Probabilistic Zeeman-Doppler imaging of stellar magnetic fields

I. Analysis of τ Scorpii in the weak-field limit

J.R. Andersson¹, O. Kochukhov², Z. Zhao³, and J. Sjölund¹

¹ Department of Information Technology, Uppsala University, Box 337, SE-75120 Uppsala, Sweden e-mail: jennifer.andersson@it.uu.se

² Department of Physics and Astronomy, Uppsala University, Box 516, SE-75120 Uppsala, Sweden

³ Department of Computer and Information Science, Linköping University, SE-58183 Linköping, Sweden

Received ; accepted

ABSTRACT

Context. Zeeman-Doppler imaging (ZDI) is used to study the surface magnetic field topology of stars, based on high-resolution spectropolarimetric time series observations. Multiple ZDI inversions have been conducted for the early B-type star τ Sco, which has been found to exhibit a weak but complex non-dipolar surface magnetic field.

Aims. The classical ZDI framework suffers from a significant limitation in that it provides little to no reliable uncertainty quantification for the reconstructed magnetic field maps, with essentially all published results being confined to point estimates. To fill this gap, we propose a Bayesian framework for probabilistic ZDI. Here, the proposed framework is demonstrated on τ Sco in the weak-field limit. *Methods.* We propose three distinct statistical models, and use archival ESPaDOnS high-resolution Stokes V observations to carry out the probabilistic magnetic inversion in closed form. The surface magnetic field is parameterised by a high-dimensional spherical-harmonic expansion.

Results. We provide mean magnetic field distributions along with the corresponding surface uncertainty maps for τ Sco. By comparing three different prior distributions over the latent variables in the spherical-harmonic decomposition, our results showcase the ZDI sensitivity to various hyperparameters. The mean magnetic field maps are qualitatively similar to previously published point estimates, but analysis of the magnetic energy distribution indicates high uncertainty and higher energy content at low angular degrees *l*. *Conclusions.* Our results effectively demonstrate that, for stars in the weak-field regime, reliable uncertainty quantification of recovered magnetic field maps can be obtained in closed form with natural assumptions on the statistical model. Future work will explore

magnetic field maps can be obtained in closed form with natural assumptions on the statistical model. Future work will explore extending this framework beyond the weak-field approximation and incorporating prior uncertainty over multiple stellar parameters in more complex magnetic inversion problems.

Key words. Methods: numerical – Techniques: polarimetric – Stars: activity – Stars: magnetic field – Stars: individual: τ Sco

1. Introduction

Magnetic fields are understood to play a significant role in the formation and evolution of stars. For example, magnetic interlocking of young stars with their accretion disks and winds is the main mechanism behind stellar rotational braking (Shu et al. 1994; Bouvier et al. 2014), and is therefore a key component of the early stellar evolution. At all stellar ages, a multitude of time-dependent surface phenomena and processes – spots, flares, chromospheric emission – is believed to be driven by magnetic fields (Reiners 2012). Extension of the surface field into the circumstellar environment – the stellar magnetosphere – shapes stellar winds, thereby governing the stellar angular momentum evolution (Weber & Davis 1967; Matt & Pudritz 2008; Vidotto et al. 2009), and strongly impacts the interaction of stars and nearby planets (Lanza 2018; Strugarek 2016; Vidotto 2020).

Two basic types of magnetic fields are known to exist in stars (e.g. Donati & Landstreet 2009). On the one hand, all cool, lowmass stars, including the Sun and similar objects, continuously generate magnetic fields in their interiors through the so-called dynamo process. The resulting fields emerge at stellar surfaces in the form of evolving structures with moderate average strengths (10–100 G) and spatial scales ranging from a tiny fraction of the stellar surface (small-scale fields) all the way to the scales comparable to stellar radii (large-scale or global fields). On the other hand, the majority of hot, intermediate-mass and massive stars are not immediately recognised as magnetic. Only about 10% or less, depending on the stellar mass, of these objects possess detectable magnetic fields (Grunhut et al. 2017; Sikora et al. 2019). Characteristics of these fields are drastically different from those in cool stars. Magnetic fields of hot stars are topologically simple and strong (~ kG), lack significant small-scale components, and are constant on the observable time scales.

Irrespective of the type of stellar magnetic field, information about its strength and surface geometry represents a critical stellar parameter highly sought after by theoretical and observational studies alike. Direct detection of stellar magnetic fields typically relies on the Zeeman effect in spectral lines (e.g. Landi Degl'Innocenti & Landolfi 2004). In the presence of a field, atomic energy levels split into magnetic sub-levels. This produces the corresponding splitting of spectral lines. Additionally, lines become polarised, with the direction of polarisation dependent on the field vector orientation and the amplitude of the polarisation signal linked to the field strength. This Zeeman-induced polarisation is particularly useful for detecting and characterising stellar magnetic fields.

Polarisation observations of stars are typically carried out using a combination of a high-resolution spectrograph and a circular polarimeter (e.g. Donati et al. 1997). Such instrumentation can measure the total intensity spectrum (Stokes *I* parameter) together with the circular polarisation spectrum (Stokes *V* parameter). The latter is sensitive to the line of sight component of the magnetic field. A complementary measurement of linear polarisation spectra (Stokes *Q* and *U* parameters), which provide information on the transverse field component, is also possible, but rarely accomplished due to a factor of 10 weaker linear polarisation signals (Wade et al. 2000; Kochukhov et al. 2011; Rosén et al. 2015).

Quantitative interpretation of the observed polarisation spectra of stars other than the Sun faces a significant challenge due to the unresolved nature of these observations. Polarisation signatures recorded by a distant observer represent an average of Stokes V profiles originating from different locations on the visible stellar hemisphere, which makes it impossible to infer the magnetic field geometry from a single observation. Considerably more information is contained in a time sequence of polarisation observations of rotating magnetic stars. First, the rotational Doppler shift redistributes signals from different parts of the stellar surface within a spectral line profile, allowing one to relate the position in the line profile to the longitude on the stellar surface. Effectively, this makes a single line profile observation equivalent to a one-dimensional projection of the stellar surface map (Vogt et al. 1987). Second, the spectral line profile changes due to stellar rotation as different parts of the stellar surface come in and out of view. Combining these two effects in the inversion technique known as Doppler Imaging (DI) offers the possibility to reconstruct two-dimensional maps of the surfaces of unresolved distant stars by fitting a set of spectral line profiles recorded at different stellar rotation phases (e.g. Kochukhov 2016).

DI was initially applied to map spots on different kinds of rapidly rotating stars with inhomogeneous surfaces using intensity spectra (Khokhlova et al. 1986; Vogt et al. 1999). The same inversion principles were subsequently extended to the problem of reconstructing the magnetic geometry of both cool stars with complex fields (Brown et al. 1991; Donati 1999) and hot stars with simple field topologies (Piskunov & Kochukhov 2002; Kochukhov et al. 2002). These early Zeeman-Doppler imaging (ZDI) studies approximated magnetic field distributions as a set of three independent images corresponding to the three magnetic field vector components, and employed either maximum entropy or Tikhonov regularisation to ensure convergence to a stable and unique solution. The majority of more recent ZDI applications use the spherical-harmonic decomposition to describe the stellar surface magnetic field (Donati et al. 2006b; Kochukhov et al. 2014; Folsom et al. 2018). This approach ensures that the reconstructed magnetic field is solenoidal and provides a convenient framework to characterise poloidal/toroidal as well as axisymmetric/nonaxisymmetric magnetic field components as a function of the spatial scale (angular degree *l*). In this case, spherical-harmonic coefficients, rather than the field vector component values at particular coordinates on the stellar surface, represent the free parameters of the inversion problem. Unless the stellar field is restricted to a very simple configuration (e.g. a pure dipole with only l = 1 harmonics), regularisation is still required for stable recovery, due to the ill-posedness of the inverse problem.

Regardless of the specific implementation of ZDI, existing magnetic inversion methods suffer from one major shortcoming: it is difficult or outright impossible to provide realistic uncertainties of the derived magnetic field maps. Essentially all published mapping results represent point estimates, without an attempt to explore the full range of solutions compatible with observations. Due to the high dimensionality of the spectropolarimetric inversion problem and the complexity and computational intensity of the synthetic forward model, obtaining a probabilistic solution is challenging. Occasionally, uncertainties are assessed empirically by comparing magnetic field distributions computed from independent line profile observations (Kochukhov et al. 2019, 2022). Some ZDI studies also considered formal uncertainties obtained from the diagonal of a Hessian matrix (Piskunov & Kochukhov 2002). The former approach is, however, feasible only for strongly magnetic stars, and the latter yields biased uncertainties since parameter correlations are not considered. Additionally, several important stellar parameters (including rotational period, projected rotational velocity and inclination of the rotational axis) enter ZDI calculations. Each of these parameters is known to a limited precision, but their uncertainties are rarely propagated to magnetic maps (Petit et al. 2008). All these problems make it difficult to judge the reliability of published ZDI results, hindering their interpretation and wide utilisation.

In this paper, we address the problem of deriving realistic ZDI uncertainties by presenting a fully Bayesian framework for probabilistic ZDI. Previously, similar methods have only been applied to determine properties of stellar magnetic fields from circular spectropolarimetric observations under an oblique dipole model assumption (Petit & Wade 2012), which is a restrictive assumption not suitable for many magnetic stars. Here, we present a general framework for probabilistic ZDI, using a probabilistic model based on a general spherical-harmonic field parameterisation. Compared to the dipolar modelling approach, our probabilistic ZDI framework applies to arbitrarily complex magnetic field geometries. To present the proposed framework, we start with the comparatively simple problem of probabilistic ZDI reconstruction for a target star with fixed nuisance parameters and a weak magnetic field, which ensures a linear relation between the free parameters describing the magnetic map and the modelled polarisation signal. More specifically, we apply the Bayesian framework for uncertainty quantification of the surface magnetic field of the star τ Sco (HD 149438, HR 6165). In a future paper, we will extend our approach to stars exhibiting stronger magnetic fields, corresponding to a non-linear line profile response, and incorporate a Bayesian treatment of input stellar parameters.

Our target, τ Sco, is a bright, massive star with spectral type B0.2V. This star is located in the ~ 10 Myr-old (Feiden 2016) Upper Sco stellar association. This star exhibits unusual characteristics, such as a slow rotation, a younger age than the surrounding stellar population, and a nitrogen excess (e.g. Przybilla et al. 2010; Nieva & Przybilla 2014), which distinguishes it from the other massive stars in Upper Sco. Even more remarkably, Donati et al. (2006b) demonstrated through spectropolarimetric observations and ZDI modelling the presence of a relatively weak $(\leq 600 \text{ G})$ and unusually complex magnetic field at the surface of that star. Unlike the majority of hot, magnetic stars, which possess a roughly dipolar (i.e. spherical-harmonic modes with l = 1) field geometries (Shultz et al. 2018), τ Sco was found to exhibit a surface magnetic structure containing significant power in the harmonic modes of up to l = 5, with the largest contribution coming from $l \in [3, 4]$.

These unusual chemical, rotational, and magnetic characteristics of τ Sco spurred a series of studies to explain this star as a product of a stellar merger (Nieva & Przybilla 2014; Schneider et al. 2016; Keszthelyi et al. 2021). In particular, Schneider et al. (2019) presented 3D magnetohydrodynamical simulations tailored to τ Sco that demonstrated generation of a powerful magnetic field during a stellar merger event. This field decays quickly after the merger, settling into the stable non-axisymmetric configuration observed today (Schneider et al. 2020; Braithwaite 2008).

So far, τ Sco remains the only well-established magnetic star that is likely a merger product. Therefore, it plays the role of a key benchmark object for investigations into both binarity and magnetism of massive stars. At the same time, a follow-up ZDI analysis of τ Sco by Kochukhov & Wade (2016), using an extended observational dataset, suggested that the outcome of magnetic inversions based on circular polarisation spectra strongly depends on the assumptions of the magnetic field parameterisation. This non-uniqueness problem may arise in various applications of ZDI (e.g. Donati & Brown 1997; Kochukhov et al. 2002), but is particularly significant for τ Sco due to its field complexity and small amount of rotational Doppler broadening. Consequently, Kochukhov & Wade (2016) obtained vastly different magnetic field geometries and strengths with equally good fits to available observations of τ Sco by using different, but equally plausible, harmonic parameterisations. These issues, and the significance of τ Sco for wider massive-star research, position this star as a particularly interesting target for probabilistic ZDI.

This paper is organised in the following way: in Sect. 2, the spectropolarimetric observations of τ Sco used for probabilistic spectropolarimetric inversion are introduced. Sect. 3 presents a detailed description of the proposed Bayesian framework for probabilistic ZDI. We propose several probabilistic models, some of which aim to incorporate uncertainties in design parameters corresponding to hyperparameters in the regularisation function often used in conventional ZDI. Then, in Sect. 4, we present summary statistics of the magnetic field map distributions obtained using our framework for probabilistic ZDI. Finally, in Sect. 5, we discuss and summarise the results and conclusions.

2. Observational data and stellar parameters

We use the same observations of τ Sco as analysed by Kochukhov & Wade (2016). The spectropolarimetric data available for this star comprises 49 circular polarisation observations collected in 2004–2009 with the ESPaDOnS instrument (Donati et al. 2006a) at the 3.6m Canada-France-Hawaii telescope. The spectra, publicly available in the PolarBase archive¹ (Petit et al. 2014), were processed with the least-squares deconvolution (LSD, Kochukhov et al. 2010) method, yielding high-quality mean Stokes *I* and *V* profiles.

We refer the readers to Kochukhov & Wade (2016) for further details on the data reduction and pre-processing. Following that study, as well as the earlier work by Donati et al. (2006b), we adopt a rotational period of $P_{\rm rot} = 41.033$ d, a projected rotational velocity of $v_{\rm e} \sin i = 6$ km s⁻¹ and an inclination angle of $i = 70^{\circ}$ for the ZDI analysis, without accounting for possible errors in these parameters.

The observational data used in our work includes uncertainties for each velocity bin in the Stokes *I* and *V* parameter profiles. For data with high signal-to-noise ratio, as considered here, the errors are uncorrelated and follow a Gaussian distribution. For a detailed discussion of how these errors are derived by the data reduction pipelines from raw spectropolarimetric exposures and then propagated through the derivation of LSD profiles, we refer the readers to Donati et al. (1997) and Kochukhov et al. (2010).

3. Methods

3.1. Magnetic field map parameterisation

There are multiple ways of parameterising the magnetic field distribution across the surface of a target star. Earlier ZDI in-

versions have adopted direct parameterisation techniques (e.g. Brown et al. 1991; Kochukhov et al. 2002), in which each magnetic field vector component map is represented by the nodes of a discrete surface grid, and fitted to the data independently. Since then, alternative parameterisations of the surface magnetic field, employing spherical-harmonic expansions, have often been preferred (e.g. Donati et al. 2006b; Kochukhov et al. 2014). In contrast to the direct approach, this magnetic field parameterisation allows us to impose restrictions on the magnetic geometry to follow Maxwell's equations. Specifically, excluding the l = 0 harmonic component ensures no net signed magnetic flux through the closed stellar surface. There are several other appealing properties of the spherical-harmonic field parameterisation, see e.g. the discussion by Kochukhov & Wade (2016).

In this work, we use the spherical-harmonic representation of the surface magnetic field defined as follows:

$$B_{r}(\theta,\phi) = -\sum_{l=1}^{l_{\max}} \sum_{m=-l}^{l} \alpha_{l,m} Y_{l,m}(\theta,\phi),$$

$$B_{\theta}(\theta,\phi) = -\sum_{l=1}^{l_{\max}} \sum_{m=-l}^{l} \beta_{l,m} Z_{l,m}(\theta,\phi) + \gamma_{l,m} X_{l,m}(\theta,\phi), \qquad (1)$$

$$B_{\phi}(\theta,\phi) = -\sum_{l=1}^{l_{\max}} \sum_{m=-l}^{l} \beta_{l,m} X_{l,m}(\theta,\phi) - \gamma_{l,m} Z_{l,m}(\theta,\phi).$$

Here, $B_r(\theta, \phi), B_{\theta}(\theta, \phi)$, and $B_{\phi}(\theta, \phi)$ denote the surface magnetic field components in the radial, meridional, and azimuthal direction, respectively, with θ and ϕ corresponding to the surface colatitude and longitude. The spherical-harmonic functions and their derivatives for each angular degree l and azimuthal order *m* are denoted by $Y_{l,m}(\theta, \phi), Z_{l,m}(\theta, \phi)$ and $X_{l,m}(\theta, \phi)$. Finally, the amplitudes of the spherical-harmonic modes, $\alpha_{l,m}$, $\beta_{l,m}$ and $\gamma_{l,m}$, delineate characteristics of the poloidal and toroidal magnetic field components. Together, these coefficients represent the free parameters in the spectropolarimetric inversion problem. See Kochukhov et al. (2014) for a more detailed description of the field parameterisation and relevant notations. One difference compared to Kochukhov et al. (2014) is that the constant $C_{l,m}$, defined by their Eq. (7), was modified² here to ensure that equal values of the harmonic coefficients yield equal total magnetic energy, independent of l and m.

3.2. Forward line profile model

In this study, we restrict our attention to stars for which the weak-field assumption (Landi Degl'Innocenti & Landolfi 2004) is a good approximation for the forward Stokes I and V profile model. This approach is typically valid for sub-kG fields studied with spectral lines at optical wavelengths and is applicable to a large group of stars, including solar-like stars, the majority of low-mass stars, and a few weak-field hot stars such as τ Sco.

In practice, we follow Petit & Wade (2012) in representing the local Stokes *I* profile using a Gaussian function. Its strength (d = 0.115), position (radial velocity shift $v_0 = -0.48 \text{ km s}^{-1}$), and width (FWHM = 14.3 km s⁻¹ corresponding to $\sigma = 6.1 \text{ km s}^{-1}$) are adjusted to fit the observed LSD Stokes *I* spectrum. The adopted Gaussian line width includes the effect of intrinsic stellar line broadening as well as the smearing by the instrumental profile of ESPaDOnS ($R \approx 68000$).

http://polarbase.irap.omp.eu/

² This was accomplished by multiplying $C_{l,m}$ by $\sqrt{(l+1)/l}$ for m = 0 and by $\sqrt{2(l+1)/l}$ for $m \neq 0$ modes.

The local Stokes *V* profile is derived from the velocity derivative of the Gaussian function

$$V(v) = -1.4 \times 10^{-3} \lambda_0 g_{\text{eff}} B_{\parallel} \frac{\partial I}{\partial v}$$

= $1.4 \times 10^{-3} \lambda_0 g_{\text{eff}} B_{\parallel} \frac{d(v - v_0)}{\sigma^2} \exp\left[-\frac{(v - v_0)^2}{2\sigma^2}\right].$ (2)

Here, the central wavelength adopted for the LSD profile is $\lambda_0 = 500$ nm and the effective Landé factor is $g_{\text{eff}} = 1.2$ (see Kochukhov & Wade 2016). The quantity B_{\parallel} denotes the line of sight (longitudinal) component of the magnetic field (in kG) at the θ, ϕ location on the stellar surface, as seen by the observer at the rotational phase *t*, defined over the interval [0, 1). This magnetic field projection can be calculated with

$$B_{\parallel} = \left[\cos\theta\cos i + \sin\theta\sin i\cos\left(\phi + 2\pi t\right)\right]B_{r} -\left[\sin\theta\cos i - \cos\theta\sin i\cos\left(\phi + 2\pi t\right)\right]B_{\theta}$$
(3)
$$-\sin i\sin\left(\phi + 2\pi t\right)B_{\phi}.$$

To obtain the disc-integrated Stokes *I* and *V* profiles, we divide the stellar surface into $\approx 10^4$ elements, pre-calculate the areas *S* of these elements, and evaluate the magnetic field using Eq. (1). Then, for each rotational phase *t*, we determine the cosine μ of the angle between the observer's line of sight and the surface normal (the limb angle) using

$$\mu = \cos\theta \cos i + \sin\theta \sin i \cos\left(\phi + 2\pi t\right) \tag{4}$$

and locate visible surface elements with the condition $\mu > 0$. Eq. (3) is then applied to determine B_{\parallel} for visible surface elements. Following this, the local Stokes *I* and *V* spectra are calculated as described above. This calculation is carried out on the velocity grid of observations, v_{obs} , Doppler-shifted according to the local component of the projected rotational velocity, $v_e \sin i$,

$$v = v_{obs} - v_e \sin i \sin \theta \sin (\phi + 2\pi t).$$
⁽⁵⁾

Finally, we add together all profiles from the visible stellar hemisphere using the product of the projected surface area and a linear continuum limb-darkening function with the coefficient u = 0.3as an integration weight W

$$W = S\mu(1 - u + u\mu) \tag{6}$$

Due to the weak-field assumption expressed by Eq. (2), this sequence of operations represents a linear transformation from a set of the spherical harmonic coefficients, α , β and γ , to the phase-dependent disk-integrated Stokes *V* profiles. We take advantage of this linearity of the problem in this paper. At the same time, we anticipate a generalisation to the strong-field situation in future studies by replacing Eq. (2) with another relation, not limited to weak fields.

While we have used the observed Stokes I LSD profiles to adjust the Gaussian function parameters and determine a radial velocity offset, the magnetic field does not impact Stokes I in the weak-field limit. Consequently, the classical and probabilistic ZDI analyses presented in this paper consider only the Stokes V observations.

3.3. Standard ZDI

We define the vector $\mathbf{z} = (\alpha_{l,m}, \beta_{l,m}, \gamma_{l,m})$, where $m \in \{-l, -l + 1, \ldots, l - 1, l\}$ for each $l \in \{1, 2, \ldots, l_{\max}\}$, as a collection of the amplitudes of the spherical-harmonic modes, $\alpha_{l,m}, \beta_{l,m}$, and $\gamma_{l,m}$. We truncate the expansion by $l_{\max} = 10$, corresponding to 360

free parameters in the resulting magnetic field parameterisation. In a nutshell, standard ZDI amounts to solving the weighted least-squares (LS) problem

$$\hat{\mathbf{z}} = \arg\min_{\mathbf{z}} \left(\|\Lambda^{\frac{1}{2}}(\mathbf{y} - \mathbf{f}(\mathbf{z}))\|_{2}^{2} + \eta r(\mathbf{z}) \right), \tag{7}$$

with respect to the spherical-harmonic coefficients \mathbf{z} . In the above formulation, $\mathbf{f}(\mathbf{z})$ denotes the forward mapping function (corresponding to the computations described in Sect. 3.2) between \mathbf{z} and the observed Stokes *V* spectra denoted \mathbf{y} . Moreover, Λ denotes the inverse covariance matrix of the measurement noise (i.e. the precision matrix). Another important component is the regulariser

$$r(\mathbf{z}) = \sum_{l=1}^{l_{\max}} \sum_{m=-l}^{l} l^2 (\alpha_{l,m}^2 + \beta_{l,m}^2 + \gamma_{l,m}^2),$$
(8)

with regularisation strength η . The regulariser ensures that, despite the ill-posedness of the Stokes *V*-only ZDI inversion (Piskunov & Kochukhov 2002), a unique solution exists, that avoids the introduction of higher-order modes unless the added flexibility is motivated by the data. The regularisation structure presented in Eq. (8), where the regularisation objective is to minimize the l^2 -weighted magnetic energies of the spherical-harmonic coefficients, is commonly used in some ZDI studies (e.g. Hussain et al. 2000; Kochukhov et al. 2014; Kochukhov 2016) while others adopt the *l*-weight for this regularisation (Morin et al. 2008, 2010).

The first part of the optimisation objective in Eq. (7) is hereon referred to as the weighted LS objective, and the second part is referred to as the regularisation objective. Since η determines the trade-off between these two optimisation objectives, the choice of η has a significant impact on the properties of the solution. Larger η may result in information loss by favouring the generation of smoother magnetic field maps. On the other hand, small η may introduce spurious high-order modes, resulting in magnetic field maps containing unjustified surface details. To choose the regularisation parameter, we follow the empirical approach of choosing η such that the quotient between the weighted LS objective value and the regularisation objective value at the optimum is between 2 and 4. To briefly summarise the underlying motivation, this region defines a breaking point after which the fit quality improves slowly as the regularisation strength decreases. This approach to choosing the regularisation parameter is discussed in more detail by Kochukhov (2017).

This classical ZDI framework views the model parameters \mathbf{z} as fixed quantities to infer. To express uncertainty about the resulting point estimates $\hat{\mathbf{z}}$ of these quantities, the distribution of possible datasets can be considered. However, data collection in the context of ZDI is very costly and stellar magnetic fields often evolve between different observing runs, rendering this approach impractical. Moreover, due to the ill-posedness of the problem, formal error bars on $\hat{\mathbf{z}}$ become strongly dependent on the regularisation strength η . For these reasons, and due to computational limitations of previous ZDI approaches, such error bars have generally not been presented in previous ZDI studies. In the following section, we describe a procedure for quantifying uncertainty in the parameters \mathbf{z} according to the Bayesian viewpoint.

3.4. Probabilistic ZDI

In this section, we propose several statistical models to extend the standard ZDI framework, discussed in Sect. 3.3, into a fully

Bayesian setting. The Bayesian setting (e.g. Gelman et al. 2020) provides a framework for updating prior beliefs about the uncertainty in the spherical-harmonic coefficients, represented as a probability distribution over z, as new data becomes available. To this end, we treat the spherical-harmonic coefficients z as latent random variables. Given a set of spectropolarimetric observations y, the resulting posterior distribution, $p(\mathbf{z}|\mathbf{y})$, captures our informed state of knowledge about the parameters \mathbf{z} conditioned on y. Conveniently, this framework brings inherent regularisation into the inversion process; an especially attractive property in light of the ill-posedness of the problem at hand. We begin by constructing a probabilistic model and describing how to calculate the posterior distribution in the context of probabilistic ZDI. At the end of the section, we dive deeper into the choice of prior distribution, aiming to capture uncertainty propagated through its hyperparameters.

3.4.1. Model formulation

To model the likelihood, the observational noise is assumed to be Gaussian with diagonal covariance (see Sect. 2), such that $p(\mathbf{y}|\mathbf{z}) = \mathcal{N}(\mathbf{y}; \mathbf{f}(\mathbf{z}), \Lambda^{-1})^3$. Then, given a prior distribution $p(\mathbf{z})$, the posterior distribution $p(\mathbf{z}|\mathbf{y})$ can be obtained following Bayes' theorem, according to

$$p(\mathbf{z}|\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{z})p(\mathbf{z})}{p(\mathbf{y})}.$$
(9)

Given $p(\mathbf{z}|\mathbf{y})$, the posterior magnetic field distribution can then be obtained by applying the transformation in Eq. (1). In general, exact determination of the posterior distribution according to Eq. (9) is challenging since the marginal likelihood $p(\mathbf{y})$ is often intractable. When the posterior distribution has been inferred, the predictive distribution over the modelled Stokes V profiles \mathbf{y}_* can be obtained by marginalisation:

$$p(\mathbf{y}_*|\mathbf{y}) = \int p(\mathbf{y}_*|\mathbf{z}) p(\mathbf{z}|\mathbf{y}) d\mathbf{z}.$$
 (10)

We model the prior distribution by a Gaussian distribution, $p(\mathbf{z}) = \mathcal{N}(\mathbf{z}; \mathbf{0}, \mathbf{\Omega}^{-1})$, where $\mathbf{\Omega}^{-1}$ is a diagonal covariance matrix. Despite the Gaussian likelihood and Gaussian prior distribution, Eq. (9) is in general intractable if the forward mapping function $\mathbf{f}(\mathbf{z})$ is non-linear.

In the initial analysis, we use the parameters $\Omega_{i,i} = \eta l_i^2$ in the marginal prior distribution for each spherical-harmonic coefficient z_i . While it is difficult to construct a fully physics-informed prior distribution over the spherical-harmonic coefficients, this choice of prior distribution captures our prior belief that spherical-harmonic coefficients of high degree l are less prevalent in the solution. In addition, the prior variance adds implicit regularisation to the solution, and with the proposed likelihood model and prior distribution, the resulting maximum a posteriori probability (MAP) estimate, $\hat{\mathbf{z}}_{MAP} = \arg \max_{\mathbf{z}} \log p(\mathbf{z}|\mathbf{y})$, can be shown to solve the regularised weighted LS problem arising in the standard ZDI formulation given by Eq. (7). In that sense, the proposed probabilistic model allows us to generalise the regularisation structure often used in standard ZDI inversion (see Eq. (8)), with the hyperparameter η in the prior distribution corresponding to the regularisation strength in the standard ZDI formulation. This means that, if the mean of the full posterior distribution coincides with the MAP estimate, the Bayesian formulation results in an uncertainty quantification centred around the point estimate obtained from standard ZDI.

3.4.2. Posterior inference

In the weak-field regime, adopted in this paper, the forward model can be approximated by a function linear in the parameters \mathbf{z} such that $\mathbf{f}(\mathbf{z}) = \mathbf{A}\mathbf{z}$, where **A** is the transformation matrix. This linearity follows from the set of equations presented in Sect. 3.2. In this case, the Gaussian prior in our probabilistic model is conjugate to the likelihood, and a closed-form expression for the posterior distribution exists (e.g. Murphy 2022; Bishop 2007; Sjölund et al. 2018). In fact, following the derivation in Murphy (2022, Sect. 3.3.2), it can be shown that the posterior distribution is also a Gaussian $p(\mathbf{z}|\mathbf{y}) = \mathcal{N}(\mathbf{z}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ with $\boldsymbol{\Sigma}^{-1} = \boldsymbol{\Omega} + \mathbf{A}^T \boldsymbol{\Lambda} \mathbf{A}$ and $\mu = \Sigma \mathbf{A}^T \Lambda \mathbf{y}$. Since $p(\mathbf{z}|\mathbf{y})$ is Gaussian and, according to Eq. (1), the magnetic field B is a linear function of the sphericalharmonic coefficients z, the posterior magnetic field distribution is also Gaussian. This distribution can be expressed in closed form according to the linear transformation theorem for the multivariate Gaussian distribution: if $\mathbf{z} \sim p(\mathbf{z}|\mathbf{y}) = \mathcal{N}(\mathbf{z};\mu,\Sigma)$, then $B_r = \tilde{B}_r \mathbf{z} \sim \mathcal{N}(\tilde{B}_r \mu, \tilde{B}_r \Sigma \tilde{B}_r^T)$, $B_\theta = \tilde{B}_\theta \mathbf{z} \sim \mathcal{N}(\tilde{B}_\theta \mu, \tilde{B}_\theta \Sigma \tilde{B}_\theta^T)$ and $B_\phi = \tilde{B}_\phi \mathbf{z} \sim \mathcal{N}(\tilde{B}_\phi \mu, \tilde{B}_\phi \Sigma \tilde{B}_\phi^T)$. Here, $\tilde{B}_r, \tilde{B}_\theta$ and \tilde{B}_ϕ are defined according to the linear transformations in Eq. (1).

In this setting, the predictive distribution, see Eq. (10), is also available in closed form. We use \mathbf{y}_* to denote the vector representing the modelled LSD Stokes V profiles, whereas \mathbf{y} continues to denote the specific observations used to fit the probabilistic model. The corresponding transformation matrix used to make predictions is then denoted by A^* . This transformation matrix can be different from the transformation matrix A used to fit the model, for example if we wish to increase the discretisation of the predicted Stokes V profiles compared to the observations \mathbf{y} . Given the transformation matrix A^* , the predictive distribution over \mathbf{y}_* is then given by $p(\mathbf{y}_*|\mathbf{y}) = \mathcal{N}(\mathbf{y}_*;\mu_*,\Sigma_*)$, where $\mu_* = A^*\mu$ and $\Sigma_* = A^*\Sigma(A^*)^T + \Lambda_*^{-1}$. Here, Λ_*^{-1} is a diagonal matrix containing the estimated observational noise corresponding to y_* . This result follows from Eq. 2.115 in Bishop (2007).

3.4.3. Prior hyperparameter selection

In the statistical model described in the previous subsections, a specific set of hyperparameters $\Omega_{i,i}$ is assumed to parameterise the variance in the prior distribution. In our initial analysis, we use the hyperparameters $\Omega_{i,i} = \eta l_i^2$ with $\eta = 100$. Here, η is determined according to the empirical approach described in Sect. 3.3, motivated by the connection to the regularisation strength in the standard ZDI formulation. However, there is no consensus regarding the choice of regularisation in existing ZDI literature. Other authors favour a regularisation function that, in the probabilistic model formulation, corresponds to the hyperparameters $\Omega_{i,i} = \eta l_i$ (without the square) in the prior variance (Morin et al. 2008). Previous studies also adopt different approaches to choosing the hyperparameter η or use an entirely different form of regularisation altogether (Folsom et al. 2018). Even using a single approach, like the empirical approach described in Sect. 3.3, there is an inherent uncertainty in the choice of η . Since the proposed framework for probabilistic ZDI quantifies uncertainty within the chosen model class, the uncertainty estimate can become more reliable by broadening the class of models under consideration. This can be achieved by modifying the prior distribution. In Sect. 3.4.4-3.4.6, we describe three ways of selecting hyperparameters in the prior distribution; either as a point estimate or by incorporating prior uncertainty over these parameters.

³ The notation $p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mu, \Sigma)$ represents a Gaussian distribution over the random variable \mathbf{x} , with mean vector μ and covariance matrix Σ .

3.4.4. Empirical Bayes

One method for determining hyperparameters in the prior distribution of a probabilistic model is called *empirical Bayes*, or *evidence maximisation* (e.g. Bishop 2007). With this approach, a point estimate $\hat{\eta}$ of the hyperparameter η is determined by maximising the marginal likelihood $p(\mathbf{y}|\eta) = \int p(\mathbf{y}|\mathbf{z})p(\mathbf{z}|\eta)d\mathbf{z}$, i.e.

$$\hat{\eta} = \arg\max\,\log p(\mathbf{y}|\boldsymbol{\eta}). \tag{11}$$

To make the dependence on the hyperparameter η explicit here, $p(\mathbf{z}|\eta)$ denotes the prior distribution over the model parameters \mathbf{z} , parameterised by η . The marginal likelihood quantifies how well the model using the hyperparameter η , with latent variables \mathbf{z} governed by the prior distribution $p(\mathbf{z}|\eta)$, explains the observed data \mathbf{y} . The empirical Bayes hyperparameter estimate can then be viewed as a maximum likelihood estimate given a likelihood function defined over a specified model space, from which the random variables \mathbf{z} have been marginalised out. In relation to the fully Bayesian approach to hyperparameter optimisation, we can view the empirical Bayes approach as an approximation of the posterior distribution $p(\eta|\mathbf{y})$ by a point function at its mode. This point estimate is equivalent to the point estimate obtained assuming an improper prior distribution according to

$$\hat{\eta} = \arg\max_{\eta} p(\eta | \mathbf{y}) = \arg\max_{\eta} p(\mathbf{y} | \eta) p(\eta) = \arg\max_{\eta} p(\mathbf{y} | \eta),$$
(12)

where, in the last step, we let $p(\eta) = 1$. As we can see, the point estimates in Eq. (11) and Eq. (12) are identical.

We determine the hyperparameter η given the probabilistic model described in Sect. 3.4.1 by solving the maximisation problem in Eq. (11). Given the conjugacy between the likelihood and our choice of prior, a closed-form expression for the marginal likelihood exists in this case, namely $p(\mathbf{y}|\eta) = \mathcal{N}(\mathbf{y}; \mathbf{0}, \Lambda^{-1} + \mathbf{A}\mathbf{\Omega}^{-1}(\eta)\mathbf{A}^T)$. This result is derived using Eq. 2.115 in Bishop (2007).

Using a prior distribution with parameters $\Omega_{i,i} = \eta l_i^2$ in the statistical model described in Sect. 3.4.1, the empirical Bayes estimate $\hat{\eta}$ is in the order of 10⁻³, yielding a quotient exceeding a magnitude of 15 between the weighted LS objective value and the regularisation objective value. Fig. 1 shows the log evidence as a function of η near the maximum. In other words, the resulting quotient is far outside the interval [2,4] at the mean of the posterior distribution $p(\mathbf{z}|\mathbf{y})$. The resulting magnetic field strength becomes too large, violating the weak-field assumption on the magnetic field used to model the likelihood. Thus, estimating $\hat{\eta}$ using empirical Bayes gives a magnetic field with properties inconsistent with the underlying physical assumptions even at the mean of the posterior distribution, suggesting that the marginal likelihood, and by extension empirical Bayes, is not useful in this particular context. We conclude that maximum likelihood estimation of this hyperparameter is ill-posed in this case, and we observe that the marginal likelihood is sensitive to small variations in η . Analysing the well-posedness of maximum likelihood estimation in the regression setting is non-trivial (see, e.g., Karvonen & Oates 2024, for a deeper exposition). Consequently, maximum likelihood estimates in such settings may not reliably respect underlying assumptions.

3.4.5. Hierarchical model

In a fully Bayesian setting, we treat hyperparameters in the prior distribution $p(\mathbf{z})$ as latent random variables in an extended proba-



Fig. 1. The unnormalised log evidence as a function of η on a region close to $\eta = 0$. The log evidence is maximised for $\eta \approx 0.002$. As η grows outside of the displayed region, the log evidence continues to decrease.

bilistic model. Let's consider the case in which we treat the prior hyperparameter η , corresponding to the regularisation strength in the standard ZDI framework, as a random variable. Prior knowledge of the distribution over η can then be incorporated into the joint prior distribution, $p(\mathbf{z}, \eta)$. The posterior distribution $p(\mathbf{z}, \eta | \mathbf{y})$ is then

$$p(\mathbf{z}, \eta | \mathbf{y}) = \frac{p(\mathbf{y} | \mathbf{z}) p(\mathbf{z}, \eta)}{p(\mathbf{y})} = \frac{p(\mathbf{y} | \mathbf{z}) p(\mathbf{z} | \eta) p(\eta)}{p(\mathbf{y})}.$$
 (13)

Here, $p(\eta)$ is referred to as the *hyperprior distribution*. Treating hyperparameters of the prior distribution as random variables results in a so-called *hierarchical probabilistic model*. The sought posterior distribution is then given by marginalising over η according to

$$p(\mathbf{z}|\mathbf{y}) = \int p(\mathbf{z}, \eta|\mathbf{y}) d\eta, \qquad (14)$$

which is equivalent to assuming the prior distribution

$$p(\mathbf{z}) = \int p(\mathbf{z}, \eta) d\eta = \int p(\mathbf{z}|\eta) p(\eta) d\eta$$
(15)

in Eq. (9). Defining a suitable prior distribution $p(\eta)$ can be challenging, and this distribution will also be defined in terms of its own set of hyperparameters that need to be chosen. We can extend the probabilistic model with another hierarchical layer by treating the hyperparameters in the hyperprior distribution as random variables, but the hyperparameters in the final hierarchical layer eventually need to be fixed. The limiting distribution, however, often becomes invariant to the specific form of the hyperprior distributions, and the specific hyperparameters in the hyperprior distributions therefore become less and less important as the number of hierarchical layers grows, as long as the distributions are sufficiently broad (see, e.g. Roberts & Rosenthal 2001).

We restrict our attention to the one-layer hierarchical model in Eq. (13), and begin by modelling the hyperprior distribution $p(\eta)$. Its exact form is a modelling choice, and one possible approach is to let the hyperprior distribution be informed by the empirical approach to selecting η described in Sect. 3.3. Recall that this approach suggests that a reasonable point estimate of η can be obtained in the range within which the quotient between the weighted LS objective value and the regularisation objective value is between 2 and 4 using standard ZDI. For τ Sco, this region



Fig. 2. The trade-off between the least-squares fit and the regularisation objective as a function of η . The region where $\eta \in [16, 421]$ is shaded in purple.

is illustrated in Fig. 2. The figure shows the two components of the objective function in Eq. (7), i.e. the weighted LS objective and the regularisation objective, evaluated at the solution of the optimisation problem for a range of η -values. Specifically, for each η , we solve the optimisation problem given in Eq. (7), and plot the resulting weighted LS (data fidelity) terms $\|\Lambda^{\frac{1}{2}}(\mathbf{y} - \mathbf{f}(\hat{\mathbf{z}}))\|_{2}^{2}$ and regularisation terms $\eta r(\hat{\mathbf{z}})$ at the solutions $\hat{\mathbf{z}}$ as a function of η . The region over η where the regularisation term is between 2 and 4 times smaller than the LS objective is shaded in purple, and corresponds to $\eta \in [16, 421]$. As can be seen, the improvement in fit quality (reflected by a decreasing weighted LS objective value) begins to slow down in the vicinity of this region, as η decreases from the right in the plot. Moreover, the difference in magnitude between the weighted LS objective value and the regularisation term increases. Informed by this approach to hyperparameter selection in standard ZDI, we can model $p(\eta)$ as a distribution with non-zero probability on the interval $\eta \in [16, 421]$.

However, exactly inferring the posterior distribution as defined in Eq. (13) is generally intractable, as it is rare to have conjugacy in a hierarchical model even with a single hyperprior distribution. Thus, adopting a hierarchical modelling approach to incorporating uncertainty over the hyperparameters in our prior distribution requires approximate solution methods, such as Markov chain Monte Carlo (MCMC) methods (see, e.g. Meyn & Tweedie 2009; Cotter et al. 2013). While MCMC methods can provide statistically exact sample approximations of the posterior distribution, the computational demand can be prohibitive, especially in high-dimensional settings such as those arising from high-order spherical-harmonic magnetic field parameterisations. This challenge is further compounded when the forward model is even mildly computationally intensive (Brooks et al. 2011). Since we are operating in the weak-field regime in our analysis of τ Sco, under which our probabilistic model formulation conveniently allows for a closed-form expression of the posterior distribution in Eq. (9), we are interested in finding an alternative probabilistic model incorporating uncertainty in the hyperparameter η , allowing for a closed-form expression of the posterior distribution in Eq. (14). This can be achieved by approximating the continuous hyperprior distribution over η by a discrete distribution, obtained by evenly discretising the interval of non-zero hyperprior probability, i.e. $\eta \in [16, 421]$, in C components, and assigning each component a prior probability, or weight, ω_c . This is equivalent to modelling $p(\mathbf{z})$ as a *mixture prior* to incorporate our prior belief

about η directly into the prior distribution, in place of explicitly extending the model with a hyperprior distribution $p(\eta)$. A detailed account of this approach is provided in the following section.

3.4.6. Mixture priors

In this section, we extend the probabilistic model in two ways compared to the initial model described in Sect. 3.4.1, both of which adjusts the prior distribution to incorporate prior uncertainty over hyperparameters in the prior variance using mixture priors. A mixture prior can be constructed according to

$$p(\mathbf{z}) = \sum_{c=1}^{C} \omega_c p_c(\mathbf{z}), \tag{16}$$

where *C* is the number of mixture components and $p_c(\mathbf{z})$ denotes the *c*-th prior component. In this formulation, each prior component is a properly normalised probability density function and $\sum_{c=1}^{C} \omega_c = 1$, such that the mixture prior $p(\mathbf{z})$ is also properly normalised. With this notation, the posterior distribution can be expressed as

$$p(\mathbf{z}|\mathbf{y}) \propto \sum_{c=1}^{C} \omega_c p_c(\mathbf{z}) p(\mathbf{y}|\mathbf{z}),$$
(17)

where $p(\mathbf{y}|\mathbf{z})$ is the likelihood. Now, we define the marginal likelihood for each component *c* according to $p_c(\mathbf{y}) = \int p_c(\mathbf{z})p(\mathbf{y}|\mathbf{z})d\mathbf{z}$. We can then rewrite Eq. (17) according to

$$p(\mathbf{z}|\mathbf{y}) \propto \sum_{c=1}^{C} \omega_c p_c(\mathbf{y}) \frac{p_c(\mathbf{z}) p(\mathbf{y}|\mathbf{z})}{p_c(\mathbf{y})},$$
(18)

where $p_c(\mathbf{z})p(\mathbf{y}|\mathbf{z})/p_c(\mathbf{y})$ is the posterior distribution obtained using the prior distribution component $p_c(\mathbf{z})$. Thus, using a mixture prior in the form of Eq. (16), the posterior distribution is a mixture of the posterior distributions obtained from each component in the mixture prior distribution. The unnormalised weights in the posterior mixture are given by $\omega_c p_c(\mathbf{y})$ for each component c, where ω_c are the weights of the components in the mixture prior distribution. Properly normalised, the mixture posterior distribution becomes

$$p(\mathbf{z}|\mathbf{y}) = \left(\sum_{c=1}^{C} \omega_c p_c(\mathbf{y})\right)^{-1} \sum_{c=1}^{C} \omega_c p_c(\mathbf{y}) \frac{p_c(\mathbf{z})p(\mathbf{y}|\mathbf{z})}{p_c(\mathbf{y})}.$$
(19)

The connection between this model and the hierarchical model with a prior defined according to Eq. (15) is now clear. The mixture prior can be viewed as a discretisation of the integral arising when marginalising out η , with the mixture weights ω_c forming a discrete hyperprior distribution $p(\eta)$. Equal weights for each prior component then corresponds to a discrete uniform hyperprior distribution over η . We can also view each mixture component as a separate *model*, and interpret the prior weights as an expression of our prior belief in each model. With this view, our approach to incorporating uncertainty of prior hyperparameters into the probabilistic model is equivalent to *Bayesian model averaging* (see, e.g Hoeting et al. (1999)).

We begin by using a mixture prior to capture our prior uncertainty over the hyperparameter η in the prior variance term. The contribution of each component *c* to the prior distribution is modelled as a Gaussian according to $p_c(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{\Omega}_c^{-1})$, where $\mathbf{\Omega}_c^{-1}$ is a diagonal covariance matrix. We use the parameter $\Omega_{ci,i} = \eta_c l_i^2$ in the marginal prior distribution for each spherical-harmonic coefficient z_i , but now the prior variance depends on c. We use a mixture prior with C = 1000 components, where the parameters $\eta_c \in [\eta_1, \eta_C]$ are evenly distributed on the interval [16, 421], according to the earlier motivation.

It remains to determine the weights $\omega_c \in [\omega_1, \omega_c]$. According to Eq. (19), the weights in the mixture posterior distribution are scaled by the corresponding marginal likelihoods $p_c(\mathbf{y})$ compared to the weights in the mixture prior. The component-wise marginal likelihood is heavily dependent on η , and varies significantly in magnitude depending on the value of η_c . In fact, using equal prior weights ω_c for all components $c \in [1, 1000]$ results in negligible weights for all components in the mixture posterior, except the component c = 1. This prior component uses η_1 , the lowest possible value of η within our prior interval of non-zero probability. In light of this reflection, we note that the use of a mixture prior with equal weights in this case gives a result that is approximately equivalent to selecting η using *empirical Bayes* as discussed in Sect. 3.4.4, with an additional requirement constraining η to the region [16, 421] when solving the optimisation problem in Eq. (11). Thus, the choice of using uniform weights may not provide a sufficiently informative uncertainty quantification. Here, we choose a slightly different way of selecting the weights, where we instead incorporate a prior belief that, within the given interval of non-zero prior probability, the prior probability of choosing a model with a larger value of η is higher, reflecting our prior belief of to what extent parameters drawn from each prior component increases the risk of overfitting in the predictive model. To model this prior distribution, the unnormalised weights in the mixture prior are set to the inverse of the corresponding marginal likelihood:

$$\omega_c = \frac{p_c(\mathbf{y})^{-1}}{\sum_{c=1}^{C} p_c(\mathbf{y})^{-1}}.$$
(20)

We are also interested in incorporating prior uncertainty over the exponent of the angular degree *l* in the prior variance term. To do this, we use a mixture prior with two components. These components, $p_1(\mathbf{z})$ and $p_2(\mathbf{z})$, are both Gaussian, with a specific set of hyperparameters. Specifically, $p_1(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{\Omega}_1^{-1})$ and $p_2(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{\Omega}_2^{-1})$. Each covariance matrix, $\mathbf{\Omega}_1^{-1}$ and $\mathbf{\Omega}_2^{-1}$, is diagonal with parameters $\mathbf{\Omega}_{1i,i} = \eta_1 l_i$ and $\mathbf{\Omega}_{2i,i} = \eta_2 l_i^2$, respectively. We fix the prior variance of each component using $\eta_1 = 275.47$ and $\eta_2 = 100.00$, respectively. These values are chosen using the empirical approach discussed previously, based on independent fits using two separate models with priors adopting each set of hyperparameters, respectively. They provide comparable fits to the data in terms of mean deviation from observed values at the mean of the component-wise posterior distributions. We use the weights $\omega_1 = \omega_2 = 0.5$ in the mixture prior.

4. Results

In this section, we present the magnetic field distributions and corresponding uncertainty quantifications obtained using a) fixed prior hyperparameters chosen based on the empirical approach commonly used in standard ZDI, and b) the two probabilistic models extended by a mixture prior, as described in the previous section. For code and additional implementation details, we refer the reader to Appendix A.



Fig. 3. Comparison between the mean LSD Stokes V profiles (red) of the predictive distribution and the observed LSD Stokes V profiles (black), for a subset of the rotational phases. The profiles are offset vertically according to the rotational phase, as indicated in blue. The marginal predictive uncertainty is shaded in light grey and depicts three standard deviations. The corresponding posterior predictive uncertainty is shaded in light blue, and the blue bars represent the corresponding observational uncertainties.

4.1. Fixed hyperparameters

This section presents the posterior magnetic field distribution obtained using the statistical model described in Sect. 3.4.1 with fixed prior hyperparameters. Specifically, η is determined according to Sect. 3.3. We use the prior parameters $\Omega_{i,i} = \eta l_i^2$ with $\eta = 100$. At the mean of the posterior distribution, this value of η yields a quotient of 2.9 between the optimised weighted LS objective value and the regularisation objective value, evaluated according to the objective function in Eq. (7). A comparison between the observed and modelled LSD Stokes V profiles is illustrated in Fig. 3, which shows the best fit at the mean of the predictive distribution, along with the predictive uncertainty in terms of the marginal standard deviation of the predictive distribution. These quantities are derived from the predictive distribution given by Eq. (10), with closed form expressions for the mean vector and covariance matrix as detailed in Sect. 3.4.2. Note that the predictive uncertainty arises from two sources: the posterior uncertainty in the latent variables and the observational noise. As expected, the contribution from the posterior predictive uncertainty is generally larger towards the centre of the Stokes V profile at each rotational phase. See Appendix B for the corresponding results for all rotational phases. The mean marginal standard deviation, derived from the predictive distribution, is



Fig. 4. Rectangular maps of the reconstructed magnetic field and the corresponding uncertainty quantification. The left column shows the mean of the posterior magnetic field distribution across the stellar surface, in terms of the radial (top), meridional (middle) and azimuthal (bottom) magnetic field components. The field strength is given in kG. The grey transparent rectangle indicates the part of the stellar surface that is obscured from view. The right column shows the corresponding standard deviation maps. The results are obtained using a statistical model with hyperparameters $\Omega_{i,i} = \eta l_i^2$ and $\eta = 100$ in the marginal prior distribution for each spherical-harmonic coefficient z_i . Note that the contour plots in this paper consistently use the same number of levels between the minimum and maximum values in all figures.

 $4.98 \cdot 10^{-5}$ and the fit quality, in terms of mean deviation between the best fit and the observed Stokes V profiles at the mean of the predictive distribution, is $9.17 \cdot 10^{-5}$. These results account for all rotational phases.

The maps of the mean and standard deviation of the posterior magnetic field distribution are displayed in Fig. 4. As evident from these images, the surface distribution of the ZDI uncertainties exhibits a distinct latitudinal pattern, reflecting the different sensitivity of the disc-integrated Stokes V profiles to the three components of the magnetic field vector. Specifically, the radial and azimuthal fields are determined most precisely around the sub-observer latitude whereas the meridional field has the least error at the visible rotational pole. At the same time, the precision of the azimuthal field drops off towards the visible pole less sharply than for the radial field. The lack of longitudinal variation of the standard deviation maps is due to a dense rotational phase sampling of the particular observational data set analysed in our paper. In other situations, for example when large phase gaps exist in the data, we would expect our method to yield a noticeable longitudinal variation of uncertainties corresponding to this uneven phase coverage.

In terms of the average standard deviation, we find values in the 0.056–0.067 kG range. This corresponds to 10.4–14.0% of the peak (95 percentile) values of the respective mean magnetic maps. If we instead consider maximum errors (95 percentile for the visible part of the surface), we obtain 0.065–0.074 kG corresponding to 12.9–15.7% of the magnetic field amplitudes. These two fractional error estimates illustrate the average (most representative) and maximum (conservative) errors of the ZDI inversion for τ Sco with fixed hyperparameters.



Fig. 5. Comparison of predicted (mean) LSD Stokes *V* profiles from independent inversions using $\eta = 16$ (blue) and $\eta = 421$ (red) for a subset of the rotational phases. The mean deviations between the means of the predictive distributions and the observed Stokes *V* profiles (black) is $7.87 \cdot 10^{-5}$ and $1.05 \cdot 10^{-4}$, respectively, across all rotational phases.

4.2. Mixture priors

We analyse the resulting posterior magnetic field distributions obtained using two different mixture priors, capturing prior uncertainty over the exponent of the angular degree l and η , respectively, in the parameterisation of the prior variance.

First, we present the posterior magnetic field distribution using a mixture prior with 1000 components capturing the prior uncertainty over η . According to Sect. 3.4.6, each prior component $p_c(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{\Omega}_c^{-1})$ has η -dependent parameters $\Omega_{ci,i} = \eta_c l_i^2$, with parameters η_c evenly distributed on the interval [16, 421]. To illustrate the impact of η on the predicted LSD Stokes V profiles, the computed means based on independent inversions using $\eta = 16$ and $\eta = 421$, respectively, are presented in Fig. 5. Using the prior weights given in Eq. (20), we obtain equal posterior weights for all components in the mixture posterior distribution according to Eq. (19). The final inversion results in a mean deviation of $9.6 \cdot 10^{-5}$ between the mean of the predictive distribution and the observed LSD Stokes V profiles, similar in magnitude to the fixed hyperparameter case. The mean marginal standard deviation derived from the predictive distribution has increased slightly, reaching $5.16 \cdot 10^{-5}$. Appendix B shows a comparison between the observed LSD Stokes V profiles and the prediction at the mean of the predictive distribution, together with the predictive uncertainty.

The mean and standard deviation of the posterior magnetic field distribution is calculated as that of the mixture distribution



Fig. 6. Same as Fig. 4 but for results obtained using a statistical model with a mixture prior consisting of 1000 η -dependent components $p_c(\mathbf{z})$.

obtained by calculating the posterior magnetic field distribution corresponding to each component in the mixture posterior distribution given in Eq. (19). Fig. 6 shows the mean of the resulting mixture posterior magnetic field distribution across the stellar surface, together with the corresponding standard deviation maps. Four samples from the resulting distribution are displayed in Fig. 7. It can be noted that with C = 1000components, the obtained mean magnetic field maps and corresponding uncertainty maps do not change significantly when the number of components increases. Thus, the obtained results closely correspond to using the corresponding continuous hyperprior $p(\eta)$ in a hierarchical model.

As one can see in Fig. 6, the surface pattern of the standard deviation changes drastically compared to the results presented in the previous section. Instead of the smooth latitudinal variation seen in Fig. 4, we find a highly structured standard deviation map, with the largest scatter roughly corresponding to the strongest (by absolute value) features in the mean magnetic maps. This is to be expected since decreasing η generally leads to an increase in the contrast of each map, with sharper details and more prominent small-scale features. Quantifying the standard deviation distribution, we derive the average values 0.063–0.088 kG and maximum values 0.087–0.139 kG. This corresponds to 15.9–18.0% and 22.2–28.8.0% of the 95 percentile mean magnetic map amplitudes, respectively – significantly larger than in the fixed- η case.

Next, we present the magnetic field distribution corresponding to the two-component mixture prior. Recall that $p_1(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{\Omega}_1^{-1})$ and $p_2(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{\Omega}_2^{-1})$, with parameters $\Omega_{1i,i} = \eta_1 l_i$ and $\Omega_{2i,i} = \eta_2 l_i^2$. With the specific choice of hyperparameters defined in Sect. 3.4.6, and prior weights $\omega_1 = \omega_2 = 0.5$, the marginal likelihoods coincide, i.e. $p_1(\mathbf{y}) = p_2(\mathbf{y})$. Consequently, the normalised weights in the resulting mixture posterior becomes $\left(\sum_{c=1}^2 \omega_c p_c(\mathbf{y})\right)^{-1} \omega_1 p_1(\mathbf{y}) = \left(\sum_{c=1}^2 \omega_c p_c(\mathbf{y})\right)^{-1} \omega_2 p_2(\mathbf{y}) = 0.5$. The mean deviation between the mean of the resulting predictive distribution and the observed LSD Stokes *V* profiles is 9.09 $\cdot 10^{-5}$. In this case, the mean marginal standard deviation of the predictive distribution corresponding to the computed LSD Stokes *V* profiles is 5.00 $\cdot 10^{-5}$. See Appendix B for a comparison between the observed LSD Stokes *V* profiles and the prediction at the observed LSD Stokes *V* profiles and the prediction at the subserved LSD Stokes *V* profiles and the prediction at the subserved LSD Stokes *V* profiles and the prediction at the subserved LSD Stokes *V* profiles and the prediction at the subserved LSD Stokes *V* profiles and the prediction at the subserved LSD Stokes *V* profiles and the prediction at the subserved LSD Stokes *V* profiles and the prediction at the subserved LSD Stokes *V* profiles and the prediction at the subserved LSD Stokes *V* profiles and the prediction at the subserved LSD Stokes *V* profiles and the prediction at the subserved LSD Stokes *V* profiles and the prediction at the subserved LSD Stokes *V* profiles and the prediction at the subserved LSD Stokes *V* profiles and the prediction at the subserved LSD Stokes *V* profiles and the prediction at the subserved LSD Stokes *V* profiles and the prediction at the subserved LSD Stokes *V* profiles and the prediction at the subserved LSD Stokes *V* profiles and the prediction at the subserved LSD S

mean of the predictive distribution. Fig. 8 shows the mean of the posterior magnetic field distribution across the stellar surface, together with the corresponding standard deviation maps. Four samples from the resulting distribution are displayed in Fig. 9. In this case, the standard deviation maps retain some of the small-scale structures highlighted in the discussion of Fig. 6, but these structures are no longer closely associated with the strongest features in the magnetic maps. The representative amplitude of the standard deviation is comparable to the previous test: 0.072–0.098 kG on average and 0.082–0.137 kG if one considers the 95 percentiles, i.e. 15.3–19.5% and 19.7–27.4% of the magnetic map amplitudes, respectively.

4.3. Magnetic energy spectrum

Using our spherical-harmonic field parameterisation, it is straightforward to assess the strength of the axisymmetric vs. nonaxisymmetric field, as well as the strength of the poloidal vs. toroidal field components. Specifically, the $\gamma_{l,m}$ coefficients specify the strength of the toroidal field component, whereas the $\alpha_{l,m}$ and $\beta_{l,m}$ coefficients specify the vertical and horizontal components of the poloidal field, respectively. Given the posterior distribution obtained from each independent choice of prior distribution, we have analysed the numerical distribution of the magnetic field energy contributions over different harmonic modes (*l*-modes). These distributions, together with a comparison of the distributions of the magnetic energy of the poloidal and toroidal field components over the harmonic modes, are illustrated in Fig. 10–12. Fig. 10 shows the distributions corresponding to the fixed hyperparameter case, with $\Omega_{i,i} = \eta l_i^2$ and $\eta = 100$. These magnetic energies correspond to the posterior magnetic field distribution in Fig. 4. As we can see, the total magnetic energy is spread over modes $l \in [1, 5]$, with a clear peak at l = 1. Considering the poloidal and toroidal field components separately, we note a significant peak in the poloidal field component at l = 1, followed by a smaller peak at l = 4. In contrast, the peak of the toroidal field component ranges over $l \in [1, 2]$, whereafter the contribution decreases with l. Fig. 11 shows the distribution of magnetic energy over the harmonic modes corresponding to the posterior magnetic field distribution presented in Fig. 6, obtained from a mixture prior with with 1000 η -dependent components $p_c(\mathbf{z})$. Finally, Fig. 12 shows the corresponding results for the posterior magnetic field distribution given in Fig. 8, obtained from a mixture prior with components $p_1(\mathbf{z})$ and $p_2(\mathbf{z})$.

The magnetic energy distributions over harmonic modes illustrated in Fig. 11–12 are qualitatively similar to the distributions in the fixed hyperparameter case illustrated in Fig. 10. Considering the comparatively large scatter in the resulting magnetic energy based on statistical models with mixture priors, our results indicate high uncertainty in the relative magnetic energies recovered in the ZDI analysis of τ Sco.

In addition to analysing the magnetic energy distribution for poloidal vs toroidal field components, we divide the spherical-harmonic modes into two other groups: those with |m| < l/2 and those with $|m| \ge l/2$ (Fares et al. 2009). A statistical summary of the magnetic energy distributions in terms of a) the total contribution from the poloidal field component across all harmonic modes, and b) the isolated contribution from the |m| < l/2 field component, is presented in Table 1 for each of the three cases.

While the magnetic energy, as discussed here, is more interpretable and comparable between studies, we have included additional visualisations of the posterior distribution over z in Appendix C for completeness.



Fig. 7. Four samples of surface magnetic field vector maps drawn from the posterior distribution. The distribution is obtained through posterior inference using a statistical model with a mixture prior consisting of 1000 η -dependent components $p_c(\mathbf{z})$. From the left, the samples are drawn from components of the mixture posterior distribution with parameters $\eta \approx 34$, $\eta \approx 123$, $\eta \approx 62$ and $\eta \approx 258$, respectively.



Fig. 8. Same as Fig. 4 but for results obtained using a statistical model with a mixture prior consisting of the two components $p_1(\mathbf{z})$ and $p_2(\mathbf{z})$.

Table 1. Quantiles of the posterior magnetic energy fractions.

	Poloidal field			m < l/2		
quantile	0.05	0.50	0.95	0.05	0.50	0.95
Case 1	0.38	0.42	0.46	0.09	0.12	0.15
Case 2	0.38	0.50	0.61	0.09	0.13	0.16
Case 3	0.39	0.46	0.53	0.09	0.15	0.19

Notes. Quantiles of the posterior distribution of the fraction of the total magnetic energy for a) the poloidal field components across harmonic modes and b) the |m| < l/2 field component. Case 1 refers to the fixed hyperparameter case, Case 2 refers to the distribution obtained from a mixture prior with 1000 η -dependent components $p_c(\mathbf{z})$, and Case 3 refers to the distribution obtained from a mixture prior with components $p_1(\mathbf{z})$ and $p_2(\mathbf{z})$.

5. Conclusions and discussion

In this study, we have presented magnetic field distributions and corresponding uncertainty maps resulting from a probabilistic ZDI analysis of the early-type star τ Sco. We considered three distinct prior distributions over the latent variables in the sphericalharmonic field parameterisation. When the hyperparameters in the proposed prior distribution were chosen using the empirical Bayes approach, our investigation showed that the amplitude in the posterior magnetic field distribution violated the weak-field approximation motivating our choice of forward mapping function $\mathbf{f}(\mathbf{z})$. We concluded that maximum likelihood estimation of the hyperparameter η is ill-posed in this case. The uncertainty maps obtained with empirically fixed hyperparameters exhibited a smooth latitudinal variation across the stellar surface. By introducing mixture priors, we accounted for prior uncertainty over two hyperparameters in the original prior distribution, capturing the sensitivity of the magnetic field distribution to the specific choice of prior and, by extension, to the hyperparameters used in the classical ZDI framework. A mixture prior over η increased the structure in the uncertainty maps and raised the overall uncertainty level, particularly around the features in the mean magnetic map with the strongest amplitudes. A similar uncertainty level was observed when introducing a mixture prior over the exponent on the angular degree l. Although small-scale structures remained, the correlation with the strongest magnetic features in the mean magnetic field distribution was less prominent.

We also analysed the magnetic energy spectrum, with emphasis on the posterior magnetic energy distribution across *l*-modes. Compared to previous ZDI inversions targeting τ Sco, a star assumed to exhibit a complex, non-dipolar surface magnetic field, the fact that the magnetic energy contribution is dominated by l = 1 is unexpected. This result alone would indicate a predominantly dipolar field geometry, in line with the majority of hot, magnetic stars, but in contrast to previous studies of τ Sco. We investigated this issue by an in-depth comparison between our results and the magnetic field map obtained by Kochukhov & Wade (2016) with the same harmonic parameterisation ("model 2" in



Fig. 9. Same as Fig. 7 but the four samples are drawn from the posterior magnetic field distribution obtained through posterior inference using a statistical model with a mixture prior consisting of the two components $p_1(z)$ and $p_2(z)$. From the left, the first and third column represent samples drawn from the first component of the mixture posterior distribution, and the second and fourth column represent samples drawn from the second component.



Fig. 10. Violin plots illustrating the distribution of magnetic energy over different harmonic modes. In these plots the widths of the coloured regions represent the smoothed probability density of the fraction of the total magnetic energy, numerically obtained by sampling from the posterior distribution $p(\mathbf{z}|\mathbf{y})$ corresponding to the results illustrated in Fig. 4. Horizontal lines mark the 0.05-, 0.50- and 0.95-quantiles of the distribution at each mode, respectively. The left panel compares the energy of the toroidal contribution (blue) and the poloidal components (red) as a function of angular degree *l*, with values normalised by the total energy, i.e., $(E_t, E_p)/E_{tot}$. The right panel shows the total magnetic energy of the poloidal and toroidal field components as a function of angular degree *l*, also normalised by the total energy, i.e., $(E_p + E_t)/E_{tot}$.

that paper). It turns out that the overall morphology of the surface distributions of the three vector magnetic components is qualitatively similar in the two studies. However, their relative strengths differ. Here we recover about 20–30% less relative energy in the radial and azimuthal components, which are dominated by higher l modes. Conversely, our meridional field compared to the results by Kochukhov & Wade (2016). This meridional field component features a simpler, dipolar-like geometry (see Fig. 4). Consequently, the contribution of the l = 1 harmonics is significantly higher in our results. This modification of the outcome of the ZDI analysis is sufficient to noticeably alter the resulting harmonic energy spectrum. While Kochukhov & Wade (2016) found the poloidal field energy to be spread over $l \in [1, 4]$ and observed a



Fig. 11. Same as Fig. 10 but for the results obtained from the posterior distribution $p(\mathbf{z}|\mathbf{y})$ corresponding to the magnetic field in Fig. 6.



Fig. 12. Same as Fig. 10 but for the results obtained from the posterior distribution $p(\mathbf{z}|\mathbf{y})$ corresponding to the magnetic field in Fig. 8.

peak in the toroidal field contribution at $l \in [2, 3]$, we recover a stronger toroidal component peaking at l = 1 with a less important poloidal field peaking at both l = 1 and l = 4.

Since there are differences in modelling choices between the ZDI reconstructions, we do not expect to exactly reproduce the results at the mean of the posterior magnetic field distribution in this case. In general, it is worth remarking that the uncertainty quantification obtained from the proposed framework for probabilistic ZDI is conditioned on our choice of probabilistic model, i.e. $p(\mathbf{y}|\mathbf{z})$ and $p(\mathbf{z})$, and the information content in the observa-

tions v. As noted in Sect. 1, earlier studies of the surface magnetic field of τ Sco have compared the field geometries obtained using different spherical-harmonic magnetic field parameterisations, ultimately concluding that the topological details in the surface magnetic field, as well as the average magnetic field strengths, vary significantly depending on the parameterisation (Kochukhov & Wade 2016). Since the Bayesian framework naturally extends to Bayesian model selection and Bayesian model averaging, an interesting avenue for future work is to perform probabilistic ZDI analyses of τ Sco using competing field parameterisations to model the likelihood.

While this study has focused on ZDI inversion in the weakfield limit, another important direction for future research is to investigate the feasibility of employing probabilistic ZDI to generate reliable uncertainty maps when the line-profile response cannot be modelled under this assumption. Addressing this challenge will require advanced computational techniques, such as MCMC methods or variational Bayes (Blei et al. 2017), since the posterior distribution $p(\mathbf{z}|\mathbf{y})$ cannot be expressed in closed form for non-linear response profiles under standard model assumptions. Additionally, the proposed framework for probabilistic ZDI can be expanded by extending the hierarchy in the statistical model (see Sect. 3.4.5) to account for prior uncertainty in stellar parameters beyond the spherical-harmonic coefficients. Such an extension would enable quantification of uncertainty in the surface magnetic field of more complex ZDI targets, including stars in eclipsing binary systems and equator-on hosts of transiting exoplanets, where degeneracy in the latent variables within the spherical-harmonic formulation is likely to occur. Both of these avenues will be pursued in future work.

To summarise our study, we have proposed a Bayesian framework for probabilistic ZDI allowing for formal uncertainty quantification of the obtained stellar magnetic field maps, given a single set of spectropolarimetric observations. Coupling the magnetic field distributions with interpretable uncertainty maps makes it possible to reason about the uncertainty in the obtained magnetic field distributions in a meaningful way. This is something that has not previously been possible, as classical ZDI has generally been restricted to point estimates. We have demonstrated that, for stars exhibiting relatively weak surface magnetic fields such that the weak-field approximation can be used to model the line-profile response, a closed-form solution for the posterior distribution over the spherical-harmonic coefficients exists under specific model assumptions. Our choice of probabilistic model is convenient in the sense that it makes it easy to quickly conduct Bayesian analyses of the surface magnetic fields of a large group of stars, despite the high-dimensional spherical-harmonic magnetic field parameterisation. Since we obtain an uncertainty quantification centered around the point estimate obtained using standard ZDI, our probabilistic ZDI formulation essentially provides uncertainty quantification as a byproduct compared to standard ZDI, without loss of information.

Acknowledgements. This work was supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. Computations were enabled by the Berzelius resource provided by the Knut and Alice Wallenberg Foundation at the National Supercomputer Centre. O.K. acknowledges support by the Swedish Research Council (grant agreements no. 2019-03548 and 2023-03667), the Swedish National Space Agency, and a sabbatical fellowship from AI4Research at Uppsala University.

References

Bishop, C. M. 2007, Pattern Recognition and Machine Learning (Information Science and Statistics), 1st edn. (Springer)

- Blei, D. M., Kucukelbir, A., & McAuliffe, J. D. 2017, Journal of the American Statistical Association, 112, 859
- Bouvier, J., Matt, S. P., Mohanty, S., et al. 2014, in Protostars and Planets VI, ed.
- H. Beuther, R. S. Klessen, C. P. Dullemond, & T. Henning, 433–450 Braithwaite, J. 2008, MNRAS, 386, 1947 Brooks, S., Gelman, A., Jones, G., & Meng, X.-L. 2011, Handbook of Markov
- Chain Monte Carlo (CRC press)
- Brown, S. F., Donati, J.-F., Rees, D. E., & Semel, M. 1991, A&A, 250, 463 Cotter, S., Roberts, G., Stuart, A., & White, D. 2013, Statistical Science, 28, 424
- Donati, J., Catala, C., Landstreet, J. D., & Petit, P. 2006a, in Astronomical Society
- of the Pacific Conference Series, Vol. 358, Astronomical Society of the Pacific Conference Series, vol. 558, Astronomical Society of the Pacific Conference Series, ed. R. Casini & B. W. Lites, 362–368 Donati, J.-F. 1999, MNRAS, 302, 457 Donati, J.-F. & Brown, S. F. 1997, A&A, 326, 1135 Donati, J. F., Howarth, I. D., Jardine, M. M., et al. 2006b, MNRAS, 370, 629 Donati, J.-F. & Landstreet, J. D. 2009, ARA&A, 47, 333 Donati, J. F., Semel, M., Carter, B. D., Rees, D. E., & Collier Cameron, A. 1997,

- MNRAS, 291, 658
- Fares, R., Donati, J. F., Moutou, C., et al. 2009, MNRAS, 398, 1383 Feiden, G. A. 2016, A&A, 593, A99 Folsom, C. P., Bouvier, J., Petit, P., et al. 2018, MNRAS, 474, 4956

- Gelman, A., Carlin, J., Stern, H., et al. 2020, Bayesian Data Analysis Grunhut, J. H., Wade, G. A., Neiner, C., et al. 2017, MNRAS, 465, 2432
- Hoeting, J. A., Madigan, D., Raftery, A. E., & Volinsky, C. T. 1999, Statistical Science, 14, 382 Hussain, G. A. J., Donati, J.-F., Collier Cameron, A., & Barnes, J. R. 2000,
- MNRAS, 318, 961
- Karvonen, T. & Oates, C. J. 2024, J. Mach. Learn. Res., 24 Keszthelyi, Z., Meynet, G., Martins, F., de Koter, A., & David-Uraz, A. 2021, MNRAS, 504, 2474
- Khokhlova, V. L., Rice, J. B., & Wehlau, W. H. 1986, ApJ, 307, 768 Kochukhov, O. 2016, in Lecture Notes in Physics, Vol. 914, Lecture Notes in Physics, ed. J.-P. Rozelot & C. Neiner, 177-204
- Kochukhov, O. 2016, in Cartography of the Sun and the stars, ed. J.-P. Rozelot & C. Neiner (Springer International Publishing), 177–204 Kochukhov, O. 2017, A&A, 597, A58
- Kochukhov, O., Lüftinger, T., Neiner, C., Alecian, E., & MiMeS Collaboration. 2014, A&A, 565, A83 Kochukhov, O., Makaganiuk, V., & Piskunov, N. 2010, A&A, 524, A5 Kochukhov, O., Makaganiuk, V., Piskunov, N., et al. 2011, ApJ, 732, L19 Kochukhov, O., Papakonstantinou, N., & Neiner, C. 2022, MNRAS, 510, 5821 Kochukhov, O., Piskunov, N., Ilyin, I., Ilyina, S., & Tuominen, I. 2002, A&A,

- - 389, 420

 - Kochukhov, O., Shultz, M., & Neiner, C. 2019, A&A, 621, A47 Kochukhov, O. & Wade, G. A. 2016, A&A, 586, A30 Landi Degl'Innocenti, E. & Landolfi, M. 2004, Astrophysics and Space Science Library, Vol. 307, Polarization in Spectral Lines (Kluwer Academic Publishers) Lanza, A. F. 2018, in Handbook of Exoplanets, ed. H. J. Deeg & J. A. Belmonte,
 - 17

 - Matt, S. & Pudritz, R. E. 2008, ApJ, 678, 1109 Meyn, S. & Tweedie, R. L. 2009, Markov Chains and Stochastic Stability, 2nd edn. (USA: Cambridge University Press)

 - Morin, J., Donati, J. F., Petit, P., et al. 2008, MNRAS, 390, 567 Morin, J., Donati, J. F., Petit, P., et al. 2010, MNRAS, 407, 2269 Murphy, K. P. 2022, Probabilistic Machine Learning: An introduction (MIT Press) Nieva, M. F. & Przybilla, N. 2014, A&A, 566, A7 Petit, P., Dintrans, B., Solanki, S. K., et al. 2008, MNRAS, 388, 80 Petit, P., Louge, T., Théado, S., et al. 2014, PASP, 126, 469 Petit, V. & Wade, G. A. 2012, MNRAS, 420, 773

 - Piskunov, N. & Kochukhov, O. 2002, A&A, 381, 736 Przybilla, N., Firnstein, M., Nieva, M. F., Meynet, G., & Maeder, A. 2010, A&A, 517, A38

 - Reiners, A. 2012, Living Reviews in Solar Physics, 9, 1 Roberts, G. O. & Rosenthal, J. S. 2001, Bernoulli, 7, 453
 - Rosén, L., Kochukhov, O., & Wade, G. A. 2015, ApJ, 805, 169
 - Schneider, F. R. N., Ohlmann, S. T., Podsiadlowski, P., et al. 2020, MNRAS, 495, 2796
 - Schneider, F. R. N., Ohlmann, S. T., Podsiadlowski, P., et al. 2019, Nature, 574, 211
 - Schneider, F. R. N., Podsiadlowski, P., Langer, N., Castro, N., & Fossati, L. 2016, MNRAS, 457, 2355

 - Shu, F., Najita, J., Ostriker, E., et al. 1994, ApJ, 429, 781 Shultz, M. E., Wade, G. A., Rivinius, T., et al. 2018, MNRAS, 475, 5144 Sikora, J., Wade, G. A., Power, J., & Neiner, C. 2019, MNRAS, 483, 2300

 - Sjölund, J., Eklund, A., Özarslan, E., et al. 2018, NeuroImage, 175, 272
 - Strugarek, A. 2016, ApJ, 833, 140 Vidotto, A. A. 2020, in IAU Symposium, Vol. 354, Solar and Stellar Magnetic
 - Fields: Origins and Manifestations, ed. A. Kosovichev, S. Strassmeier, & M. Jardine, 259-267
- Vidotto, A. A., Opher, M., Jatenco-Pereira, V., & Gombosi, T. I. 2009, ApJ, 699, 441

 - Vogt, S. S., Hatzes, A. P., Misch, A. A., & Kürster, M. 1999, ApJS, 121, 547 Vogt, S. S., Penrod, G. D., & Hatzes, A. P. 1987, ApJ, 321, 496 Wade, G. A., Donati, J. F., Landstreet, J. D., & Shorlin, S. L. S. 2000, MNRAS, 313.823
 - Weber, E. J. & Davis, Leverett, J. 1967, ApJ, 148, 217

Article number, page 13 of 18

Appendix A: Implementation details

Our probabilistic ZDI framework is implemented in Python, utilizing NumPy for general numerical computations and JAX for efficient linear algebra and automatic differentiation. In particular, JAX's automatic differentiation functionality is employed to compute the design matrix A in the forward model, a key component in the Bayesian formulation. For further implementation details, see our probabilistic ZDI code made available on GitHub⁴.

Appendix B: Predictive distribution

In Fig. 3, we showed a comparison between the observed LSD Stokes V profiles and the reconstruction at the mean of the predictive distribution for a subset of the rotational phases, together with the posterior and predictive uncertainties. Fig. B.1 shows the corresponding results for all rotational phases.



Fig. B.1. Same as Fig. 3 but for all rotational phases. Recall that the observed profiles are depicted in black, and the reconstruction at the mean is depicted in red. The shaded light grey area shows the marginal predictive uncertainty, whereas the shaded light blue area shows the corresponding posterior predictive uncertainty.

Fig. B.2 and Fig. B.3 show the corresponding results obtained from statistical models using the two different mixture priors explored in this paper (referred to as Case 2 and Case 3). We observe that the posterior predictive uncertainty is significantly larger using the mixture prior with 1000 η -dependent components (Case 2).

⁴ https://github.com/jenan007/ProbabilisticZDI/



Fig. B.2. Same as Fig. B.1 but for results obtained using a statistical model with **a**) a mixture prior consisting of 1000 η -dependent components $p_c(\mathbf{z})$. and **b**) a mixture prior consisting of the two components $p_1(\mathbf{z})$ and $p_2(\mathbf{z})$.



Fig. B.3. Same as Fig. 3 but for results obtained using a statistical model with **a**) a mixture prior consisting of 1000 η -dependent components $p_c(\mathbf{z})$. and **b**) a mixture prior consisting of the two components $p_1(\mathbf{z})$ and $p_2(\mathbf{z})$.

Appendix C: Results for spherical harmonic coefficients

In the main paper, we focus on the posterior magnetic field distribution and the magnetic energy distributions as a function of *l*. The posterior distribution over the spherical harmonic coefficients z is also readily available but more difficult to interpret. In addition, the values of the coefficients are generally not comparable between studies due to different formulations in terms of e.g. real and imaginary representations of spherical harmonic expansions and different normalisation approaches. For completeness, we present the covariance matrices and triangular plots for the three posterior distributions p(z|y) presented in this paper. These results can be found in Fig. C.1, Fig. C.2 and Fig. C.3, corresponding to Case 1, Case 2 and Case 3, respectively. The uncertainties over the coefficients are dominated by the trend with *l*. The apparent trends with *m* (e.g. a higher uncertainty for the odd-*m* toroidal modes for l = 1 but a lower uncertainty for the same modes for l = 2) reflect the amplitude of recovered harmonic coefficients and disappear when one considers uncertainties normalised by the absolute value of the coefficients.



Fig. C.1. Graphical representation of the spherical harmonic coefficients and their covariances. This plot corresponds to the magnetic field reconstruction results in Fig. 4. **a**) Covariance matrix with the three groups of spherical harmonic coefficients, α , β , γ , stored sequentially in the order of increasing *l* and *m* numbers. Dashed lines highlight parts of the covariance matrix corresponding to each group of harmonic coefficients. **b**) Mean value of the spherical harmonic coefficients. **c**) Standard deviation of the spherical harmonic coefficients.



Fig. C.2. Same as Fig. C.1 but for the spherical harmonic coefficients corresponding to the magnetic field reconstruction in Fig. 6.



Fig. C.3. Same as Fig. C.1 but for the spherical harmonic coefficients corresponding to the magnetic field reconstruction in Fig. 8.