












Accurate Machine Learning Atmospheric Retrieval via a Neural Network Surrogate Model for Radiative Transfer

MICHAEL D. HIMES ¹, JOSEPH HARRINGTON ¹, ADAM D. COBB ², ATILIM GÜNEŞ BAYDIN ²,
FRANK SOBOCZENSKI ³, MOLLY D. O'BEIRNE ⁴, SIMONE ZORZAN ⁵, DAVID C. WRIGHT ¹, ZACCHAEUS SCHEFFER ¹,
SHAWN D. DOMAGAL-GOLDMAN ⁶, AND GIADA N. ARNEY ⁶

¹*Planetary Sciences Group, Department of Physics, University of Central Florida*

²*Department of Engineering Science, University of Oxford*

³*SPHES, King's College London*

⁴*Department of Geology and Environmental Science, University of Pittsburgh*

⁵*ERIN Department, Luxembourg Institute of Science and Technology*

⁶*NASA Goddard Space Flight Center, Greenbelt, MD*

(Received March 4, 2020; Revised —; Accepted —)

Submitted to PSJ

ABSTRACT

Atmospheric retrieval determines the properties of an atmosphere based on its measured spectrum. The low signal-to-noise ratio of exoplanet observations require a Bayesian approach to determine posterior probability distributions of each model parameter, given observed spectra. This inference is computationally expensive, as it requires many executions of a costly radiative transfer (RT) simulation for each set of sampled model parameters. Machine learning (ML) has recently been shown to provide a significant reduction in runtime for retrievals, mainly by training inverse ML models that predict parameter distributions, given observed spectra, albeit with reduced posterior accuracy. Here we present a novel approach to retrieval by training a forward ML surrogate model that predicts spectra given model parameters, providing a fast approximate RT simulation that can be used in a conventional Bayesian retrieval framework without significant loss of accuracy. We demonstrate our method on the emission spectrum of HD 189733 b and find Bhattacharyya coefficients of 97.74 – 99.74% between our 1D marginalized posterior distributions and those of the Bayesian Atmospheric Radiative Transfer (BART) code. Our retrieval method is $\sim 20\times$ faster than BART when run on an Intel i7-4770 central processing unit (CPU). Neural-network computation using an NVIDIA Titan Xp graphics processing unit is $\sim 600\times$ faster than BART on that CPU.

Keywords: techniques: retrieval — techniques: machine learning — methods: statistical — planets and satellites: atmospheres — planets and satellites: individual (HD 189733 b)

1. INTRODUCTION

Over the past decades, exoplanet studies have expanded from their detection to include characterization of their atmospheres via retrieval (Deming & Seager 2017). Retrieval is the inverse modeling technique whereby forward models of a planet's spectrum are compared to observational data in order to constrain the

model parameters (Madhusudhan 2018). These typically include the shape of the thermal profile, abundances of species, and condensate properties. While some solar-system objects can be characterized with simpler approaches (such as Levenberg-Marquardt minimization) due to their high signal-to-noise ratios (e.g., Koskinen et al. 2016), retrieval on noisy exoplanet spectra require Bayesian methods to provide a distribution of models that can explain the observed data. The posterior distribution resulting from a Bayesian retrieval places limits on each model parameter (constrained within some range, an upper or lower limit, or uncon-

Corresponding author: Michael D. Himes
mhimes@knights.ucf.edu

strained), informing the statistical significance of the result.

Bayesian retrieval methods involve evaluating thousands to millions of spectra, integrating over the observational bandpasses, and comparing to observations. Depending on model complexity, this requires hundreds to thousands of parallelizable compute hours, resulting in hours to days of runtime. Calculating the model spectra by solving the radiative transfer (RT) equation takes the vast majority of compute time.

Machine learning (ML) encompasses algorithms that learn representations of and uncover relationships within a collection of data samples. Deep learning (Goodfellow et al. 2016) is a subfield of ML that is based on neural networks (NNs), which are highly-flexible differentiable functions that can be fit to data. NNs can classify images (He et al. 2016), recognize speech (Amodei et al. 2016), and translate between languages (Wu et al. 2016). NNs consist of a hierarchy of layers that contain nodes performing weighted non-linear transformations of their inputs, through a series of hidden layers, to a desired output. For example, for a retrieval, one might have the input layer receive the observed spectrum, hidden layers extract features, and the output layer predict the underlying atmospheric parameters. NN training conventionally uses gradient-based optimization, iteratively adjusting the weights of the connections between nodes to minimize the error between the NN’s prediction and the desired output (Rumelhart et al. 1986).

Recent applications of ML to atmospheric retrieval reduced compute time from hundreds of hours to minutes or less. Márquez-Neila et al. (2018, HELA) presented a random forest of regression trees to build predictive distributions comparable to the posterior distributions of traditional, Bayesian retrievals. Zingales & Waldmann (2018, ExoGAN) utilized a generative adversarial network (GAN, Goodfellow et al. 2014) to retrieve distributions for model parameters. Waldmann & Griffith (2019, PlanetNet) used a convolutional NN (CNN) to map spatial and spectral features across Saturn. In Cobb et al. (2019), we introduced **plan-net**, an ensemble of Bayesian neural networks that uses parameter correlations to inform the uncertainty on retrieved parameters. Hayes et al. (2019) demonstrated a new approach to ML retrieval by applying k -means clustering to a principal component analysis of the observed spectrum to inform a standard Bayesian retrieval. Johnsen & Marley (2019) showed that a dense NN can provide quick estimations of atmospheric properties.

While these approaches are promising, all except Hayes et al. (2019) suffer from a common deficiency: the reduction in computational time is accompanied by

a reduction in posterior accuracy because they are not performing a true Bayesian inference. For ML to become an integral part of atmospheric retrieval, the posterior accuracy must be preserved.

The solution lies in simulation-based inference methods (Cranmer et al. 2019). While directly using a simulator (e.g., RT code) requires a consistent amount of compute time for each new inference (e.g., retrieval), surrogate models that emulate the simulator (e.g, NNs) allow new data to be quickly evaluated after an up-front computational cost to train the surrogate (Kasim et al. 2020; Munk et al. 2019). ML and simulation-based inference approaches have been successfully applied to a variety of tasks ranging from quantum chemistry (Gilmer et al. 2017) to particle physics (Brehmer et al. 2018; Baydin et al. 2019), resulting in significant reductions in compute cost with minimal loss in accuracy.

Here we present a novel application of this approach to retrieval, which uses an NN model of RT within a Bayesian framework, and apply it to the emission spectrum of HD 189733 b. Our general method is to (1) generate a data set over some parameter space, (2) train a surrogate forward model on the generated data, and (3) infer the inverse process via a Bayesian sampler (Figure 1). Our approach circumvents the existing limitations of ML retrieval methods, which seek to directly learn the inverse process, by learning the forward, deterministic process (RT) and using the simulator surrogate in a standard inference pipeline. This approach preserves the accuracy of the Bayesian inference and, while slower than direct ML retrieval, is still much faster than computing RT.

In Section 2 we introduce two software packages that implement our approach for ML atmospheric retrieval. Section 3 details our application of these packages to HD 189733 b, and Section 4 describes our results. Finally, Section 5 presents conclusions.

2. SOFTWARE

We have developed two Python packages implementing our retrieval method for this investigation. Both are open-source software, with full documentation, under the Reproducible Research Software License¹. We encourage users to contribute to the code via pull requests on Github.

2.1. MARGE

At its core, our ML retrieval approach relies on an NN surrogate model for RT. Training an NN requires a

¹ <https://planets.ucf.edu/resources/reproducible-research/software-license/>

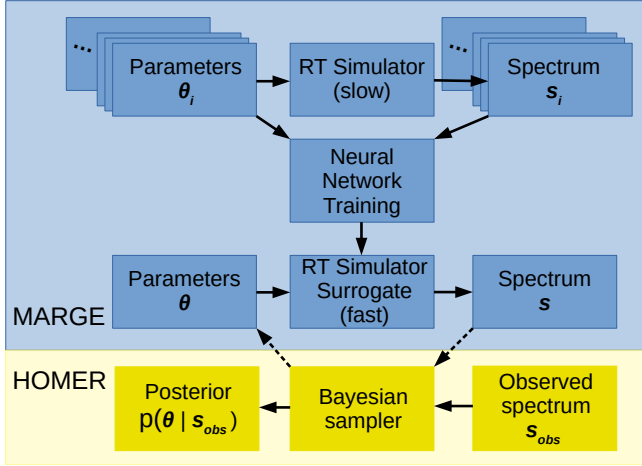


Figure 1. Schematic diagram of our inverse modeling method. MARGE generates a data set based on a deterministic, forward process (e.g., RT) and trains a surrogate model to approximate that process. Using the trained surrogate, HOMER infers the inverse process (e.g., atmospheric retrieval) by simulating many forward models and comparing them to the target data (e.g., an observed spectrum) in a Bayesian framework.

large data set encompassing the parameter space where the NN will be applied. We have therefore developed the Machine learning Algorithm for Radiative transfer of Generated Exoplanets² (MARGE), which (1) generates a specified number of atmospheric models within some parameter space and computes their emission spectra, (2) processes the generated data, and (3) trains, validates, and tests a user-specified NN architecture on the generated data set (Figure 1). The software package allows independent execution of any of the three modes, enabling a wide range of applications beyond exoplanet retrieval.

At present, MARGE generates data using a modified version of the Bayesian Atmospheric Radiative Transfer code (BART, Harrington et al. 2020; Cubillos et al. 2020; Blecic et al. 2020), though it is designed to easily use other software packages for this purpose. BART’s Bayesian sampler is the Multi-core Markov chain Monte Carlo (MCMC) code (MC3, Cubillos et al. 2017). We have extended MC3 with a random uniform sampler to generate collections of forward models. BART’s RT package, *transit*, handles the calculation of emission spectra.

We implement NN model training in Keras (version 2.2.4, Chollet et al. 2015), using Tensorflow (version 1.13.1, Abadi et al. 2016) for the backend. MARGE enables early stopping by default to prevent overfitting,

and the user can halt or resume training. MARGE allows for cyclical learning rates for more efficient training (Smith 2015). Users specify the model architecture details and the data location, and the software handles the data normalization, training, validation, and testing. MARGE pre-processes the data into Tensorflow’s TFRecords format for efficient handling. Users have multiple options when pre-processing the data, which include taking the logarithm of the inputs and/or outputs, standardizing the data according to its mean and standard deviation, and/or scaling the data to be within some range. The mean and standard deviation of the data set are computed using Welford’s method (Welford 1962) to avoid the need to load the entire data set into memory at once. MARGE computes the root-mean-squared error (RMSE) and coefficient of determination (R^2) for predictions on the validation and test sets to evaluate model performance. Finally, users may specify cases from the test set to plot the predicted and true spectra, with residuals (e.g., Figure 2).

2.2. HOMER

The Helper Of My Eternal Retrievals³ (HOMER), utilizes a MARGE-trained model in place of BART’s *transit* RT code (Figure 1). HOMER currently uses MC3 for its Bayesian sampler. The user specifies data, uncertainties, observational filters, a parameter space, and a few related inputs, which are passed to MC3 to perform the inference. A graphics processing unit (GPU) calculates spectra, though the central processing unit (CPU) can do this at the cost of increased runtime. For each iteration of the MCMC, the trained NN predicts spectra for the proposed input parameters, which are modified as necessary (de-scale, de-normalize, divide by the stellar spectrum, unit conversions, etc.) and integrated over the observational bandpasses. MC3 determines convergence of parallel chains via the test of Gelman & Rubin (1992). HOMER produces plots of the 1D marginalized posteriors, 2D pairwise posteriors, and parameter history traces. To estimate the accuracy of credible regions, HOMER calculates the steps per effective independent sample (SPEIS) and effective sample size (ESS) as described in Harrington et al. (2020). For easy comparison with other retrieval results, HOMER can also overplot the 1D and 2D posteriors for multiple MCMC runs (e.g., Figure 3). HOMER computes the Bhattacharyya coefficient (Bhattacharyya 1943; Aherne et al. 1998) to compare the similarity of 1D marginalized

² MARGE is available at <https://github.com/exosports/marge>

³ HOMER is available at <https://github.com/exosports/homer>

Table 1. Forward Model Parameter Space

Parameter	Minimum	Maximum
$\log(\kappa)$	-5.0	1.0
$\log(\gamma_1)$	-2.0	2.0
$\log(\gamma_2)$	-1.3	1.3
α	0.0	1.0
β	0.5	1.3
$\log(\text{H}_2\text{O})$	-13	-0.5
$\log(\text{CO}_2)$	-13	-0.5
$\log(\text{CO})$	-13	-0.5
$\log(\text{CH}_4)$	-13	-0.5

posteriors, where a value of 0 indicates no overlap and a value of 1 indicates identical distributions.

3. METHODS

We first generate emission spectra across the retrieval parameter space (Table 1). We use the five-parameter temperature–pressure profile, $T(p)$, parameterization of [Line et al. \(2013\)](#): κ , the Planck mean infrared opacity; γ_1 and γ_2 , the ratios of the Planck mean visible and infrared opacities for each of two streams; α , which controls the contribution of the two streams; and β , which represents albedo, emissivity, and energy recirculation. We also include a free parameter for the uniform vertical abundance profiles of H_2O , CO_2 , CO , and CH_4 . We allow a wide range of values without regard for physical plausibility, except by enforcing that (1) the H_2/He ratio remains constant, (2) the total relative abundances of molecules in the atmosphere equals 1, and (3) the $T(p)$ profile does not exceed the temperature range of the line lists. For example, this could lead to models with H_2O at conditions where it dissociates ([Arcangeli et al. 2018](#)), though such models would be rejected with a high probability in the case of HD 189733 b.

For opacities, we use HITEMP for H_2O , CO , and CO_2 ([Rothman et al. 2010](#)), HITRAN for CH_4 ([Rothman et al. 2013](#)), and collision-induced absorptions of $\text{H}_2\text{-H}_2$ and $\text{H}_2\text{-He}$ ([Borysow et al. 2001](#); [Borysow 2002](#); [Abel et al. 2012](#); [Richard et al. 2012](#)). While there are newer line lists available with a greater number of lines, these tests are meant to demonstrate consistency between BART and HOMER; we therefore use the setup described in [Harrington et al. \(2020\)](#), which uses this set of line lists to compare with previous studies.

To train our NN surrogate model, we generate 315,626 spectra, which are subdivided into 227,736 ($\sim 70\%$) for training, 61,811 ($\sim 20\%$) for validation, and 26,079 ($\sim 10\%$) for testing. Model parameters come from the uniform distribution bound by the limits listed in Table

1. Each spectrum spans $280 - 7100 \text{ cm}^{-1}$ at a resolution of 0.25 cm^{-1} and corresponds to the planet’s emitted flux in $\text{erg s}^{-1} \text{ cm}^{-1}$. We normalize the data by taking the logarithm of the target output values, standardizing the data by subtracting the training mean and dividing by the training standard deviation, and scaling the data based on the standardized training set extrema to be in the range $[-1, 1]$. The NN model’s hidden layers consist of Conv1d(64) – Dense(4096) – Dense(4096) – Dense(4096), with ReLU activation functions for each hidden layer. The Dense output layer has 27,281 nodes, corresponding to the emitted spectrum over the defined wavenumber grid. We train with a batch size of 256 using a mean squared error loss function, the Adam optimizer, and early stopping with a patience of 20 epochs based on the validation loss. We employ a cyclical learning rate that increases from 10^{-4} to 10^{-3} over 4 epochs, then decreases over the same window. During each complete cycle (8 epochs), the maximum learning rate decays by half the difference between the maximum and minimum learning rates. The boundaries were chosen according to the method described in [Smith \(2015\)](#), except that we consider the loss instead of accuracy. To evaluate the model’s performance, we compute RMSE and R^2 between the normalized data and the predictions.

Following the setup of [Harrington et al. \(2020\)](#), we perform a retrieval of the dayside atmosphere of HD 189733 b based on the measurements by the *Hubble Space Telescope* Near Infrared Camera MultiObject Spectrograph ([Swain et al. 2009](#)); *Spitzer Space Telescope* Infrared Spectrograph (IRS [Todorov et al. 2014](#)); *Spitzer* InfraRed Array Camera (IRAC) channels 1 and 2 ([Knutson et al. 2012](#)); IRAC channel 3, IRS 16 μm photometry, and Multiband Imaging Photometer for *Spitzer* ([Charbonneau et al. 2008](#)); and IRAC channel 4 ([Agol et al. 2010](#)). We use a K2 solar-abundance Kurucz stellar model for the host star’s emission ([Castelli & Kurucz 2003](#)). Using the differential evolution with snooker updating algorithm of [ter Braak & Vrugt \(2008\)](#), 2,000,000 iterations are spread across 10 parallel chains, with a burn-in of 5000 iterations per chain.

For this investigation, we focus on MARGE+HOMER as a faster replacement for an analytical retrieval code; we therefore only compare the results of HOMER and BART here. For a discussion of these results in the context of previous retrievals of HD 189733 b’s dayside atmosphere, see [Harrington et al. \(2020\)](#).

Table 2. Model Evaluation

Metric	Minimum	Median	Mean	Maximum
RMSE	0.0046	0.0061	0.0066	0.0382
R^2	0.9918	0.9998	0.9997	0.9999

4. RESULTS & DISCUSSION

Table 2 details the normalized RMSE and R^2 metrics for the MARGE-trained model on the test set. The small RMSE and $R^2 \sim 1$ indicate an accurate model for RT over the parameter space. Figure 2 shows example comparisons between the spectra predicted by MARGE and true spectra calculated with *transit*. While residuals tend to be around 5%, they fluctuate around 0; when band integrated over the observational filters, these errors effectively cancel. There are occasionally regions where the spectrum is consistently over- or underestimated (e.g., near 7000 cm^{-1} in the left panel of Figure 2), thereby introducing error in the band-integrated value. However, the small deviations appear to have little effect on this retrieval’s result.

When applying HOMER to the emission spectrum of HD 189733 b, the results are consistent with BART. The retrieved $T(p)$ profiles (Figure 3, right panel) agree in the regions probed by the observations. They deviate in the lower atmosphere, where little to no signal is measured, as determined by the contribution functions (Figure 10, top right panel, of Harrington et al. 2020). By nature, HOMER cannot calculate contribution functions, as the MARGE model does not solve RT. We could have included contribution functions for each case in the training set, but this would require significantly more compute resources. Computing the contribution functions for the single best-fit case using the RT code that trained MARGE is much faster.

Table 3 compares HOMER’s retrieved 68.27% (“1 σ ”), 95.45% (“2 σ ”), and 99.73% (“3 σ ”) credible regions with BART’s retrieved credible regions (Harrington et al. 2020). While the 68.27% regions for CH_4 and CO disagree by a noticeable amount, the 95.45% and 99.73% regions are in close agreement. This is due to the errors in the NN model as well as the mostly-flat nature of the posterior for those parameters (Figure 3, left panel). Table 4 compares the SPEIS, ESS values, and associated uncertainties in the 1, 2, and 3 σ credible regions for HOMER and BART; their values differ by $\sim 10\%$, attributable to randomness among the sampler’s chains. The Bhattacharyya coefficients between the 1D marginalized posteriors of HOMER and BART indicate agreement in the range 97.74 – 99.74% (Table 5).

HOMER’s accuracy is, by nature, bound by the accuracy of the NN model. Model inaccuracies may slightly bias the results, as seen in the small peak near -4 for CH_4 present in HOMER’s posterior but not BART’s. This discrepancy will shrink as NN model accuracy improves and does not significantly affect the scientific conclusions at the spectral resolution of these observations for the current NN accuracy.

For the architecture presented here, MARGE predicts each spectrum in ~ 0.0002 seconds on an NVIDIA Titan Xp, compared to ~ 0.006 seconds on an Intel i7-4770 – a $\sim 30\times$ speedup. The combination of MARGE + HOMER required under 400 CPU hours (mostly spent on data generation and NN model training), compared to BART’s ~ 360 CPU hours. While additional retrievals with BART within this parameter space would take a similar length of time, additional retrievals with HOMER using a GPU would require ~ 0.6 CPU and GPU hours – a $\sim 600\times$ speedup. Further, if using a GPU, retrievals with HOMER can utilize hundreds of parallel chains. Additional chains enable faster exploration of the parameter space, and, if executed for the same number of iterations per chain, increases the ESS, which reduces the uncertainty in the bounds of credible regions (Harrington et al. 2020). While BART’s compute cost scales linearly with the number of chains for a given number of iterations per chain, HOMER’s does not. For example, using 64 chains instead of 10 requires $\sim 4.6\times$ as much compute time, compared to $6.4\times$ as much for BART. Thus, the combination of MARGE and HOMER saves valuable compute resources when performing multiple retrievals over a similar parameter space.

5. CONCLUSIONS

This paper presents a novel technique for ML retrieval that uses an NN model of RT within a Bayesian framework to reduce the runtime of a retrieval. Our open-source codes, MARGE and HOMER, provide the community with an easy-to-use implementation of this approach.

Our method enables fast retrievals that are consistent with traditional, analytical retrieval algorithms. The approach circumvents limitations of current ML retrieval models by using an RT surrogate in place of the analytical RT code found in traditional retrieval algorithms, thereby preserving the accuracy of the Bayesian inference. On our hardware, HOMER achieves a $\sim 600\times$ speedup in performing a retrieval that yields 1D marginalized posteriors that overlap with $>97.7\%$ of BART’s posteriors for the case of HD 189733 b. This reduction in compute time justifies using more realistic

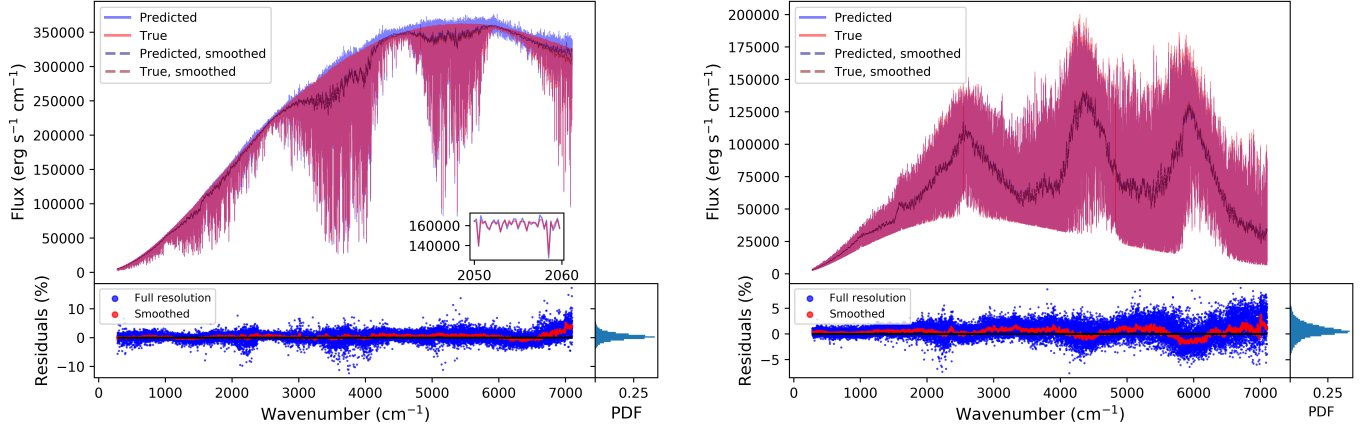


Figure 2. Two comparisons of MARGE and `transit` spectra. The smoothed curves use a Savitzky-Golay filter with a third-order polynomial across a window of 101 elements (25 cm^{-1}). The purple color arises due to a detailed match between the red and blue spectra lines. For the residuals, a black line is plotted at 0 to show regions where the NN consistently over- or underpredicts the spectrum. A histogram of the residuals appears to the right of the residual scatter plot, where the x-axis shows the probability density function (PDF) for the range of residual percentages. Left: case with mainly absorption. Inset plot shows a zoomed-in region. Right: case with emission and absorption.

Table 3. Retrieved Credible Regions

Log Parameter	Code	68.27%	95.45%	99.73%
H ₂ O	HOMER	[-2.1, -1.1]	[-2.4, -0.6]	[-2.7, -0.5]
	BART	[-1.9, -0.8]	[-2.3, -0.5]	[-2.6, -0.5]
CO ₂	HOMER	[-2.6, -1.7]	[-2.9, -1.3]	[-3.2, -0.9]
	BART	[-2.5, -1.5]	[-2.8, -1.1]	[-3.2, -0.7]
CO	HOMER	[-4.4, -0.5]	[-12.4, -0.5]	[-12.9, -0.5]
	BART	[-6.6, -0.5]	[-12.0, -0.5]	[-12.9, -0.5]
CH ₄	HOMER	[-11.8, -7.2] \cup [-5.0, -3.6]	[-12.8, -3.5]	[-13.0, -2.7]
	BART	[-11.7, -5.5]	[-12.9, -3.7]	[-13.0, -2.6]

Table 4. Credible Region Accuracy

Code	SPEIS	ESS	1σ Uncertainty	2σ Uncertainty	3σ Uncertainty
HOMER	537	3631	0.772%	0.346%	0.086%
BART	486	4012	0.735%	0.329%	0.082%

(and computationally expensive) RT models. Additionally, 3D retrievals with 180 cells could be completed in a matter of days.

Our approach is particularly well suited to planning studies for future observations, telescopes, and instruments, like the James Webb Space Telescope and the Large UltraViolet Optical InfraRed Surveyor (e.g., [Rocchetto et al. 2016](#); [Feng et al. 2018](#)). Using a single MARGE model trained over the desired parameter space, HOMER can perform dozens to hundreds of retrievals in the time it takes to run a single retrieval with an RT solver.

More generally, our technique and tools can be applied to problems beyond the scope of this investigation. MARGE provides a generalized method to train an NN to model a deterministic process, while HOMER uses a MARGE-trained model to infer the inverse process. MARGE models could be trained for cloud/haze formation or photochemistry. MARGE and HOMER could also be used to map gravitationally-lensed galaxies (e.g., [Perreault Levasseur et al. 2017](#)).

With the plethora of ML retrieval algorithms that have emerged in recent years, standard data sets should be created and used for benchmarking. Ideally, such




Figure 3. Comparisons between HOMER and BART posteriors. Left: Normalized probability density functions (PDF) of the 1D marginalized posteriors retrieved for HD 189733 b. Right: Posterior median, 1 σ from the median, and 2 σ from the median $T(p)$ profile.

Table 5. Bhattacharyya Coefficients

Parameter	Value
κ	97.87%
γ_1	97.74%
γ_2	99.74%
α	99.55%
β	98.30%
H_2O	98.42%
CO_2	97.97%
CO	98.83%
CH_4	99.45%

a data set would cover a wide range of wavelengths at high resolution and include all available opacity sources, scattering, clouds/hazes, and, in the case of terrestrial

planets, surface properties. This would allow easy comparisons among current and future ML retrieval codes.

The Reproducible Research Compendium for this work is available for download⁴. It includes all of the code, configuration files, data, and plots used in support of this work.

ACKNOWLEDGMENTS

We thank contributors to NumPy, SciPy, Matplotlib, Tensorflow, Keras, and the Python Programming Language, the free and open-source community, and the NASA Astrophysics Data System for software and services. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research. This research was supported by the NASA Fellowship Activity under NASA Grant 80NSSC20K0682 and NASA Exoplanets Research Program grant NNX17AB62G.

REFERENCES

- Abadi, M., Barham, P., Chen, J., et al. 2016, in OSDI, Vol. 16, 265–283
- Abel, M., Frommhold, L., Li, X., & Hunt, K. L. C. 2012, The Journal of Chemical Physics, 136, 044319, doi: [10.1063/1.3676405](https://doi.org/10.1063/1.3676405)
- Agol, E., Cowan, N. B., Knutson, H. A., et al. 2010, ApJ, 721, 1861, doi: [10.1088/0004-637X/721/2/1861](https://doi.org/10.1088/0004-637X/721/2/1861)
- ⁴ Available at <https://physics.ucf.edu/~jh/RRC-HimesEtal2020-MARGE-HOMER/>
- Aherne, F. J., Thacker, N. A., & Rockett, P. I. 1998, Kybernetika, 34, 363
- Amodei, D., Ananthanarayanan, S., Anubhai, R., et al. 2016, in International Conference on Machine Learning, 173–182
- Arcangeli, J., Désert, J.-M., Line, M. R., et al. 2018, ApJ, 855, L30, doi: [10.3847/2041-8213/aab272](https://doi.org/10.3847/2041-8213/aab272)
- Baydin, A. G., Heinrich, L., Bhimji, W., et al. 2019, in Advances in Neural Information Processing Systems 33
- Bhattacharyya, A. 1943, Bull. Calcutta Math. Soc., 35, 99

- Blecic, J., Harrington, J., Cubillos, P., et al. 2020, ApJS, in review
- Borysow, A. 2002, A&A, 390, 779, doi: [10.1051/0004-6361:20020555](https://doi.org/10.1051/0004-6361:20020555)
- Borysow, A., Jorgensen, U. G., & Fu, Y. 2001, JQSRT, 68, 235, doi: [10.1016/S0022-4073\(00\)00023-6](https://doi.org/10.1016/S0022-4073(00)00023-6)
- Brehmer, J., Cranmer, K., Louppe, G., & Pavez, J. 2018, Phys. Rev. D, 98, 052004, doi: [10.1103/PhysRevD.98.052004](https://doi.org/10.1103/PhysRevD.98.052004)
- Castelli, F., & Kurucz, R. L. 2003, in IAU Symposium, Vol. 210, Modelling of Stellar Atmospheres, ed. N. Piskunov, W. W. Weiss, & D. F. Gray, A20. <https://arxiv.org/abs/astro-ph/0405087>
- Charbonneau, D., Knutson, H. A., Barman, T., et al. 2008, ApJ, 686, 1341, doi: [10.1086/591635](https://doi.org/10.1086/591635)
- Chollet, F., et al. 2015, Keras, <https://github.com/keras-team/keras>, GitHub
- Cobb, A. D., Himes, M. D., Soboczinski, F., et al. 2019, AJ, 158, 33, doi: [10.3847/1538-3881/ab2390](https://doi.org/10.3847/1538-3881/ab2390)
- Cranmer, K., Brehmer, J., & Louppe, G. 2019, arXiv preprint arXiv:1911.01429
- Cubillos, P., Harrington, J., Lored, T. J., et al. 2017, AJ, 153, 3, doi: [10.3847/1538-3881/153/1/3](https://doi.org/10.3847/1538-3881/153/1/3)
- Cubillos, P., Harrington, J., Blecic, J., et al. 2020, ApJS, in review
- Deming, L. D., & Seager, S. 2017, Journal of Geophysical Research (Planets), 122, 53, doi: [10.1002/2016JE005155](https://doi.org/10.1002/2016JE005155)
- Feng, Y. K., Robinson, T. D., Fortney, J. J., et al. 2018, AJ, 155, 200, doi: [10.3847/1538-3881/aab95c](https://doi.org/10.3847/1538-3881/aab95c)
- Gelman, A., & Rubin, D. B. 1992, Statistical Science, 7, 457
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., & Dahl, G. E. 2017, in Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML17 (JMLR.org), 12631272
- Goodfellow, I., Bengio, Y., & Courville, A. 2016, Deep Learning (MIT Press)
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., et al. 2014, in Advances in Neural Information Processing Systems, 2672–2680
- Harrington, J., Himes, M., Cubillos, P., et al. 2020, ApJS, in review
- Hayes, J. J. C., Kerins, E., Awiphan, S., et al. 2019, arXiv e-prints, arXiv:1909.00718. <https://arxiv.org/abs/1909.00718>
- He, K., Zhang, X., Ren, S., & Sun, J. 2016, in Proceedings of the IEEE conference on computer vision and pattern recognition, 770–778
- Johnsen, T. K., & Marley, M. S. 2019, arXiv e-prints, arXiv:1911.03997. <https://arxiv.org/abs/1911.03997>
- Kasim, M., Watson-Parris, D., Deaconu, L., et al. 2020, arXiv preprint arXiv:2001.08055
- Knutson, H. A., Lewis, N., Fortney, J. J., et al. 2012, ApJ, 754, 22, doi: [10.1088/0004-637X/754/1/22](https://doi.org/10.1088/0004-637X/754/1/22)
- Koskinen, T. T., Moses, J. I., West, R. A., Guerlet, S., & Jouchoux, A. 2016, Geophys. Res. Lett., 43, 7895, doi: [10.1002/2016GL070000](https://doi.org/10.1002/2016GL070000)
- Line, M. R., Wolf, A. S., Zhang, X., et al. 2013, ApJ, 775, 137, doi: [10.1088/0004-637X/775/2/137](https://doi.org/10.1088/0004-637X/775/2/137)
- Madhusudhan, N. 2018, Atmospheric Retrieval of Exoplanets (Springer International Publishing AG), 104, doi: [10.1007/978-3-319-55333-7_104](https://doi.org/10.1007/978-3-319-55333-7_104)
- Márquez-Neila, P., Fisher, C., Sznitman, R., & Heng, K. 2018, Nature Astronomy, 2, 719, doi: [10.1038/s41550-018-0504-2](https://doi.org/10.1038/s41550-018-0504-2)
- Munk, A., cibior, A., Baydin, A. G., et al. 2019, arXiv preprint arXiv:1910.11950
- Perreault Levasseur, L., Hezaveh, Y. D., & Wechsler, R. H. 2017, ApJL, 850, L7, doi: [10.3847/2041-8213/aa9704](https://doi.org/10.3847/2041-8213/aa9704)
- Richard, C., Gordon, I., Rothman, L., et al. 2012, Journal of Quantitative Spectroscopy and Radiative Transfer, 113, 1276, doi: <https://doi.org/10.1016/j.jqsrt.2011.11.004>
- Rocchetto, M., Waldmann, I. P., Venot, O., Lagage, P. O., & Tinetti, G. 2016, ApJ, 833, 120, doi: [10.3847/1538-4357/833/1/120](https://doi.org/10.3847/1538-4357/833/1/120)
- Rothman, L. S., Gordon, I. E., Barber, R. J., et al. 2010, JQSRT, 111, 2139
- Rothman, L. S., Gordon, I. E., Babikov, Y., et al. 2013, JQSRT, 130, 4
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. 1986, Nature, 323, 533, doi: [10.1038/323533a0](https://doi.org/10.1038/323533a0)
- Smith, L. N. 2015, arXiv e-prints, arXiv:1506.01186. <https://arxiv.org/abs/1506.01186>
- Swain, M. R., Vasisht, G., Tinetti, G., et al. 2009, ApJL, 690, L114, doi: [10.1088/0004-637X/690/2/L114](https://doi.org/10.1088/0004-637X/690/2/L114)
- ter Braak, C. J. F., & Vrugt, J. A. 2008, Statistics and Computing, 18, 435, doi: [10.1007/s11222-008-9104-9](https://doi.org/10.1007/s11222-008-9104-9)
- Todorov, K. O., Deming, D., Burrows, A., & Grillmair, C. J. 2014, ApJ, 796, 100, doi: [10.1088/0004-637X/796/2/100](https://doi.org/10.1088/0004-637X/796/2/100)
- Waldmann, I. P., & Griffith, C. A. 2019, Nature Astronomy, 3, 620, doi: [10.1038/s41550-019-0753-8](https://doi.org/10.1038/s41550-019-0753-8)
- Welford, B. P. 1962, Technometrics, 4, 419, doi: [10.1080/00401706.1962.10490022](https://doi.org/10.1080/00401706.1962.10490022)
- Wu, Y., Schuster, M., Chen, Z., et al. 2016, arXiv preprint arXiv:1609.08144
- Zingales, T., & Waldmann, I. P. 2018, AJ, 156, 268, doi: [10.3847/1538-3881/aae77c](https://doi.org/10.3847/1538-3881/aae77c)