

---

# Bayesian Uncertainty Quantification in High-dimensional Stellar Magnetic Field Models

---

Jennifer R. Andersson<sup>1</sup> Oleg Kochukhov<sup>2</sup> Zheng Zhao<sup>1</sup> Jens Sjölund<sup>1</sup>

## Abstract

Spectropolarimetric inversion techniques, known as Zeeman Doppler imaging (ZDI), have become the standard tools for reconstructing surface magnetic field maps of stars. Accurate and efficient uncertainty quantification of such magnetic field maps is an open problem in current research, and the high dimensionality of the spherical-harmonic magnetic field parameterization makes inference inherently difficult. We propose a probabilistic machine learning framework for stellar surface magnetic field reconstruction using a gradient-based Metropolis-adjusted Langevin algorithm. By efficient implementation in JAX, our framework allows for reliable uncertainty quantification of the global stellar magnetic field topology. We test the proposed scheme on the bright, massive star  $\tau$  Scorpii, and show that our approach enables accurate computation of the posterior magnetic field distribution with fast convergence.

## 1. Introduction

Magnetic fields drive many key stellar processes throughout a star’s lifetime. Since distant stellar surfaces cannot be spatially resolved using imaging techniques, spectropolarimetric time series observations are necessary in the reconstruction of surface magnetic fields of stars other than the Sun. For such stars, Zeeman Doppler imaging (ZDI, Kochukhov, 2016) can be used to find the inverse mapping between two-dimensional stellar surface magnetic field maps and a set of spectropolarimetric line profile observations recorded at different rotational phases of the target star (see, e.g., Boro Saikia et al., 2022; Grunhut et al., 2022;

Petit et al., 2022). However, despite the significant role of magnetic fields in stellar evolutionary processes, essentially all published magnetic field maps represent point estimates with no formal quantification of the uncertainty, including recent ZDI inversions by, e.g., Marsden et al. (2023) and Fréour et al. (2023). This is a major shortcoming of modern ZDI implementations. While some attempts at quantifying the field uncertainty have been made, they do not capture the family of field topologies capable of fitting a given set of observations. For example, Kochukhov et al. (2019) and Kochukhov et al. (2022) provide uncertainty quantifications by empirical comparison of magnetic field maps obtained from independent spectropolarimetric time series — a noisy approach only feasible for strongly magnetic stars. Piskunov & Kochukhov (2002) do the same by considering the diagonal of the Hessian matrix, which ignores the correlation between parameters. Petit & Wade (2011) use Bayesian analysis to determine stellar magnetic field properties from circular spectropolarimetric observations, but their model is restricted to an oblique dipole.

In this work, we formalize a probabilistic ZDI framework focusing on high-dimensional spherical-harmonic field representations, by extending the standard ZDI formulation to a fully Bayesian setting. To obtain the posterior magnetic field distribution, we leverage Markov chain Monte Carlo (MCMC) methods. These methods come with convergence guarantees without requiring the specification of a family of variational distributions, unlike the commonly used variational Bayes approaches (Jordan et al., 1999; Blei et al., 2017). Indeed, proposing such a variational distribution family is challenging due to the limited amount of previous studies providing theoretical and numerical indications. While the high dimensionality of the inversion problem and complexity of the forward mapping process pose significant challenges in the MCMC setting, we show that fast and stable convergence can be achieved using the high-performance numerical computing and machine learning library JAX (Bradbury et al., 2018). Implementations for two different forward models will be published at <https://github.com/jenan007/ProbabilisticZDI/>. Moreover, we show that, under certain model assumptions, the Bayesian formulation collapses to basic Bayesian linear regression when the

---

<sup>1</sup>Department of Information Technology, Uppsala University, Uppsala, Sweden <sup>2</sup>Department of Physics and Astronomy, Uppsala University, Uppsala, Sweden. Correspondence to: Jennifer R. Andersson <jennifer.andersson@it.uu.se>.




Figure 1. Our pipeline for probabilistic ZDI. The approach involves fitting a probabilistic model to a set of spectropolarimetric line profiles observed at different rotational phases. From the resulting posterior distribution, we obtain samples  $\mathbf{z}'$ , which are used to reconstruct two-dimensional magnetic field maps and corresponding synthetic line profiles.

weak field approximation (Landi Degl’Innocenti & Landolfi, 2004) is used, allowing for analytical derivation of the posterior magnetic field distribution for a large group of stars.

We demonstrate our approach by reconstructing the surface magnetic field of the bright, massive star  $\tau$  Scorpii ( $\tau$  Sco, hd 149438, HR 6165), with corresponding uncertainties across the stellar surface. Previous results produce point estimates that predict an unusually complex surface magnetic field topology for  $\tau$  Sco with significant contribution in many higher-order harmonic modes (Donati et al., 2006), motivating the need for a high-dimensional field parameterization. Moreover, Kochukhov & Wade (2016) find that the magnetic field maps obtained from ZDI inversion of  $\tau$  Sco are highly sensitive to specific assumptions in the spherical-harmonic field parameterization, further motivating why it is particularly interesting to accurately quantify the uncertainty of the reconstructed magnetic field maps for this star.

## 2. ZDI of Stellar Magnetic Fields

In this section, the spherical-harmonic magnetic field parameterization and forward model are introduced, followed by a description of the standard ZDI inversion.

### 2.1. Magnetic Field Model

Our stellar surface magnetic field model is parameterized by a spherical-harmonic expansion, which has been widely used in recent research (e.g., Kochukhov et al., 2014; Folsom et al., 2018; Grunhut et al., 2022; Marsden et al., 2023). The inversion problem resulting from this field representation is inherently high-dimensional, with hundreds of coefficients truncated from an infinite series expansion. When the exact solution is available, however, it comes with substantial benefits. For example, the spherical-harmonic representation guarantees that the net signed magnetic flux through the closed stellar surface is zero, fulfilling the corresponding divergence-free constraint imposed by Maxwell’s laws. This property in general does not hold for the magnetic field generated by direct field parameterization. For a detailed exposition of the spherical-harmonic field representation, see, e.g., Kochukhov & Wade (2016).

Let  $B_r(\theta, \phi)$ ,  $B_\theta(\theta, \phi)$ , and  $B_\phi(\theta, \phi)$  denote the surface magnetic field components in the radial, meridional, and azimuthal direction, respectively, as a function of the surface latitude  $\theta$  and longitude  $\phi$ . Our spherical-harmonic field representation is defined as follows:

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

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

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

Here,  $l$  and  $m$  denote the degree and order, respectively, of the spherical-harmonic functions  $X_{l,m}(\theta, \phi)$ ,  $Y_{l,m}(\theta, \phi)$ , and  $Z_{l,m}(\theta, \phi)$ . See Kochukhov et al. (2014) for a more detailed description of these equations. The free parameters in the resulting spectropolarimetric inversion problem are the amplitudes of the spherical-harmonic modes,  $\alpha_{l,m}$ ,  $\beta_{l,m}$ , and  $\gamma_{l,m}$ . For ease of notation, we collect the spherical-harmonic coefficients into a vector  $\mathbf{z} = (\alpha_{l,m}, \beta_{l,m}, \gamma_{l,m})$ , where  $m \in \{-l, -l+1, \dots, l\} \forall l \in \{1, 2, \dots, l_{\max}\}$ . The hyperparameter  $l_{\max}$  specifies the truncation of the expansion.

### 2.2. Forward Model

Spectropolarimetry allows for measurement of not only the total intensity spectrum (Stokes I), but also the circular polarization spectrum (Stokes V), which is sensitive to the line-of-sight component of the stellar magnetic field vector. Given a set of spherical-harmonic coefficients, our forward model produces a set of synthetic Stokes V profiles. The process is based on the integration of local line profiles across the visible stellar surface on a discretized surface grid. For our purposes, this simulation process is referred to as a deterministic forward simulator denoted  $\mathbf{f}(\mathbf{z}) \in \mathbb{R}^N$ , where  $N$  is the observation profile dimension. We model the local line profiles by the Unno–Rachkovsky (UR) analytical solution of the polarized radiative transfer equation in the Milne–Eddington atmosphere (Landi Degl’Innocenti & Landolfi, 2004), using a Gaussian absorption profile. Disk-integration

of the weighted local line profiles is carried out over the visible surface elements at each observational phase, producing the corresponding synthetic Stokes V profiles.

### 2.3. Point Estimate of Field Reconstruction

To obtain a point estimate of the surface magnetic field map, standard ZDI inversion is often performed by solving the weighted nonlinear least-squares problem given in Equation (2), with a regularizer  $r(\mathbf{z}) = \eta \sum_{l=1}^{l_{\max}} \sum_{m=-l}^l l^2 (\alpha_{l,m}^2 + \beta_{l,m}^2 + \gamma_{l,m}^2)$  (see, e.g., Kochukhov, 2016). The regularization strength is denoted by  $\eta$ , and each term in the regularization function corresponds to the  $l^2$ -weighted magnitudes of the magnetic energies of the spherical-harmonic coefficients. Since ZDI inversion using spectropolarimetric information only in Stokes V is an ill-posed problem, the regularization is required in order to ensure convergence to a stable and unique solution. Formally, if  $\mathbf{y}$  is the observed Stokes V profile and  $\Lambda$  is the precision matrix of the measurement noise, then the model parameters  $\mathbf{z}$  are given by the solution to the following optimization problem:

$$\mathbf{z} = \arg \min_{\mathbf{z}} \|\Lambda^{\frac{1}{2}}(\mathbf{y} - \mathbf{f}(\mathbf{z}))\|_2^2 + r(\mathbf{z}). \quad (2)$$

## 3. Probabilistic ZDI

By extending the least-squares formulation of the spectropolarimetric inversion problem to a Bayesian nonlinear regression setting, we can quantify the magnetic field uncertainty following the standard Bayesian framework. This section details the statistical model and the proposed inference approach.

### 3.1. Statistical Model

In our Bayesian framework,  $\mathbf{z}$  is treated as a latent random variable, and the posterior distribution  $p(\mathbf{z}|\mathbf{y})$  is given by  $p(\mathbf{z}|\mathbf{y}) = p(\mathbf{y}|\mathbf{z})p(\mathbf{z})/p(\mathbf{y})$ . The corresponding Bayesian network is illustrated in Figure 1, together with a schematic illustration of our inference pipeline. Given the field parameterization provided in Equation (1), the posterior magnetic field distribution can be obtained from a linear transformation of random variable  $\mathbf{z}' \sim p(\mathbf{z}|\mathbf{y})$ . Similar to standard ZDI described in Section 2.3, we assume the observational noise to be Gaussian with diagonal covariance. Thus, our likelihood model becomes  $p(\mathbf{y}|\mathbf{z}) = \mathcal{N}(\mathbf{y}; \mathbf{f}(\mathbf{z}), \Lambda^{-1})$ . The marginal likelihood  $p(\mathbf{y})$  is, however, intractable in general, due to the nonlinearity of  $\mathbf{f}(\mathbf{z})$ . As for the prior distribution, we use  $p(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \omega^{-1}I)$ , where  $\omega^{-1}I$  is a diagonal matrix.

Although it is difficult to construct physics-informed priors on the spherical-harmonic coefficients, we can design priors

that capture the ZDI regularization structure described in the previous section, detailed by, e.g., Kochukhov et al. (2014). More specifically, we choose the parameter  $\omega_i = \eta l_i^2$  for the marginal prior distribution of each spherical-harmonic coefficient  $z_i$ . The maximum a posteriori estimate  $\hat{\mathbf{z}}_{\text{MAP}} = \arg \max_{\mathbf{z}} p(\mathbf{z}|\mathbf{y})$  then solves the the weighted least-squares problem detailed in Equation (2). Thus, if the mean of the full posterior distribution happens to coincide with the MAP estimate, the Bayesian inference provides an uncertainty quantification around the point estimate obtained from standard ZDI.

### 3.2. Computing the Posterior Distribution

Monte Carlo methods are commonly used to approximate the posterior distribution when it is not analytically tractable, as is often the case for forward mapping functions modeling the relationship between spherical-harmonic coefficients and spectropolarimetric time series. To draw samples from the posterior distribution, MCMC methods can be employed. Simple MCMC algorithms are often limited by poor sample efficiency, preventing efficient exploration of complex and high-dimensional state spaces. Since the spherical-harmonic field parameterization in Equation (1) produces a high-dimensional posterior distribution, we use the gradient-based Metropolis-adjusted Langevin algorithm (MALA, Brooks et al., 2011) for inference, achieving significant improvement in the mixing of the chain. MALA is based on a Langevin diffusion process and uses the gradient of the target log-posterior to form efficient proposal distributions under the hood of the Metropolis-Hastings algorithm, see, e.g., Luengo et al. (2020) for details.

### 3.3. Differentiable Simulator in JAX

We implement our forward simulator  $\mathbf{f}(\mathbf{z})$  in JAX, which makes the MCMC inference computationally feasible for this application. The main advantage is that JAX offers just-in-time (JIT) compilation of Python code, resulting in significant hardware acceleration on CPUs. The accelerated linear algebra (XLA) based compilation also allows for effortless GPU and TPU extensions. In addition, XLA is combined with an updated Autograd functionality, allowing for automatic differentiation of the log-posterior, removing the need for analytical derivation, implementation, or explicit numerical approximation of the gradients,  $\nabla_{\mathbf{z}} \log(p(\mathbf{z}|\mathbf{y}))$ . In light of these attractive properties, JAX provides a convenient framework for implementation of complex MCMC algorithms incorporating gradient information, while supporting flexibility in terms of forward mapping functions and prior distributions. To sample from the posterior distribution, we use the JAX-based library BlackJAX (Lao & Louf, 2020), which offers a robust baseline implementation of MALA.




Figure 2. The two columns to the left show the analytical MAP magnetic field estimates across the stellar surface obtained using the weak field approximation, and the corresponding standard deviation maps. The magnetic field maps correspond to the radial (top), meridional (middle) and azimuthal (bottom) magnetic field components (kG). The middle columns show the MAP magnetic field map estimates and standard deviation maps obtained using the nonlinear UR-solution. The right-most column shows the corresponding model Stokes V spectra (red) compared to observations (black). These profiles are offset vertically according to the corresponding rotational phase.

### 3.4. Probabilistic ZDI under the Weak Field Approximation

For many stars, in particular solar-like stars and the majority of cool, low-mass stars, the weak field approximation (Landi Degl’Innocenti & Landolfi, 2004) applies. In this special case, 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  $\mathbf{A}$  is the transformation matrix. When this is the case, our choice of prior is conjugate to the likelihood, and an analytical expression for the Gaussian posterior distribution  $p(\mathbf{z}|\mathbf{y})$  exists, such that  $p(\mathbf{z}|\mathbf{y}) = \mathcal{N}(\mathbf{z}; \mu, \Sigma)$  with  $\Sigma^{-1} = \omega\mathbf{I} + \mathbf{A}^T\Lambda\mathbf{A}$  and  $\mu = \Sigma(\mathbf{A}^T\Lambda\mathbf{y})$ . Since the magnetic field  $B$  is a linear function of  $\mathbf{z}$  according to Equation (1), and  $p(\mathbf{z}|\mathbf{y})$  is Gaussian, the posterior magnetic field distribution is also Gaussian in this case, and can be expressed in closed form.

## 4. Results

We use high-quality mean Stokes I and Stokes V spectra based on the observational data of  $\tau$  Sco analyzed by Kochukhov & Wade (2016). The original data consists of 49 circular polarization observations recorded with the ESPaDOnS instrument (Donati et al., 2006) at the 3.6m Canada-France-Hawaii telescope between 2004 and 2009. Following Kochukhov & Wade (2016), we fix the rotational period to 41.033 days, the projected rotational velocity to  $6 \text{ km s}^{-1}$ , and the inclination angle to 70 degrees. For further details on the data and preprocessing, readers are referred to Kochukhov & Wade (2016).

We let  $l_{\text{max}} = 10$  in Equation (1), resulting in 360 spherical-harmonic coefficients. With our implementation of the

proposed scheme, we achieve stable and fast convergence. Compared to the corresponding implementation in pure NumPy, where we use JAX for gradient calculations only, our implementation achieves a runtime speedup of 648x on a Tesla A100 GPU. We run the MALA sampler for  $3 \cdot 10^6$  iterations, and remove a burn-in of 500k samples. Figure 2 shows the resulting MAP estimate of the surface magnetic field maps and the corresponding Stokes V fit obtained using the forward simulator described in Section 2.2, together with the obtained standard deviation of the magnetic field across the stellar surface. The analytical MAP estimate of the magnetic field maps obtained using the weak field approximation and corresponding uncertainty are also included in Figure 2. Since the weak field approximation applies, the obtained field maps in the two cases are expected to be similar. The presented point estimates in both cases coincide with the magnetic field maps reconstructed using standard ZDI, and are similar to the previously published magnetic field maps of  $\tau$  Sco obtained by Kochukhov & Wade (2016). Our results show that the uncertainty is latitude-dependent in all three components, which is reasonable due to the stellar inclination with respect to the observer. The mean deviation between the MAP fits and observed Stokes V profiles is  $9.17 \cdot 10^{-5}$  (weak field) and  $1.15 \cdot 10^{-4}$  (UR), which is comparable to previously published point estimates.

## 5. Conclusions and Future Work

We present a Bayesian extension of standard ZDI using the high-dimensional spherical-harmonic field parameterization, allowing for accurate uncertainty quantification of stellar surface magnetic field maps where current research

is dominated by point estimates. By efficient implementation in JAX, we achieve a significant runtime speedup and show that fast and stable convergence is possible using MCMC methods. Our results formalize empirical uncertainty quantification in stellar magnetic field maps, and can be extended to a variety of ZDI targets. In future work, we plan to expand the hierarchy in the statistical model to account for uncertainties in multiple stellar parameters. This allows for uncertainty quantification of particularly challenging ZDI targets, including equator-on hosts of transiting exoplanets and stars in eclipsing binary systems, where spherical-harmonic coefficients may exhibit degeneracy.

## Acknowledgements

This research was partially supported by the *Wallenberg Artificial Intelligence, Autonomous Systems and Software Program (WASP)*, funded by Knut and Alice Wallenberg Foundation. The computations were enabled by the Berzelius resource provided by the Knut and Alice Wallenberg Foundation at the National Supercomputer Centre.

## References

- Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. Variational inference: a review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- Boro Saikia, S., Lüftinger, T., Folsom, C. P., Antonova, A., Alecian, E., Donati, J. F., Guedel, M., Hall, J. C., Jeffers, S. V., Kochukhov, O., Marsden, S. C., Metodieva, Y. T., Mittag, M., Morin, J., Perdelwitz, V., Petit, P., Schmid, M., and Vidotto, A. A. Time evolution of magnetic activity cycles in young suns: the curious case of  $\kappa$  Ceti. *Astronomy & Astrophysics*, 658:A16, 2022.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- Brooks, S., Gelman, A., Jones, G., and Meng, X.-L. (eds.). *Handbook of Markov chain Monte Carlo*. Chapman & Hall/CRC Handbooks of Modern Statistical Methods. CRC Press/Taylor & Francis, 2011.
- Donati, J. F., Catala, C., Landstreet, J. D., and Petit, P. ES-PaDONs: the new generation stellar spectro-polarimeter. Performances and first results. *Astronomical Society of the Pacific Conference Series*, 358:362–368, 2006.
- Donati, J. F., Howarth, I. D., Jardine, M. M., Petit, P., Catala, C., Landstreet, J. D., Bouret, J. C., Alecian, E., Barnes, J. R., Forveille, T., Paletou, F., and Manset, N. The surprising magnetic topology of  $\tau$  Sco: fossil remnant or dynamo output? *Monthly Notices of the Royal Astronomical Society*, 370(2):629–644, 2006.
- Folsom, C. P., Bouvier, J., Petit, P., Lèbre, A., Amard, L., Palacios, A., Morin, J., Donati, J.-F., and Vidotto, A. A. The evolution of surface magnetic fields in young solar-type stars II: the early main sequence (250–650 Myr). *Monthly Notices of the Royal Astronomical Society*, 474:4956–4987, 2018.
- Fréour, L., Neiner, C., Landstreet, J. D., Folsom, C. P., and Wade, G. A. The magnetic field of the chemically peculiar star V352 Peg. *Monthly Notices of the Royal Astronomical Society*, 520(3):3201–3217, 2023.
- Grunhut, J. H., Wade, G. A., Folsom, C. P., Neiner, C., Kochukhov, O., Alecian, E., Shultz, M., Petit, V., MiMeS Collaboration, and BinaMiCS Collaboration. The magnetic field and magnetosphere of Plaskett’s star: a fundamental shift in our understanding of the system. *Monthly Notices of the Royal Astronomical Society*, 512(2):1944–1966, 2022.
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., and Saul, L. K. An introduction to variational methods for graphical models. *Machine Learning*, 37(2):183–233, 1999.
- Kochukhov, O. Doppler and zeeman doppler imaging of stars. In Rozelot, J.-P. and Neiner, C. (eds.), *Cartography of the Sun and the stars*, pp. 177–204. Springer International Publishing, 2016.
- Kochukhov, O. and Wade, G. A. Magnetic field topology of  $\tau$  Scorpii. The uniqueness problem of Stokes V ZDI inversions. *Astronomy & Astrophysics*, 586:A30, 2016.
- Kochukhov, O., Lüftinger, T., Neiner, C., Alecian, E., and MiMeS Collaboration. Magnetic field topology of the unique chemically peculiar star CU Virginis. *Astronomy & Astrophysics*, 565:A83, 2014.
- Kochukhov, O., Shultz, M., and Neiner, C. Magnetic field topologies of the bright, weak-field Ap stars  $\theta$  Aurigae and  $\epsilon$  Ursae Majoris. *Astronomy & Astrophysics*, 621:A47, 2019.
- Kochukhov, O., Papakonstantinou, N., and Neiner, C. Magnetic field topology, chemical spot distributions, and photometric variability of the Ap star  $\varphi$  Draconis. *Monthly Notices of the Royal Astronomical Society*, 510(4):5821–5833, 2022.
- Landi Degl’Innocenti, E. and Landolfi, M. *Polarization in Spectral Lines*. Springer, 2004.

- Lao, J. and Louf, R. Blackjax: a sampling library for JAX, 2020. URL <http://github.com/blackjax-devs/blackjax>.
- Luengo, D., Martino, L., Bugallo, M., Elvira, V., and Särkkä, S. A survey of monte carlo methods for parameter estimation. *EURASIP Journal on Advances in Signal Processing*, 25, 2020.
- Marsden, S. C., Evensberger, D., Brown, E. L., Neiner, C., Seach, J. M., Morin, J., Petit, P., Jeffers, S. V., and Folsom, C. P. The magnetic field and stellar wind of the mature late-F star  $\chi$  Draconis A. *Monthly Notices of the Royal Astronomical Society*, 522(1):792–810, 2023.
- Petit, P., Böhm, T., Folsom, C. P., Lignières, F., and Cang, T. A decade-long magnetic monitoring of Vega. *Astronomy & Astrophysics*, 666:A20, 2022.
- Petit, V. and Wade, G. A. Stellar magnetic field parameters from a Bayesian analysis of high-resolution spectropolarimetric observations. *Monthly Notices of the Royal Astronomical Society*, 420(1):773–791, 2011.
- Piskunov, N. and Kochukhov, O. Doppler Imaging of stellar magnetic fields. I. Techniques. *Astronomy & Astrophysics*, 381:736–756, 2002.