Inverting Solar Spectropolarimetric Observations with Deep Learning

S. Curt Dodds
Institute for Astronomy
University of Hawaiʻi at Mānoa
Honolulu, HI 96822
dodds@hawaii.edu

Ian Cunnyngham
Institute for Astronomy
University of Hawaiʻi at Mānoa
Honolulu, HI 96822
ianmc@hawaii.edu

Lucas Tarr
National Solar Observatory
Makawao, HI 96768
ltarr@nso.edu

Sarah Jaeggli
National Solar Observatory
Makawao, HI 96768
sjaeggli@nso.edu

Tom Schad
National Solar Observatory
Makawao, HI 96768
tschad@nso.edu

Peter Sadowski
Information and Computer Sciences
University of Hawaiʻi at Mānoa
Honolulu, HI 96822
peter.sadowski@hawaii.edu

Xudong Sun
Institute for Astronomy
University of Hawaiʻi at Mānoa
Honolulu, HI 96822
xudongs@hawaii.edu

Abstract

The upcoming *Daniel K. Inouye Solar Telescope* (*DKIST*) will produce unprecedented high-cadence, high-resolution, and multi-line spectropolarimetric observations of the Sun. New computational techniques are needed to fully exploit these measurements to address long-standing questions regarding the physical processes that govern solar magnetohydrodynamics (MHD). We investigate deep learning as a method for approximate inference of the photosphere in 4D (3D spatial volume over time). While standard approaches perform iterative inference procedures on individual pixels, a deep neural network could more quickly and accurately predict the state of the photosphere by taking advantage of spatial and temporal structure. We demonstrate this approach by training on both MHD simulations and observational data from the Hinode Solar Optical Telescope.

1 Introduction

The solar photosphere hosts a wide range of plasma and magnetic conditions. Many phenomena in this layer, including magnetic field emergence, convection, and energy transfer/conversion, are of fundamental importance in solar physics. These phenomena are generally well-described by the magnetohydrodynamic (MHD) equations that describe the evolution of the photospheric *state variables*: the magnetic field $\boldsymbol{B}$, velocity field $\boldsymbol{v}$, density ρ, and pressure p (or, redundantly, temperature T) throughout the 3D volume. Because these quantities determine how photons are emitted, absorbed, and re-emitted, solar physicists infer them from the distinct spectral bands and polarization patterns (Stokes profiles) observed by solar telescopes.

The *Daniel K. Inouye Solar Telescope* [1], scheduled to begin operations in 2020, will produce unprecedented spectropolarimetric observations. The 4-m clear aperture is the largest in the world; favorable seeing conditions atop the Haleakalā volcano in Hawaiʻi and advanced adaptive optics

Second Workshop on Machine Learning and the Physical Sciences (NeurIPS 2019), Vancouver, Canada.

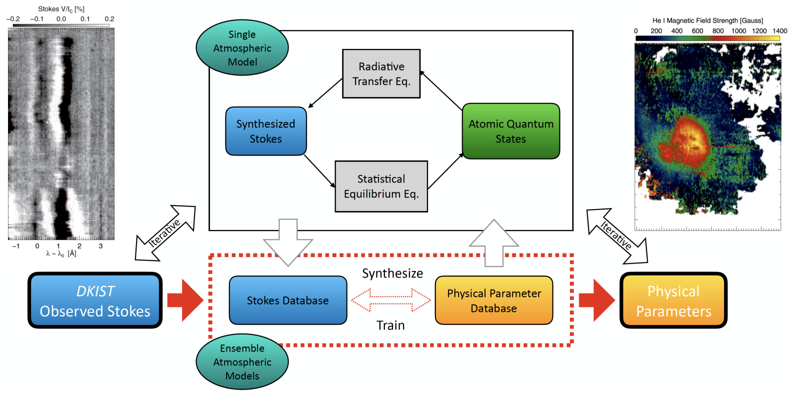

Figure 1: Stokes inversion for a single frame. The input is Stokes parameters for each pixel in a 2D map (left); the output is a sequence of 3D MHD state variables (right). Traditional iterative inversion process is illustrated by the solid box (top). The Deep Learning method is illustrated by the red dotted box (bottom), where a convolutional neural network is trained and tested with realistic MHD simulations. The proposed 4D inversion will train on a time sequence of Stokes profiles.

will enable diffraction-limited observations ($0.03''$ or 23 km at $\lambda = 500$ nm), with the polarization signal measured at an accuracy better than 0.1%. In particular, the Diffraction-Limited Near-InfraRed Spectro-Polarimeter (DL-NIRSP) will provide multi-wavelength, high-cadence spectropolarimetric observations *simultaneously* over a continuous field of view (FoV). Solar physicists hope to use this data to infer the MHD state variables in 4D (three spatial dimensions plus time) and their derivatives to help address many long-standing questions in solar physics such as: (1) What is the 3D structure of sunspots, e.g., in light bridges and penumbral filaments? (2) How do magnetic and velocity fields evolve during flux emergence and cancellation? (3) Can we estimate the vector electric current density, electric field, and Poynting flux?

Inference of the MHD state variables from spectropolarization observations is typically performed using iterative "inversion" algorithms (Figure 1). In general, these inversion algorithms start with an initial guess of the solar state variables, and a physics model is used to calculate Stokes profiles that we would expect to observe on Earth. These profiles are compared to observation data, and the error is used to update the estimate of the solar state variables — iterating until convergence. This inversion is essentially a non-convex optimization problem — with convergence difficulties, ambiguities due to multiple local minima, and computationally expensive evaluations of the forward model. Thus, the standard approaches make a number of simplifying assumptions: (1) they *invert each pixel and time step independently*, without capturing the spatial/temporal dependencies; (2) they invert each spectral line independently, without capturing the dependencies between different depths of the solar atmosphere; (3) the depth dimension of the Sun's atmosphere is often modeled at very low resolution.

Even with these simplifications, the fastest iterative inference algorithms are *extremely computationally demanding* and will have trouble keeping up with the *DKIST* data stream. DL-NIRSP alone can generate high dispersion polarized spectral profiles at a rate of 10^4 s^{-1}. Highly-optimized simple inversions, e.g. VFISV [2], will still require > 50 CPU hours per 3 hour observation while not utilizing the additional diagnostic power afforded by the high spectral dispersion and accuracy of DKIST observables. Inversions utilizing more appropriate atmospheric models, e.g. SIR [3], require 2 to 3 orders of magnitude longer computing times. **Thus, new computational methods are needed to meet the demands of modern solar astronomy.**

Data-driven methods offer a promising approach for both faster and more accurate inference. The high-dimensional data contains spatial and temporal structure that can be exploited by data-driven models trained on MHD simulations or the states inferred from historical observations by slow inversion algorithms. Previous works have employed compressed sensing [4] and lookup tables built from snapshots of MHD simulations [5]. We believe a more promising approach is *deep learning*, which has become an important approach for approximate inference in physics models where high-dimensional data is generated by a known physical process [6, 7]. The Stokes inversion problem is a natural application of deep learning given the large data volume, the spatio-temporal structure of the data, and the speed bottleneck facing solar physicists. The feasibility of this application was first demonstrated by the pioneering work of [8] using a neural network with convolution over two spatial dimensions to predict images one at a time. However, to our knowledge, no existing inversion methods have exploited the full spatial, spectral, and temporal structure of the data.

Numerical MHD simulations are now at an advanced state where many solar features are reproduced faithfully from first principles [e.g., 9, 10, 11]. As a result, they are already being utilized extensively to guide the interpretation of the observed spectropolarimetric signatures, or as a test bed for various theories. **The ability to simulate this physical process presents an opportunity to apply data-intensive methods such as deep learning for fast approximate inference.** In this extended abstract we demonstrate that the deep learning approach is able to accurately infer major structures in both sunspots and quiet sun, and confirm the hypothesis that accuracy can be increased by using more spatial and spectral information.

2 Experiments

Below, we train deep learning approximate inference models using two separate labeled data sets. The first consists of historical data from the Hinode Solar Optical Telescope for which inversions have already been calculated using standard methods. The second is created from existing MHD simulations, for which we have to simulated the forward physics model to obtain the Stokes profiles.

2.1 Training on Hinode Observations

Historical data from Hinode SOT SP [12] contains Level 1 spectropolarimetric observations (Stokes parameters) paired with Level 2 photosphere states inferred by the MERLIN inversion code [13]. The Level 2 pixel-wise inversions are not quite ground truth, but they are accurate enough to capture many of the major features of solar dynamics, and are used by solar astronomers to interpret solar behavior. Because computing these pixel-wise inversions is computationally intensive, a fast approximation would be valuable in itself so that solar astronomers could visualize the structure of the Sun in real time.

Our training set was selected from 9.2TB of available Hinode data. Level 1 data was standardized by subtracting the mean and dividing by the standard deviation, then each image was split into non-overlapping image *patches* of $64 \times 64 \approx 0.3$ arcsec pixels. For simplicity, we discarded all patches outside the central $2/3$ of the solar disc, for which the Sun's curvature becomes significant. Of 24,323 patches, 80% were used for training, 10% for validation, and 10% for testing.

The deep neural network maps four Stokes parameters (I, Q, U, V) for 112 wavelengths at 64×64 pixels to the strength, inclination, and azimuth of the magnetic field vector (B, γ, ψ) at each pixel (also 64×64). We experimented with different variations of a fully convolutional U-Net model [14] including both 2D and 3D convolution layers. The best results were obtained with a 22 layer model: four convolution modules (each containing two 3D-convolution layers with ReLU activation and 3x3x3 kernel shape; then 5% dropout; then 3D max-pooling with stride of 2x2x2), followed by four residual transpose convolution modules (one 3D conv-transpose with skip connections, two 3x3x3 Conv3D layers with ReLU activation function, then 5% dropout), and finally a 1x1x112 linear 3D-convolution layer. The mean absolute error (MAE) was minimized using ADAM with a learning rate of 1e-5. Training was performed on a single NVIDIA V100 GPU for 131 epochs, at which point no improvement was seen on the validation set.

Once trained, the CNN inversion is orders of magnitude faster ($\approx 7.4 \times 10^4$ pixel/sec) than current inversion methods such as VFISV ($\approx 6 \times 10^2$ pixel/sec) [2], and accurately approximates the slower algorithm. The inversion examples in Figures 2 and 3 show that the major features of sun spots can be clearly seen in these fast CNN inversions.

2.2 Training on MHD Simulations

DKIST's multi-line observations should enable us to cover a wider range of physical heights and better constrain the vertical gradient of MHD variables. The latter will improve our understanding of several important questions, for example, the horizontal electric current density [15] and the degree of force-freeness of the photosphere [16]. Initial line candidates include the well studied Fe I 630.2

and Fe I 1565 nm, which will be simultaneously observed with DL-NIRSP, with possible extension to the He I/Si I lines near 1083 nm.

We confirm that observing multiple lines improves performance in the deep learning approach. Starting from the sunspot simulation from [9], Stokes observations of both pairs of Fe I 630.2 and Fe I 1565 nm lines were synthesized using the software framework of [8]. A deep neural network was trained on 50% of the simulation region, while 25% was used for model development and hyperparameter-optimization with SHERPA[17], and the final 25% was set aside for testing. The best neural network architecture, a six-layer convolutional network with 175 M parameters, was then trained with two different sets of inputs: (1) only the Fe I 630.2 lines; (2) both the Fe I 630.2 and Fe I 1565 nm lines. The result show a clear performance improvement using both sets of lines (Table 1).

3 Discussion and Conclusion

Deep learning models trained for approximate inference will simultaneously learn both how to invert the forward model *and* a structured prior over the photosphere state variables. Thus, this data-driven approach depends heavily on the training data. MHD simulations generated from first principles are appealing because they provide ground truth photosphere states, but they are few in number due to the computational expense of producing them. While we have demonstrated that training on only two simulations can yield a model that generalizes well on a held-out test set, it should be possible to

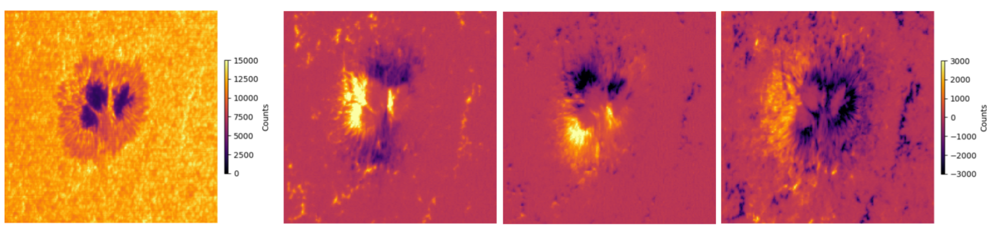

Figure 2: Image of Hinode SOT sunspot (≈ 1 arcmin^2) from test set. From left to right: observed Stokes parameters I, Q, U, V.

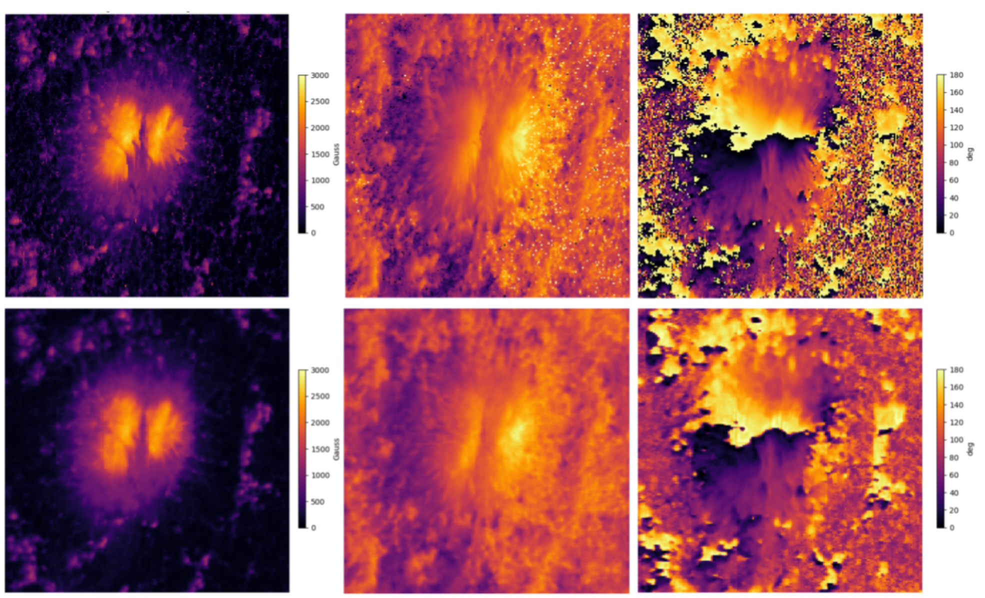

Figure 3: Traditional inversion using MERLIN (top) and deep learning inversion (bottom) for the test sunspot in Figure 2. From left to right: the inferred strength, inclination, and azimuth of the magnetic field vectors.

Property		Quiet Sun Single line	Quiet Sun Multi-line	Umbra / Penumbra Single line	Umbra / Penumbra Multi-line
Height	km	16.60	16.39	18.88	17.64
Temperature	K	77.28	80.10	74.83	79.37
Pressure	Pa	1244	909.8	1639	1489
Velocity Z	m/s	525.2	462.7	312.3	271.4
Magnetic field Q	$gauss$	110.5	95.41	235.0	157.2
Magnetic field U	$gauss$	79.05	67.68	168.8	121.3
Magnetic field Z	$gauss$	113.0	83.21	148.2	108.7

Table 1: Mean absolute error of two CNN inversion models trained using single (Fe I 630.2nm) and multiple lines (Fe I 630.2nm and Fe I 1565nm). The multi-line model outperforms the single-line model for six of the seven properties; the small performance loss in temperature is due to over-fitting on that task. Test set performance is shown separately for quiet sun regions and umbra/penumbra regions at $\log\tau = -3.0$.

make use of *both* simulations and Hinode inversions — perhaps by pre-training a model on the much larger Hinode data and fine-tuning on simulations — and we are investigating this approach.

DKIST will be the center of ground-based solar observation for several decades. The *DKIST* Level-2 effort, being implemented by the National Solar Observatory (NSO) for the community, will initially concentrate on providing routine inversion data products for a fraction of DKIST observations focused on chromospheric and coronal science use cases only at particular snapshots in time, rather than at the high time cadence afforded by photospheric use cases. This deep learning approach could accelerate scientific progress by providing faster, cheaper, larger and more accurate data products to the scientific community.

Acknowledgements

The work was made possible by hardware grants from NVIDIA to CD and PS, along with technical support and advanced computing resources from the University of Hawai'i Information Technology Services Cyberinfrastructure group. Alfred de Wijn of the High Altitude Observatory, National Center for Atmospheric Research, assisted us with downloading the Hinode data set.

References

[1] Thomas R. Rimmele, Valentin Martinez Pillet, Philip R. Goode, Michael Knoelker, Jeffrey Richard Kuhn, Robert Rosner, Roberto Casini, Haosheng Lin, Oskar von der Luehe, Friedrich Woeger, Alexand ra Tritschler, Andre Fehlmann, Sarah A. Jaeggli, Wolfgang Schmidt, Alfred De Wijn, Mark Rast, David M. Harrington, Stacey R. Sueoka, Christian Beck, Thomas A. Schad, Mark Warner, Joseph P. McMullin, Steven J. Berukoff, Mihalis Mathioudakis, and DKIST Team. Status of the Daniel K. Inouye Solar Telescope: unraveling the mysteries the Sun. In *American Astronomical Society Meeting Abstracts #232*, volume 232 of *American Astronomical Society Meeting Abstracts*, page 316.01, Jun 2018.

[2] J. M. Borrero, S. Tomczyk, M. Kubo, H. Socas-Navarro, J. Schou, S. Couvidat, and R. Bogart. VFISV: Very Fast Inversion of the Stokes Vector for the Helioseismic and Magnetic Imager. *Solar Physics*, 273(1):267–293, Oct 2011.

[3] B. Ruiz Cobo and J. C. del Toro Iniesta. SIR: Stokes Inversion based on Response functions. Astrophysics Source Code Library, December 2012.

[4] A. Asensio Ramos and J. de la Cruz Rodríguez. Sparse inversion of Stokes profiles. I. Two-dimensional Milne-Eddington inversions. *Astron. Astrophys.*, 577:A140, May 2015.

[5] T. L. Riethmüller, S. K. Solanki, P. Barthol, A. Gand orfer, L. Gizon, J. Hirzberger, M. van Noort, J. Blanco Rodríguez, J. C. Del Toro Iniesta, D. Orozco Suárez, W. Schmidt, V. Martínez Pillet, and M. Knölker. A New MHD-assisted Stokes Inversion Technique. *Astrophys. J. Suppl.*, 229(1):16, Mar 2017.

[6] Peter Sadowski and Pierre Baldi. Deep learning in the natural sciences: applications to physics. In *Braverman Readings in Machine Learning. Key Ideas from Inception to Current State*, pages 269–297. Springer, 2018.

[7] Johann Brehmer, Gilles Louppe, Juan Pavez, and Kyle Cranmer. Mining gold from implicit models to improve likelihood-free inference. *arXiv preprint arXiv:1805.12244*, 2018.

[8] A. Asensio Ramos and C. J. Díaz Baso. Stokes inversion based on convolutional neural networks. *Astron. Astrophys.*, 626:A102, Jun 2019.

[9] M. Rempel. Numerical Sunspot Models: Robustness of Photospheric Velocity and Magnetic Field Structure. *Astrophys. J.*, 750(1):62, May 2012.

[10] M. C. M. Cheung, M. Rempel, A. M. Title, and M. Schüssler. Simulation of the Formation of a Solar Active Region. *Astrophys. J.*, 720(1):233–244, Sep 2010.

[11] M. C. M. Cheung, M. Rempel, G. Chintzoglou, F. Chen, P. Testa, J. Martínez-Sykora, A. Sainz Dalda, M. L. DeRosa, A. Malanushenko, V. Hansteen, B. De Pontieu, M. Carlsson, B. Gudiksen, and S. W. McIntosh. A comprehensive three-dimensional radiative magnetohydrodynamic simulation of a solar flare. *Nat. Astron.*, 3:160–166, Jan 2019.

[12] B. W. Lites, D. L. Akin, G. Card, T. Cruz, D. W. Duncan, C. G. Edwards, D. F. Elmore, C. Hoffmann, Y. Katsukawa, N. Katz, M. Kubo, K. Ichimoto, T. Shimizu, R. A. Shine, K. V. Streander, A. Suematsu, T. D. Tarbell, A. M. Title, and S. Tsuneta. The Hinode Spectro-Polarimeter. *Solar Physics*, 283:579–599, April 2013.

[13] B. W. Lites and K. Ichimoto. The SP_PREP Data Preparation Package for the Hinode Spectro-Polarimeter. *Solar Physics*, 283:601–629, April 2013.

[14] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention*, volume 9351, pages 234–241, Nov 2015.

[15] K. G. Puschmann, B. Ruiz Cobo, and V. Martínez Pillet. The Electrical Current Density Vector in the Inner Penumbra of a Sunspot. *Astrophys. J. Lett.*, 721(1):L58–L61, Sep 2010.

[16] Thomas R. Metcalf, Litao Jiao, Alexander N. McClymont, Richard C. Canfield, and Han Uitenbroek. Is the Solar Chromospheric Magnetic Field Force-free? *Astrophys. J.*, 439:474, Jan 1995.

[17] Lars Hertel, Julian Collado, Peter Sadowski, and Pierre Baldi. Sherpa: hyperparameter optimization for machine learning models, 2018.