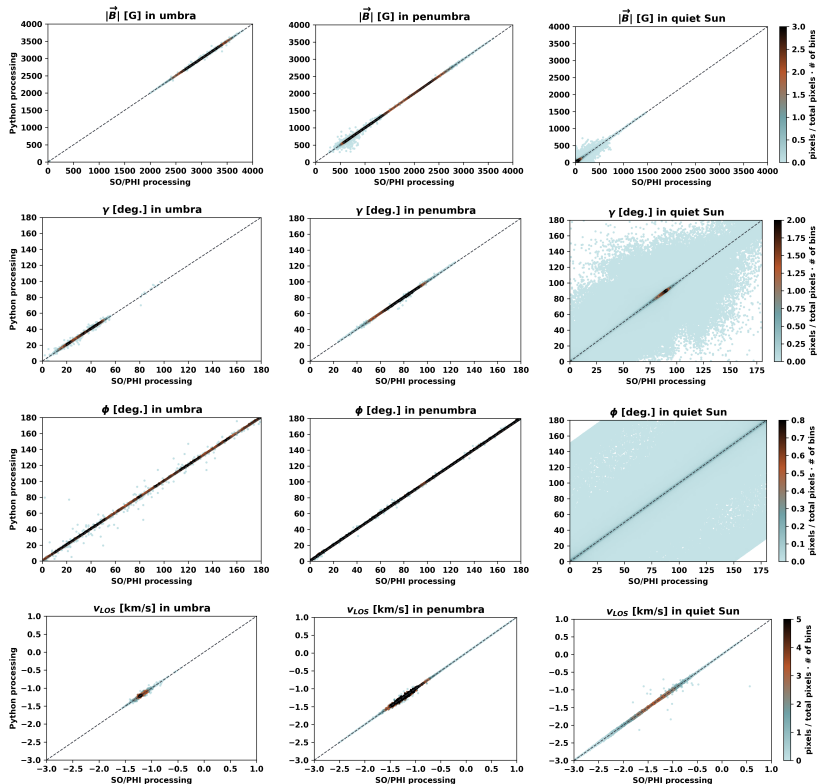


Fig. 10 The correlation of the reference and SO/PHI processing results for $|\vec{B}|$, γ , ϕ and v_{LOS} . While the Stokes vector errors are uniformly distributed in these regions, (see Fig. 7, the scatter of the results varies significantly. This is due to the different stability of the inversion in the different regions, discussed in the main text.

penumbra, reaching the lower limit of typical field strengths in penumbral regions. A few outliers appear also in the umbra, particularly above 2500 G, hinting at the difficulty in inverting complex line profiles. These profiles, on one hand, are sampled by only six points, on the other hand they are also significantly changed by the prefilter and spatial PSF convolution (see Fig. 3).

The errors in the orientation of the magnetic field vector we present to integer precision. This is due to the fact, that the RTE inversion truncates these values to integer accuracy, therefore any sub-decimal precision differences between the errors in the various regions would be an artefact of this operation. The umbral and penumbral regions show the same accuracy for both the inclination, γ , and the azimuth: ϕ , 1° and 2° , respectively. (see Fig. 9). In the quiet Sun, we see an error increase in both angles, a sign of lower signal strength. This results in noisy vector direction, with the correlation plot showing a large scatter across all possible values. The RMS of the errors in the quiet Sun regions is 13° and 38° for γ and ϕ , respectively. In the analysis of ϕ , we assume that any difference between the two results larger than 150° is

caused by the ambiguity of the angle, and therefore we change all corresponding pixels to their supplementary angles. This results in the sharp cut at these errors, showing up as empty corners, in the last panel of Fig. 10.

v_{LOS} is calculated with the best precision in the penumbra and the quiet Sun, with an error RMS between 3.0-3.5 m/s (see Fig. 9). The increase of the accuracy in the penumbra is due to the stronger velocities in this region. The calculation of v_{LOS} in the umbra produces an error with 5.5 m/s RMS. These regions, with higher magnetic fields, produce stronger Zeeman splitting: the spectral profiles widen, become more shallow and complex. Due to our spectral sampling (as seen in Fig. 4), the representation of these profiles becomes less accurate, and therefore, the sensitivity of the Stokes vector to the LOS velocity perturbations diminishes, leading to larger errors.

Using the vector magnetic field that we retrieve from the RTE inversion, we also calculate the line of sight magnetic field (B_{LOS}), for further insight. The RMS error of the B_{LOS} is 16.9 G, 7.9 G and 5.8 G for the umbra, penumbra and quiet Sun, respectively. It shows the best agreement in the quiet Sun, significantly better than what $|\vec{B}|$ provided. This is due to the fact that the granules harbour weak transverse fields (see Orozco Suárez and Bellot Rubio, 2012; Danilovic et al, 2016), which translate to small B_{LOS} values, lowering their contribution to the RMS of the error. In contrast, in the umbra and penumbra, we can observe an increase in the error. These regions harbour stronger magnetic fields, which appear at higher angles, all the way to close to vertical in the umbra. Their orientation and strength translates to strong B_{LOS} signals, creating a higher RMS error. We note, that B_{LOS} can also be computed on-board without using an inversion, with analytical formulae (see above), which is a different approach, and these values can not be applied to it.

5 Discussion and Conclusions

Our accuracy analysis shows that the on-board processing trade-offs do not compromise the accuracy of the SO/PHI data. The comparison between the results of the fully on-board processed data (with the necessary trade-offs) and the results obtained on-ground (without the trade-offs of the on-board processing) conveys that the errors in the determination of the final Stokes parameters are below 7×10^{-5} . This leaves a good margin for other sources of errors (e.g., calibration errors) to fulfil the 10^{-3} polarimetric sensitivity requirement of SO/PHI.

We present the errors in the physical parameters to give an idea of how the RTE inversion results may change by these processing inaccuracies. However, it is important to remark, that most input data sets intrinsically deviate from Milne-Eddington line profiles, and a tiny error in the input may cause the inversion to converge to a different local minimum, providing large differences in the physical parameters.

In [Albert et al \(2019\)](#) we have done a similar analysis on a data set acquired by the Helioseismic and Magnetic Imager on-board the Solar Dynamics Observatory (SDO/HMI; [Schou et al, 2012](#)), using an earlier version of the SO/PHI data reduction pipeline. In that study, we compared the errors from the processing to a statistical analysis performed with the He-Line Information Extractor inversion code (HeLLI_x+; see [Lagg et al, 2004](#)). The statistical mode of HeLLI_x+ inverts the data with varying starting conditions, and we regard the variation of the results as a measure of the inaccuracy inherent in the inversion of the data set, amounting to 21.9 G, 1.34°, 1.37°, and 14.5 m/s for $|\vec{B}|$, γ , ϕ and v_{LOS} , respectively. The errors in the inversion results, introduced to the SDO/HMI data by the on-board processing of SO/PHI ([Albert et al, 2019](#)), are 33.64 G, 2.56°, 1.92°, and 19 m/s for the same parameters, which are slightly higher than what we find in the current study. This is due to the different input data, as well as the earlier version of the processing pipeline. It is important to note, that in the [Albert et al \(2019\)](#) study we determined the errors in the magnetic field vector limiting the FOV to an area with strong polarisation signals, while for the v_{LOS} we take the full FOV into consideration.

[Cobos Carrascosa et al \(2016\)](#), while verifying the implementation of the RTE inverter on-board SO/PHI, compared inversions with the C and FPGA implementation of the code for a collection of Milne-Eddington synthetic profiles (considered as ideal input, containing only symmetric profiles). These profiles were sampled with 5 mÅ steps, and selected to range between 0 and 1500 G in $|\vec{B}|$, 0 and 180° in γ and ϕ , and -2 and 2 m/s for v_{LOS} , to which they added noise with a magnitude of $10^{-3} \times I_c$. The results agree with an error RMS of 5.3 G, 4.86°, 5.77° and 5.9 m/s for $|\vec{B}|$, γ , ϕ , and v_{LOS} , respectively. However, the same comparison, performed on observations from the Swedish Solar Telescope, results in RMS errors of 69.2 G, 6.5°, 5.47°, and 79.41 m/s. This is due to several factors, including higher noise, instrumental errors, and asymmetric data profiles, which is typical of observed solar Stokes profiles (see e.g. [Solanki, 1993](#)). The data that we analyse in this paper fall between the two tests in [Cobos Carrascosa et al \(2016\)](#): we do not introduce additional noise into our data, other than what the pipeline produces (which is in the order of 7×10^{-5}), however we do have asymmetry in some profiles. We find that the error introduced by the on-board processing pipeline is comparable to the error introduced by the inverter implementation when tested on synthetic data. The differences between the results on the different data sets indicate that the error resulting from the RTE inversion is dominated by the noise level of the data and the input data profiles, which is an inherent property of RTE inversions. This underlines the fact, that the accuracy of the pipeline can be judged best by comparing the Stokes parameters, instead of the results of the RTE inversion. In order to find a context for the errors of the physical parameters, retrieved by the RTE inversion, we must compare them to a very similar data set.

[Borrero et al \(2014\)](#) compared different Milne-Eddington inversions using data from a sunspot simulation described in [Rempel \(2012\)](#), which is very

similar to what we used in this work. They synthesised two absorption lines (Fe I 630.15 nm and Fe I 630.24 nm) over the whole MHD cube, and did not introduce any noise. The lines were sampled at 100 wavelength steps, 10 mÅ apart. This sampling provides more information than is available in SO/PHI observations, consequently a better reconstruction of the absorption lines is expected. Moreover, the authors cropped the field of view, such that it contains a comparable number of pixels in umbra, penumbra, and solar granulation (more precisely a ~ 200 G plage region surrounding the spot). Using this data set, they found that different Milne-Eddington inversion codes, executed on one input, produce values within an interval of 35 G, 1.2° and 10 m/s for $|\vec{B}|$, γ , and v_{LOS} , respectively. In this work, we find that the differences in the SO/PHI inversion of on-board and on ground reduced data are smaller than the differences introduced by different inversion codes executed on a single input in Borrero et al (2014). This means, that the accuracy of the on-board processing is higher than the accuracy of a typical Milne-Eddington inversion. We furthermore note, that Borrero et al (2014) does not discuss the errors introduced by Milne-Eddington inversions due to simplifications in the physics underlying this model. We expect these to be considerably larger than the numerical uncertainties between different Milne-Eddington codes (see Orozco Suárez et al, 2010; Castellanos Durán, 2022, and references therein).

In conclusion, the SO/PHI pipeline provides the necessary accuracy to process spectropolarimetric data with 10^{-3} polarimetric sensitivity. We show, that the data processing pipeline does not compromise the accuracy of the inversion results, since it preserves the confidence interval of the Milne-Eddington RTE inversions. Comparing the results of this study with others shows that the effect of the on-board pipeline errors on the RTE inversion is below the errors produced by the RTE inversion inherently on both simulated and observed data. In this paper, we analyse the errors introduced by the on-board data processing pipeline in comparison to on-ground processing. However, these errors can be regarded as negligible, or at least small compared with other sources of error.

Similar processing accuracy can be expected in other on-board data processing pipelines as well (e.g., calculating flat fields, or determining polarimetric ad-hoc cross-talk correction terms), since the restrictions and solutions presented here overarch all on-board implementations. This study also shows, that while requiring a significant effort, on-board reduction of solar spectropolarimetric data is a viable option for future instruments, even with stringent limitations in computational resources. It significantly reduces telemetry requirements for SO/PHI (from 24 images, to five at most, in addition to obviating the necessity to download dark and flat fields), and will be particularly valuable for spectropolarimeters on deep-space missions.

Acknowledgments. We thank Amanda Romero Avila and Philipp Löschl for their work in synthesising the MHD cube for the input data. We use data provided by M. Rempel at the National Center for Atmospheric Research (NCAR). This work was carried out in the framework of the International Max

Planck Research School (IMPRS) for Solar System Science at the Technical University of Braunschweig and the University of Göttingen. Solar Orbiter is a mission led by the European Space Agency (ESA) with contribution from National Aeronautics and Space Administration (NASA). The SO/PHI instrument is supported by the German Aerospace Center (DLR) through grants 50 OT 1201 and 50 OT 1901. The Spanish contribution is funded by AEI/MCIN/10.13039/501100011033/ (RTI2018-096886-C5, PID2021-125325OB-C5, PCI2022-135009-2) and ERDF “A way of making Europe”; “Center of Excellence Severo Ochoa” awards to IAA-CSIC (SEV-2017-0709, CEX2021-001131-S); and a Ramón y Cajal fellowship awarded to DOS. JSCD was also funded by the Deutscher Akademischer Austauschdienst (DAAD). The National Center for Atmospheric Research is sponsored by the National Science Foundation.

- Funding: German Aerospace Center (DLR), 50 OT 1201, 50 OT 1901; Spanish Research Agency, ESP2016-77548-C5, RTI2018-096886-B-C5; European FEDER funds; State Agency for Research of the Spanish MCIU, SEV-2017-0709; Deutscher Akademischer Austauschdienst (DAAD).
- Conflict of interest/Competing interests: Not applicable
- Ethics approval: Not applicable
- Consent to participate: Not applicable
- Consent for publication: Not applicable
- Availability of data and materials
- Code availability: Not applicable
- Authors’ contributions: K.A. did the analysis, wrote the main manuscript text and prepared the figures. J.H., J.S.C.D. and D.O.S. provided advice during the analysis. All authors reviewed the manuscript.

If any of the sections are not relevant to your manuscript, please include the heading and write ‘Not applicable’ for that section.

References

- Albert K, Hirzberger J, Busse D, et al (2019) Performance Analysis of the SO/PHI Software Framework for On-board Data Reduction. In: *Astronomical Data Analysis Software and Systems XXVIII*, Astronomical Society of the Pacific
- Albert K, Hirzberger J, Kolleck M, et al (2020) Autonomous on-board data processing and instrument calibration software for the Polarimetric and Helioseismic Imager on-board the Solar Orbiter mission. *Journal of Astronomical Telescopes, Instruments, and Systems* 6:048004. <https://doi.org/10.1117/1.JATIS.6.4.048004>
- Borrero JM, Lites BW, Lagg A, et al (2014) Comparison of inversion codes for polarized line formation in mhd simulations - i. milne-eddington codes. *Astronomy and Astrophysics* 572:A54. <https://doi.org/10.1051/0004-6361/>

- 24 *Accuracy analysis of the on-board data reduction pipeline for SO/PHI*
201424584, URL <https://doi.org/10.1051/0004-6361/201424584>
- Castellanos Durán JS (2022) Strong magnetic fields and unusual flows in sunspots. PhD thesis, University of Göttingen
- Cobos Carrascosa JP, Aparicio del Moral B, Ramos Mas JL, et al (2016) The RTE inversion on FPGA aboard the solar orbiter PHI instrument. In: Chiozzi G, Guzman JC (eds) Software and Cyberinfrastructure for Astronomy IV, p 991342, <https://doi.org/10.1117/12.2232332>
- Danilovic S, van Noort M, Rempel M (2016) Internetwork magnetic field as revealed by two-dimensional inversions. *Astronomy and Astrophysics* 593:A93. <https://doi.org/10.1051/0004-6361/201527842>, <https://arxiv.org/abs/arXiv:1607.00772> [astro-ph.SR]
- Dominguez-Tagle C, Appourchaux T, Fourmond JJ, et al (2014) Filtergraph Calibration for the Polarimetric and Helioseismic Imager. *Transactions of the Japan Society of Aeronautical and Space Sciences, Aerospace Technology Japan* 12:25–27. <https://doi.org/10.2322/tastj.12.Tk.25>
- Evershed J (1909) Radial movement in sun-spots. *The Observatory* 32:291–292
- Frutiger C, Solanki SK, Fligge M, et al (2000) Properties of the solar granulation obtained from the inversion of low spatial resolution spectra. *Astronomy and Astrophysics* 358:1109–1121
- Lagg A, Woch J, Krupp N, et al (2004) Retrieval of the full magnetic vector with the He I multiplet at 1083 nm. Maps of an emerging flux region. *Astronomy and Astrophysics* 414:1109–1120. <https://doi.org/10.1051/0004-6361:20031643>
- Landi Degl’Innocenti E, Landolfi M (2004) Polarization in Spectral Lines, vol 307. Kluwer Academic, <https://doi.org/10.1007/978-1-4020-2415-3>
- Lange T, Fiethe B, Michel H, et al (2017) On-board processing using reconfigurable hardware on the solar orbiter phi instrument. In: 2017 NASA/ESA Conference on Adaptive Hardware and Systems (AHS), pp 186–191, <https://doi.org/10.1109/AHS.2017.8046377>
- Müller D, St. Cyr OC, Zouganelis I, et al (2020) The Solar Orbiter mission. Science overview. *Astronomy and Astrophysics* 642:A1. <https://doi.org/10.1051/0004-6361/202038467>, <https://arxiv.org/abs/arXiv:2009.00861> [astro-ph.SR]
- Orozco Suárez D, Bellot Rubio LR (2012) Analysis of Quiet-Sun Internetwork Magnetic Fields Based on Linear Polarization Signals. *The Astrophysical Journal* 751(1):2. <https://doi.org/10.1088/0004-637X/751/1/>

2, <https://arxiv.org/abs/arXiv:1203.1440> [astro-ph.SR]

- Orozco Suárez D, Bellot Rubio LR, Vögler A, et al (2010) Applicability of Milne-Eddington inversions to high spatial resolution observations of the quiet Sun. *Astronomy and Astrophysics* 518:A2. <https://doi.org/10.1051/0004-6361/200913421>, <https://arxiv.org/abs/arXiv:1005.5012> [astro-ph.SR]
- Rees DE, Semel MD (1979) Line formation in an unresolved magnetic element: a test of the centre of gravity method. *Astronomy and Astrophysics* 74(1):1–5
- Rempel M (2012) Numerical Sunspot Models: Robustness of Photospheric Velocity and Magnetic Field Structure. *The Astrophysical Journal* 750(1):62. <https://doi.org/10.1088/0004-637X/750/1/62>, <https://arxiv.org/abs/arXiv:1203.0534> [astro-ph.SR]
- Rempel M (2015) Numerical Simulations of Sunspot Decay: On the Penumbra-Evershed Flow-Moat Flow Connection. *The Astrophysical Journal* 814(2):125. <https://doi.org/10.1088/0004-637X/814/2/125>, <https://arxiv.org/abs/arXiv:1511.01410> [astro-ph.SR]
- Sanchez Almeida J, Lites BW (1992) Observation and Interpretation of the Asymmetric Stokes Q, U, and V Line Profiles in Sunspots. *The Astrophysical Journal* 398:359. <https://doi.org/10.1086/171861>
- Schlichenmaier R, Collados M (2002) Spectropolarimetry in a sunspot penumbra. Spatial dependence of Stokes asymmetries in Fe I 1564.8 nm. *Astronomy and Astrophysics* 381:668–682. <https://doi.org/10.1051/0004-6361:20011459>
- Schou J, Scherrer PH, Bush RI, et al (2012) Design and Ground Calibration of the Helioseismic and Magnetic Imager (HMI) Instrument on the Solar Dynamics Observatory (SDO). *Solar Physics* 275:229–259. <https://doi.org/10.1007/s11207-011-9842-2>
- Semel M (1967) Contribution à l'étude des champs magnétiques dans les régions actives solaires. *Annales d'Astrophysique* 30:513–513
- Solanki SK (1987) The Photospheric Layers of Solar Magnetic Flux Tubes. PhD thesis, ETH, Zürich
- Solanki SK (1993) Smallscale Solar Magnetic Fields - an Overview. *Space Science Reviews* volume 63(1-2):1–188. <https://doi.org/10.1007/BF00749277>
- Solanki SK, del Toro Iniesta JC, Woch J, et al (2020) The Polarimetric and Helioseismic Imager on Solar Orbiter. *Astronomy and Astrophysics* 642:A11. <https://doi.org/10.1051/0004-6361/201935325>, <https://arxiv.org/abs/arXiv:1903.11061> [astro-ph.IM]