



Neural Network Radiative Transfer for Imaging Spectroscopy

Brian D. Bue¹, David R. Thompson¹, Shubhankar Deshpande², Michael Eastwood¹, Robert O. Green¹, Terry Mullen³, Vijay Natraj¹, and Mario Parente³

¹Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA

²The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA

³University of Massachusetts, Amherst, MA, USA

Correspondence: B Bue (bbue@jpl.nasa.gov)

Abstract. Visible/Shortwave InfraRed imaging spectroscopy provides valuable remote measurements of Earth's surface and atmospheric properties. These measurements generally rely on inversions of computationally-intensive Radiative Transfer Models (RTMs). RTMs' computational expense makes them difficult to use with high volume imaging spectrometers, and forces approximations such as lookup table interpolation and surface/atmosphere decoupling. These compromises limit the accuracy and flexibility of the remote retrieval; dramatic speed improvements in radiative transfer models could significantly improve the utility and interpretability of remote spectroscopy for Earth science. This study demonstrates that nonparametric function approximation with neural networks can replicate radiative transfer calculations over a relevant range of surface/atmosphere parameters. Incorporating physical knowledge into the network structure provides improved interpretability and model efficiency. We evaluate the approach in atmospheric correction of data from the PRISM airborne imaging spectrometer, and demonstrate accurate emulation of radiative transfer calculations which run several orders of magnitude faster than first-principles models. These results are particularly amenable to iterative spectrum fitting approaches, providing analytical benefits including statistically rigorous treatment of uncertainty and the potential to recover information on spectrally-broad signals.

Copyright statement. The author's copyright for this publication is transferred to the Jet Propulsion Laboratory, California Institute of Technology.

1 Introduction

Remote Visible / ShortWave InfraRed (VSWIR) imaging spectroscopy, also known as hyperspectral imaging, is a powerful approach to map the composition, health, and biodiversity of Earth's ecosystems (ESAS, 2018). Remote sensing of the solar-reflected spectrum from 380-2500nm reveals physics and chemistry of many processes in Earth's surface/atmosphere system (Schaepman et al., 2009), including: terrestrial plant health and traits (Asner et al., 2017; Ustin et al., 2004); biodiversity (Jetz et al., 2016); the condition and composition of aquatic, benthic, and near-shore ecosystems (Fichot et al., 2015; Hochberg, 2011); geology (Swayze et al., 2014); and trace greenhouse gases (Frankenberg et al., 2016). While Earth scientists aim to measure these geophysical variables, remote sensors can only measure the incident light at the sensor. Inferring geophysical



properties requires inverting the measurement with a physical model — typically one that accounts for both absorption and scattering by the atmosphere, and the fraction of light reflected from the surface at each wavelength (Schaepman-Strub et al., 2006).

Radiative Transfer Models (RTMs) such as DISORT (Stamnes et al., 1988) are a critical component of such models, and form the core of common spectroscopy analysis codes including ACORN (Kruse, 2004), ATCOR (Richter and Schlapfer, 2002), FLAASH (Perkins et al., 2012), ATREM (Gao et al., 1993), and others (Gao et al., 2000, 2007; Thompson et al., 2015). The RTM posits a stratified atmosphere populated by atmospheric gases at appropriate concentrations and temperatures, and solves the general equations of radiative transport based on a known solar input and presumed surface. This is an intensive computation, requiring special care for modern high volume imaging spectrometers that acquire thousands or millions of spectra per second.

Because imaging spectrometers produce too much data to analyze each measurement with an independent RTM, investigators use RTMs to pre-calculate lookup tables of atmospheric optical properties — such as scattered path radiance or transmission — for atmospheric states appropriate to the conditions observed at the time a given image was acquired. At runtime, the model inversion estimates the actual state from the radiance spectrum, and interpolates within the lookup table to find the associated optical properties. This informs parametric approximations of atmospheric transport, such as the formulation by Vermote (Vermote et al., 1997), permitting algebraic solutions for the remaining unknowns like surface reflectance. The sequential retrieval of atmospheric and surface properties, a process known as *atmospheric correction*, obtains a self-consistent but approximate explanation for the surface and atmosphere system.

The lookup table solution is adequate for many needs, but imposes several limitations. First, lookup tables can only model a few degrees of freedom in atmospheric state due to the “curse of dimensionality;” the number of samples necessary to adequately represent the state space increases exponentially with the number of input variables. To increase the fidelity of grid samples in high dimensions, designers can leverage representative sampling or hyperparameter search strategies such as Snoek et al. (2012) within the state space, or space-filling sampling methods like Latin Hypercube Sampling (Stein, 1987) or lattice regression methods (Gupta et al., 2015). However, such techniques are restricted by prohibitive computation and storage requirements for high-dimensional state spaces, and incur increased risks of interpolation inaccuracy. Also, because the contents of precalculated lookup tables capture atmospheric optical properties independently from the surface, lookup table-based approaches preclude strong coupling between the two. Speeding RTMs to the point at which they could run many times faster for each spectrum would obviate the lookup table compromise and enable more flexible, accurate, and statistically rigorous inversion algorithms such as the optimal estimation approach used in many atmospheric sounding missions (Thompson et al., 2018c; Rodgers, 2000).

Recent work suggested the use of nonparametric function approximators such as neural networks (Verrelst et al., 2016, 2017; Thompson et al., 2018a) or Gaussian Processes (Martino et al., 2017) for this purpose. Investigators can train such models using prior runs of radiative transfer models over relevant ranges of surface and atmospheric conditions. After learning the underlying function with sufficient accuracy, the trained model could act as an instrument-specific RTM that would not have to solve the underlying differential equations. Alternative formulations such as Jamet et al. (2005); Brajard et al. (2006) provide empirical



validation of RTM assumptions by evaluating atmospheric, transmittance, and surface interactions captured in separate models, while similar methods (e.g., Jamet et al. (2012)) permit retrieval of radiometric parameters spanning multiple wavelengths. However, to date, these methods have been demonstrated on a small number of cases with limited surfaces and atmospheres (Verrelst et al., 2017; Martino et al., 2017; Brajard et al., 2006), but not across the VSWIR range with surface state vector
5 flexibilities that would permit a functionally-useful alternative to existing atmospheric correction routines. To our knowledge, emulation has not yet been demonstrated on actual imaging spectroscopy data outside of simulations.

This study demonstrates an accurate neural network model deployed as part of an iterative model inversion, showing that emulation is a practical solution for operational atmospheric correction of imaging spectroscopy data. This opens new possible avenues of research, both for the inversion algorithm itself (to explore further expansions of the state vector beyond the tra-
10 ditional retrieved variables) and for downstream analyses (to exploit the benefits of new retrieval methods that do not require lookup tables). We begin by describing the neural network architecture and RTM emulation methodology, including several novel advances: an *analytical decomposition* of the radiance spectrum into quantities that are individually easier to model, *monochromatic subnetworks* to simplify training, and channelwise *weight propagation* to reduce training time and account for channelwise correlation. We then present a case study where we evaluate the approach for the PRISM imaging spectrometer,
15 and demonstrate high quality retrievals of surface reflectance based on the optimal estimation approach of Thompson et al. (2018c). The retrievals capture subtle atmospheric variability such as view-dependence of Rayleigh scattering not typically handled in conventional atmospheric correction codes. Finally, we describe paths for future development of neural network emulation technology.

2 Neural Networks for Radiative Transfer Modeling

20 The objective of our function approximation is to map a surface/atmosphere state vector $\mathbf{x} \in \mathbb{R}^m$ to a vector of observed at-sensor radiances $\mathbf{y} \in \mathbb{R}^n$ measured at wavelengths $\{\lambda_1, \dots, \lambda_n\}$. We use boldface notation to signify vectors and matrices, with all matrices in capital letters. In our current state space, the surface has a single free parameter per instrument channel, and we also include the following four auxiliary parameters: the atmospheric aerosol optical depth at the surface, τ ; the atmospheric water vapor content of the column in g cm^{-2} , H_2O ; the cosine of the observer zenith angle θ_v ; and the relative azimuth angle
25 between the observer and the sun, written ϕ_r . These free parameters values vary independently for every spectrum in a given flightline. Naturally, alternative parameterizations are possible, including mixture models, continuum-absorption models, and others. However, these could be mapped to our representation without loss of generality.

We exploit two features of the problem to simplify the function $F(\mathbf{x})$. First, we reduce the radiance spectrum analytically into the reflectance of the surface/atmosphere system, written $\boldsymbol{\rho}_{obs}$, defined as $\boldsymbol{\rho}_{obs} = \mathbf{y}\pi/\phi_o\mathbf{e}_o$ where ϕ_o is the cosine of
30 the solar zenith angle and \mathbf{e}_o the extra-terrestrial solar irradiance. The top of atmosphere reflectance $\boldsymbol{\rho}_{obs}$ is normalized for solar illumination and, absent extreme glint, resides conveniently in the $[0,1]$ interval, making it an easier target for function approximation. For any given observing geometry, the known values of ϕ_o and \mathbf{e}_o can be used to predict the corresponding radiances.



Second, we leverage the fact that the observed radiance at any given channel is fully specified by the observation geometry, atmospheric state, and the reflectance *in that channel*. In statistical terms, absent any prior distribution that couples neighboring reflectances, the channel radiance values become conditionally independent of each other given the atmosphere and observation geometry. This permits an exact decomposition of F into *monochromatic subnetworks* that each model a single wavelength channel. We train the monochromatic subnetworks to model $F(\mathbf{x})$ using precalculated outputs generated by a first-principles RTM at representative points in the state space. Specifically, provided an input state vector \mathbf{x} , each subnetwork predicts the top of atmosphere reflectance for its corresponding channel, which we then combine with solar irradiance to predict radiance. Figure 1 shows the topology of a single subnetwork. A side benefit of this approach is that the partial derivatives of radiance channels with respect to their surface reflectances are independent of each other, which simplifies calculations of analytical Jacobians during iterative gradient descent inversions (Thompson et al., 2018c).

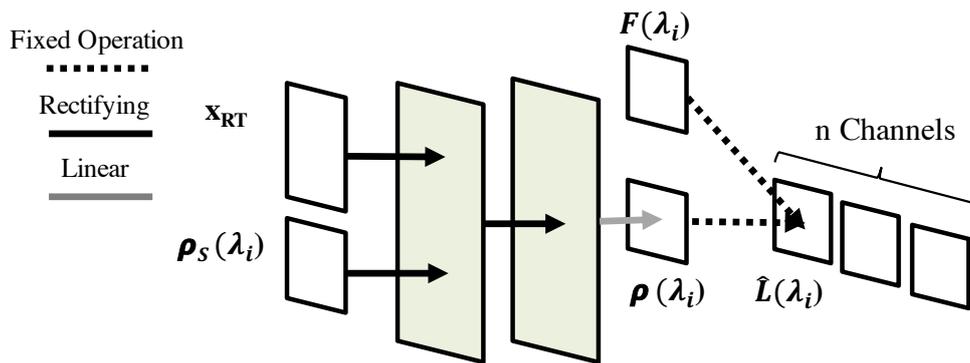


Figure 1. Topology of a single monochromatic subnetwork in our neural RTM emulator. We train input weights, output weights, and biases for both of the hidden (green) layers to predict the top of atmosphere reflectance in a single channel.

Each subnetwork uses a feed-forward architecture with two hidden layers of 50 units each, rectifying linear activation functions in the hidden layers, and a linear activation in the output layer. We use the method proposed by Glorot et al., (Glorot and Bengio, 2010) to initialize the weights for each subnetwork, and use the widely-used error back propagation algorithm (Werbos, 1982) with adaptive moment estimation (Kingma and Ba, 2014) to optimize the weights via gradient descent with respect to the L^2 -regularized mean squared error. Once trained, we combine the channelwise subnetworks into a single *neural RTM emulator* which generates a single vector valued prediction \mathbf{y} representing the top of atmosphere reflectance for all modeled instrument wavelengths.

Constructing a robust neural RTM emulator from precomputed RTM outputs faces two fundamental modeling challenges. First, the precomputed RTM outputs must provide sufficient coverage of the state space to represent the distribution of spectral responses in each channel. Second, the subnetworks must accurately predict RTM outputs for intermediate state parameter values within the bounds of the precomputed state space for all channels. Intuitively, modeling channels whose RTM outputs are stable with respect to small changes in state is easier than modeling channels whose spectral responses vary substantially. For instance, varying concentrations of atmospheric water vapor produce complex, nonlinear behavior for water absorption bands,



while other wavelengths are largely unaffected. Accurately modeling unstable channels may require generating additional RTM outputs at increased sampling density to highlight distinct responses that are poorly represented in the existing precomputed outputs, and additional computational resources to fine-tune the subnetwork to capture those distinctions may also be in order.

To ensure each subnetwork reliably models its corresponding channel, we measure prediction accuracy on a validation set of precalculated RTM outputs excluded from the training process. In our initial experiments, we performed both k-fold cross-validation and k-fold bootstrap sampling, but after we observed that the main sources of variability in the state space emerged from interactions among a small number of state parameter values, we concluded that randomized sampling of the state space lacking informed sample stratification would yield either optimistic or inconsistent estimates of test accuracy and/or convergence time in cross-validation. Ultimately, we concluded that validation using a fixed and bounded subset of the state parameter values would provide a more informative assessment of model performance. Using a bounded subset also permitted direct comparison to lookup table-based approaches, as they require upper and lower bounds on each variable to generate intermediate values via interpolation.

We train each subnetwork to within 0.1% absolute validation error per channel, and use early stopping to stop training when the requisite error is achieved. This level of accuracy is sufficient to make the approximation error a smaller contributor to total uncertainty than other unknowns in the measurement system. For example, it is generally a similar magnitude to relative calibration error of different focal plane array elements, which can vary slightly due to drift between calibrations (Thompson et al., 2018a). A longer training cycle or additional hidden units can be used to match the RTM output more precisely.

Effectively exploiting characteristics of the state space in tandem with RTM modeling assumptions can also improve model accuracy and computational efficiency. One means we use to achieve this is through a process of *weight propagation*. Rather than initializing the weights for each subnetwork from scratch, we use the converged weights of the subnetwork modeling the previous channel to initialize the weights for the subnetwork modeling the current channel, which are then fine-tuned to estimate the spectral responses for the current channel to within 0.1% validation error, as before. In comparison to training each subnetwork from scratch, using weight propagation often yields a substantial reduction in training time, along with improved accuracy for channels whose RTM outputs are relatively stable with respect to the state space, but also demands a slightly more involved training and model validation procedure. Specifically, we require that the subnetwork modeling the first channel converge to validation error below 0.1% to ensure that the propagated weights will be informative for subsequent channels. We also apply more stringent early stopping criteria for the subsequent subnetworks to ensure poorly correlated neighboring channels converge to their appropriate channelwise responses. An additional side benefit is that weight propagation provides an approximate means to account for channelwise coupling for instruments whose spectral response functions for neighboring channels partially overlap, although this issue does not affect most imaging spectrometers.

3 Neural RTM Emulation for PRISM

We define a case study based on the PRISM imaging spectrometer (Mouroulis et al., 2008, 2014). PRISM is a pushbroom design and observes a cross-track angular field of view spanning approximately 30 degrees, and is designed to observe coastal



ocean environments in the 350-1050nm spectral range at approximately 3nm spectral sampling. The instrument was mounted onboard a high altitude ER-2 aircraft which overflew Santa Monica, USA in October 2015 at 20 km above sea level (Thompson et al., 2018b; Trinh et al., 2017). At this altitude, the instrument measured the scene through nearly all of Earth’s atmospheric scattering and absorption, providing a challenging test case with relevance to future orbital instruments.

Parameter	Values
ϕ_r (radians)	$0, \frac{\pi}{8}, \dots, \frac{\pi}{2}, \dots, \frac{7\pi}{8}, \pi$
$\cos(\theta_v)$	0.94, 0.95, 0.96, 0.97 , 0.98, 0.99, 1.0
τ	0.05, 0.1, 0.2 , 0.3
H_2O (g cm^{-2})	0, 0.5, 1.0, 1.5 , 2.0, 2.5
ρ_s	0.05, 0.1, 0.25, 0.5, 1.0

Table 1. Grid values for state parameters used in LibRadTran model runs. State vectors containing the median value of each auxiliary parameter (indicated by red text) are used for model validation.

5 We begin by generating RTM outputs using the LibRadTran radiative transfer code (Emde et al., 2016; Mayer and Kylling, 2005). Table 1 lists the values of the state parameters, yielding total of $9,072 \times$ input state vectors mapped to their corresponding ρ_{obs} output spectra, each resampled to the 245 PRISM instrument wavelengths. We hold out any state vectors containing the median value of each auxiliary parameter (shown in red text in Table 1) as validation data. Figure 2 illustrates the variability in ρ_{obs} spectra with respect to the ϕ_r (top left), $\cos(\theta_v)$ (top right), τ (bottom left) and H_2O (bottom right) state parameters.

10 Unsurprisingly, the most visibly dramatic changes occur as absorption features appear with increased H_2O vapor. Of the remaining auxiliary parameters, only aerosol optical depth has an observable effect on spectral shape across the visible and near-infrared wavelengths. Changes for varying ϕ_r and θ_v are comparatively small, and predominantly observable in the visible range.

Figure 3 compares the ρ_{obs} validation error of channelwise linear regressors (top figure, black line) to the error produced by subnetworks trained from scratch (NN , blue line) versus subnetworks initialized with weight propagation (NN_{WP} , red line). In comparison to the linear regressors, the channelwise subnetworks yield an order of magnitude reduction in prediction error on all channels, clearly demonstrating the inadequacy of linear lookup table interpolation approaches. Weight propagation provides an average reduction of 64% in channelwise error, but also yields systematically higher errors in the H_2O absorption range between 890-1000nm where ρ_{obs} responses at adjacent wavelengths vary more rapidly than for the remaining channels.

20 While it is unsurprising that the H_2O absorption wavelengths are challenging to model, the fact that the two weight initialization schemes yield distinct error distributions for those wavelengths suggests model convergence issues. Figure 4 compares the number of epochs – where one epoch consists of a single pass of weight updates over all samples in the training set – required to converge for channelwise subnetworks trained from scratch versus subnetworks initialized with weight propagation. Over the set of all PRISM instrument channels, weight propagation permits convergence in $\approx 70\%$ fewer epochs over subnetworks not leveraging weight propagation. In terms of raw compute time, our scikit-learn (Pedregosa et al., 2011) implementation requires 2-3 minutes to train a single monochromatic subnetwork from scratch on a single commercial processor core, while

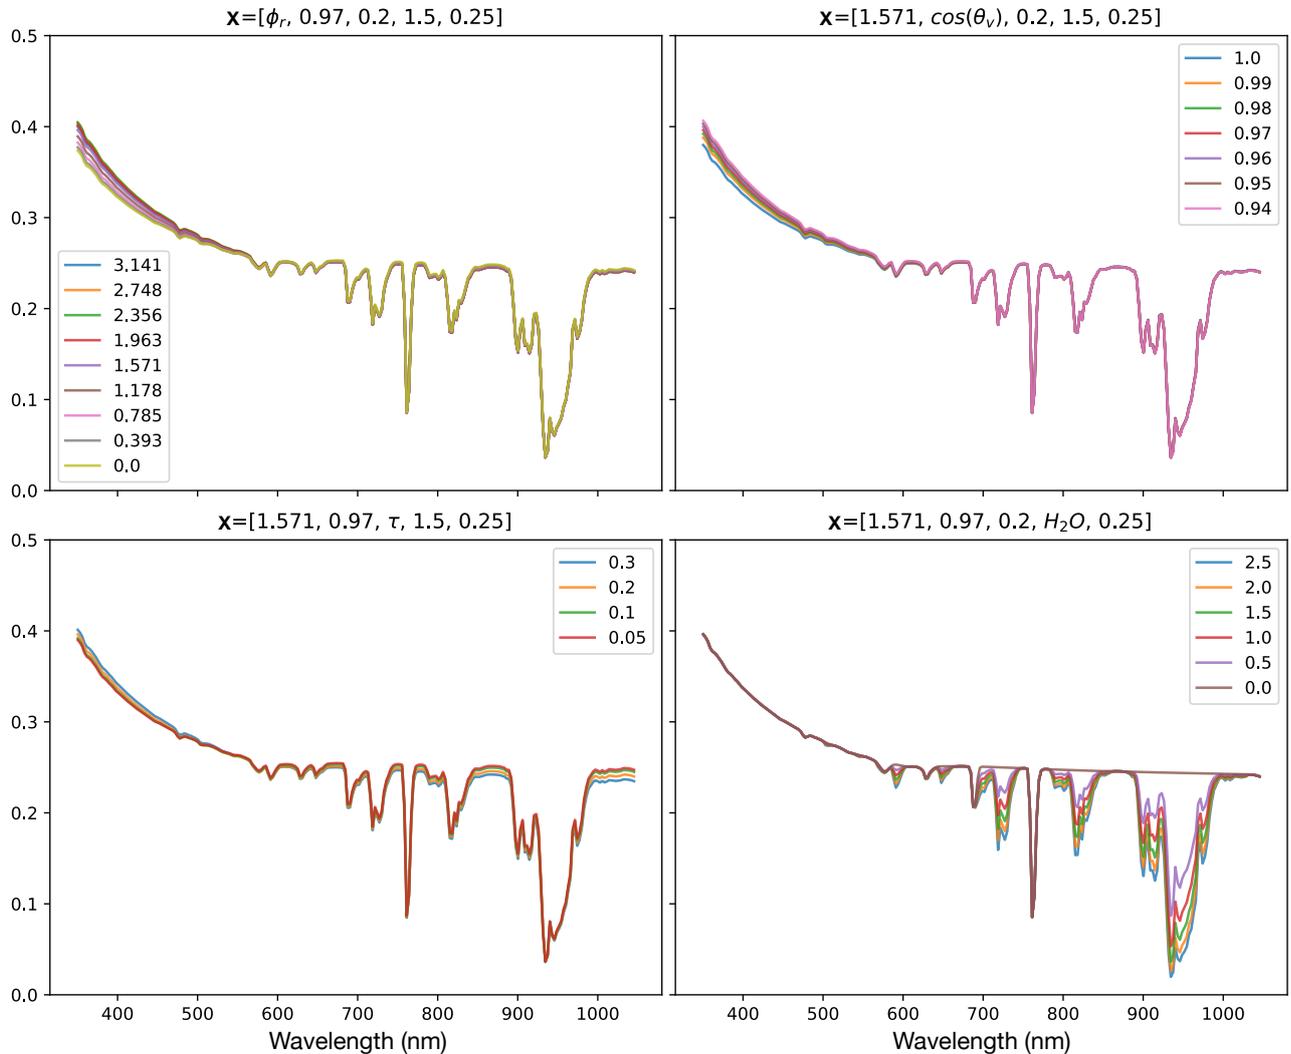


Figure 2. ρ_{obs} spectra for $\rho_s = 0.25$ spanning the range of the ϕ_r (top left), $\cos(\theta_v)$ (top right), τ (bottom left) and H_2O (bottom right) parameters with respect to the validation grid values.

subnetworks initialized with weight propagation typically require less than 30 seconds to converge. However, we note that the channelwise subnetworks trained with weight propagation converge as quickly in the 925-975nm range – where their most significant prediction errors occur – as on the remaining channels.

Investigating further, we measured the average prediction error with respect to the interactions among ρ_s and the four auxiliary parameters. Figure 5 shows the error surfaces that characterize the pairwise interactions among the state parameters. Figure 5 indicates the majority of the parameter space is relatively smooth, with the most significant variability emerging from

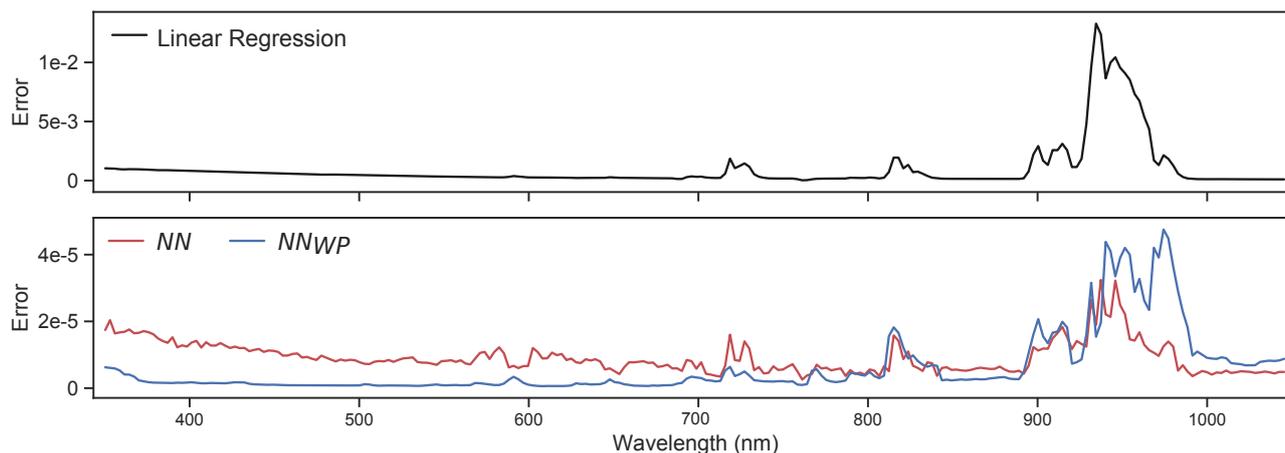


Figure 3. Top: channelwise loss in ρ_{obs} predictions using channelwise Linear Regressors (black line). Bottom: neural network prediction error using subnetworks trained from scratch (NN , red line) versus subnetworks trained weight propagation (NN_{WP} , blue line).

a small range of values in the $\rho_s \in [0.4, 0.8]$ and $H_2O \in [1.0, 2.0]$ region. The relatively high error in this regime is consistent with our earlier observation that small changes in the atmospheric water vapor parameter yield considerably different ρ_{obs} spectra, as shown in Figure 2, and the comparatively high prediction errors for the H_2O absorption bands shown in Figure 3.

4 Atmospheric Correction with the Neural RTM Emulator

- 5 We now evaluate the neural RTM emulator in the context of a surface/atmosphere retrieval problem, retrieving surface reflectance for comparison to known surface materials. To that end, we fused the Optimal Estimation (OE) formalism of Rodgers

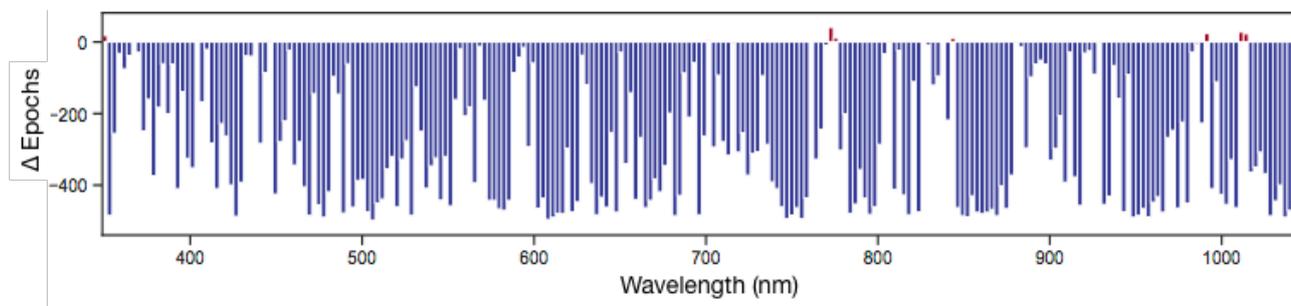


Figure 4. Difference in training epochs to converge to 0.1% validation error. Negative values (blue bars) show channelwise subnetworks that converged faster using weight propagation, while positive values (red bars) indicate channels whose subnetworks converged more quickly when trained from scratch.

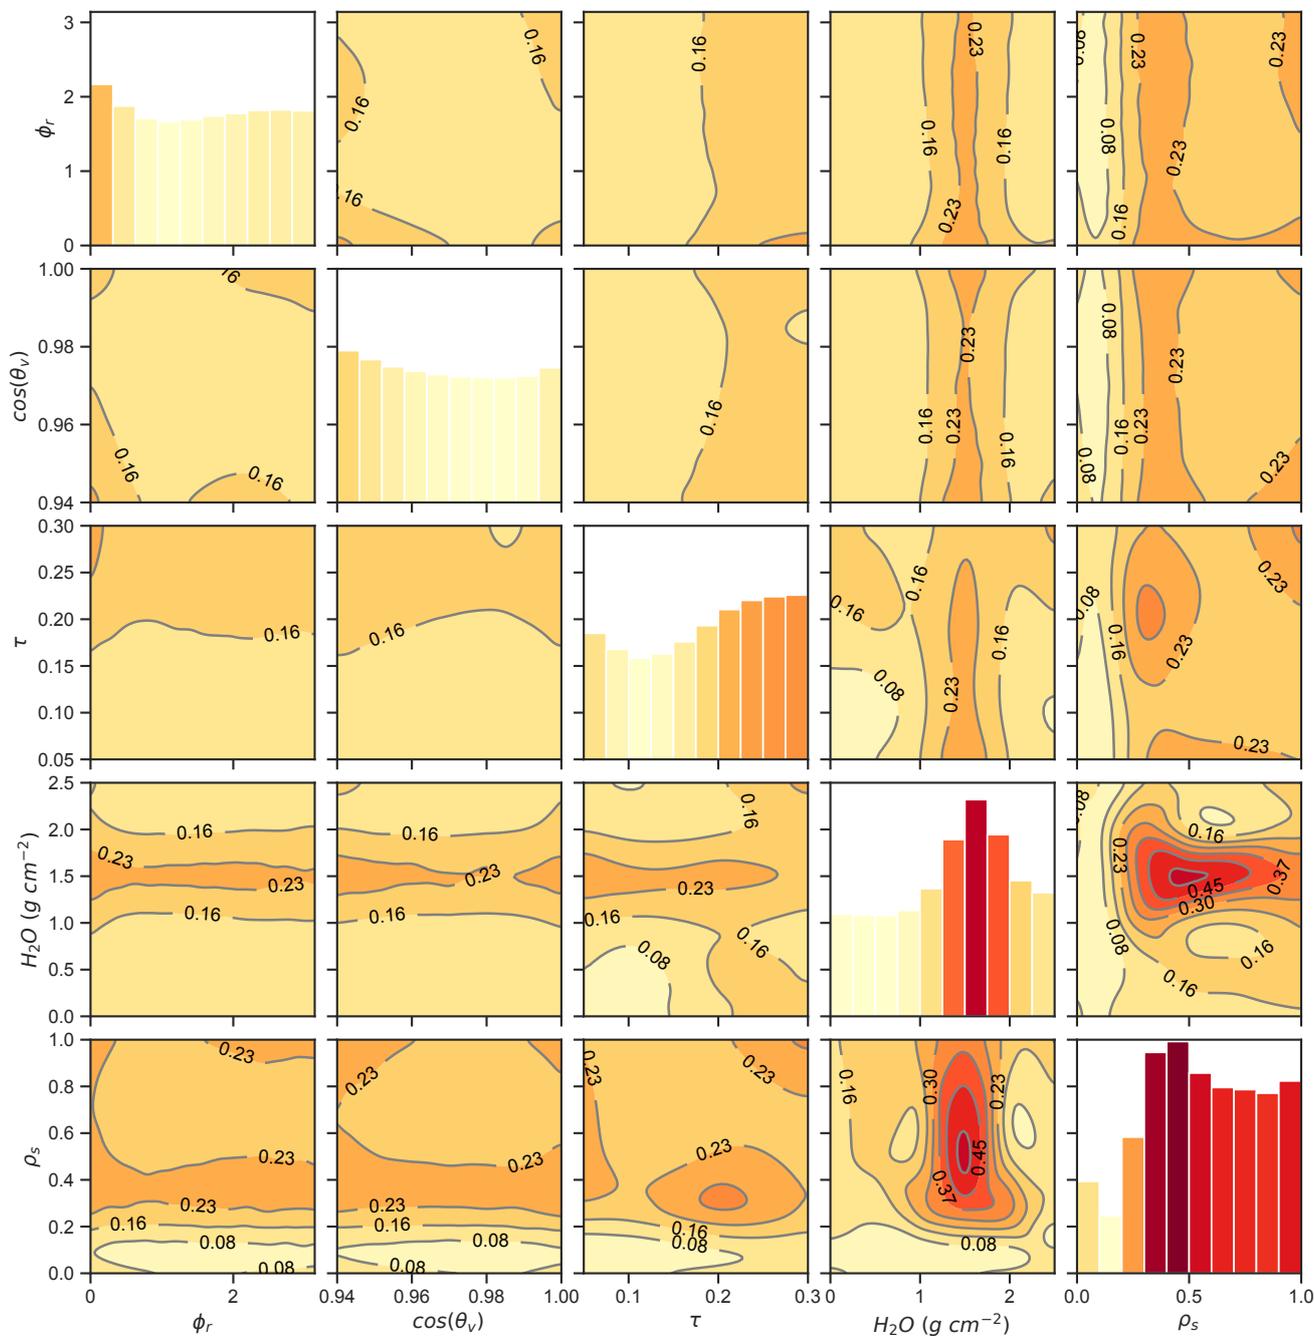


Figure 5. Pairwise error surfaces showing the mean squared ρ_{obs} validation error with respect to the grid values specified in Table 1. The diagonal shows the average error for ten uniformly spaced bins spanning the range of each parameter specified on the x-axis. The y-axes of the diagonal plots are scaled in proportion to the maximum observed error for all pairs.



(2000), following the specific approach of Thompson et al. (2018c) for application to imaging spectroscopy. The OE method estimates the atmosphere and surface state vector by an iterative least-squares optimization of the forward model's match to the measured radiances. Cost terms related to observation error and prior probabilities of state vector elements ensure rigorous propagation of uncertainties in the retrieval.

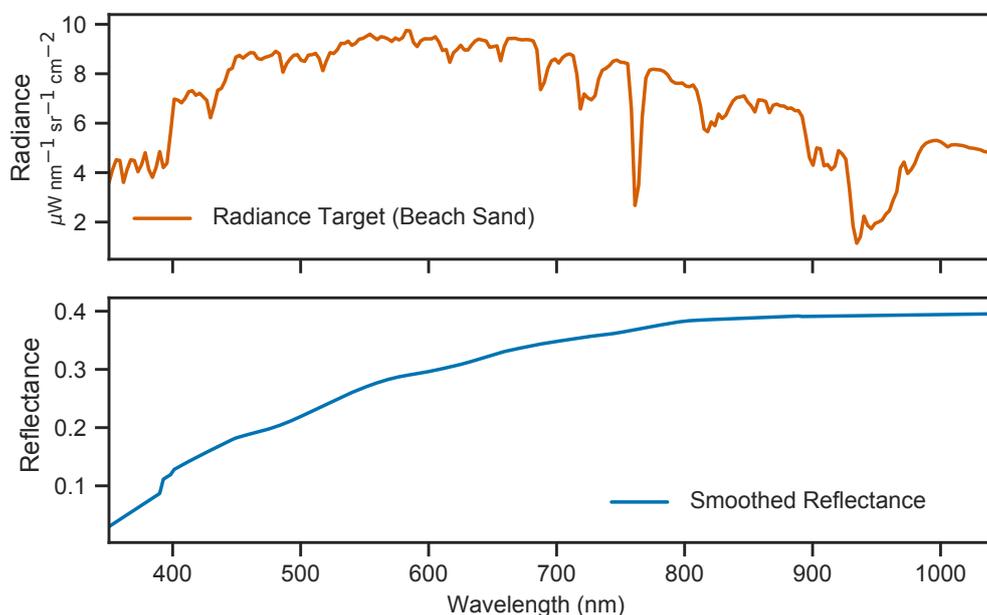


Figure 6. PRISM beach sand radiance (top) and smoothed reflectance target (bottom) used in the vicarious calibration procedure.

5 Continuing our case study, we begin by computing radiometric calibration factors for the PRISM flightline via vicarious calibration. This procedure, similar to standard practice calibration for imaging spectrometers (Thompson et al., 2018a), projects the residual error in retrieved reflectance back into radiance space where it becomes a multiplicative correction factor applied independently to each channel. We generate a “standardized” reflectance target by performing a first-principles retrieval on a manually selected beach sand spectrum from the target flightline. We smooth the resulting (noisy) reflectance spectrum to
10 suppress significant atmospheric features, and use the smoothed target reflectance to generate radiometric correction factors appropriate to the flightline. For reference, the input radiance spectrum and smoothed reflectance target are shown in Figure 6. Applying the resulting factors fine tunes the calibration for optimal results, and smooths residual errors caused by uncertainty in spectral response or RTM inaccuracy.

We applied the atmospheric correction procedure to a set of radiance spectra from the PRISM flightline representing a
15 diverse range of surface materials including grass, rooftop materials, soil, and sea foam. Figure 7 shows a successful retrieval result for a radiance spectrum representing grass on a golf course fairway. The inversion (orange line) perfectly matches the measured radiance (black dashed line) in the left panel. In the right panel, the estimated surface reflectance (blue line) is an

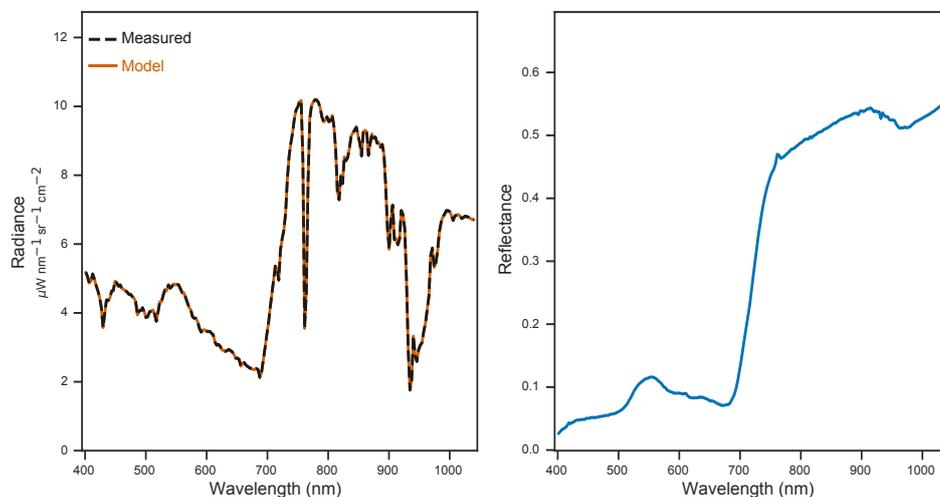


Figure 7. Example reflectance retrieval for a PRISM vegetation spectrum. The left subplot shows the measured (red) versus predicted (black) radiance spectra. The right subplot shows the retrieved reflectance spectrum (black).

extremely smooth and faithful estimate of a dark vegetation spectrum. Figure 8 shows additional radiance spectra (top), and their corresponding reflectance retrievals (bottom). The high quality reflectance estimates — evidenced by the lack of residual bumps caused by atmospheric absorption and the flat, low reflectance profiles in the aerosol-dominated interval from 400-450nm provide additional confidence in the network’s value for atmospheric correction. Our neural RTM emulator runs in less than 5 milliseconds per spectrum, whereas traditional RTMs required an average of over 10 minutes to run. This represents a five order of magnitude improvement in computational runtime.

5 Conclusions

Neural network RTM emulation offers a path to reduce both interpolation inaccuracy while simultaneously reducing runtime by several orders of magnitude. A well-parametrized neural RTM is capable of modeling state parameter spaces with significantly higher accuracy than conventional lookup table-based approaches. Such high capacity statistical models have potential for modeling state parameter spaces with much higher dimensionality than current methods.

The computational and theoretical advantages of a neural RTM are of particular benefit to iterative approaches that must recalculate the entire forward model many times for each spectrum. Equipping iterative formalisms such as optimal estimation with the neural RTM forward model also enables new retrieval approaches that jointly estimate surface and atmospheric parameters. Joint retrieval of surface and atmospheric parameters carries several advantages. It becomes possible to estimate arbitrary parameters of the atmospheric state simply by adjusting the RTM dynamically during the fitting process. A joint retrieval can represent strong coupling between surface and atmosphere, including Bidirectional Reflectance Distribution Function (BRDF) effects, and obviates parametric approximations. The ability to model strong coupling is particularly important for conditions

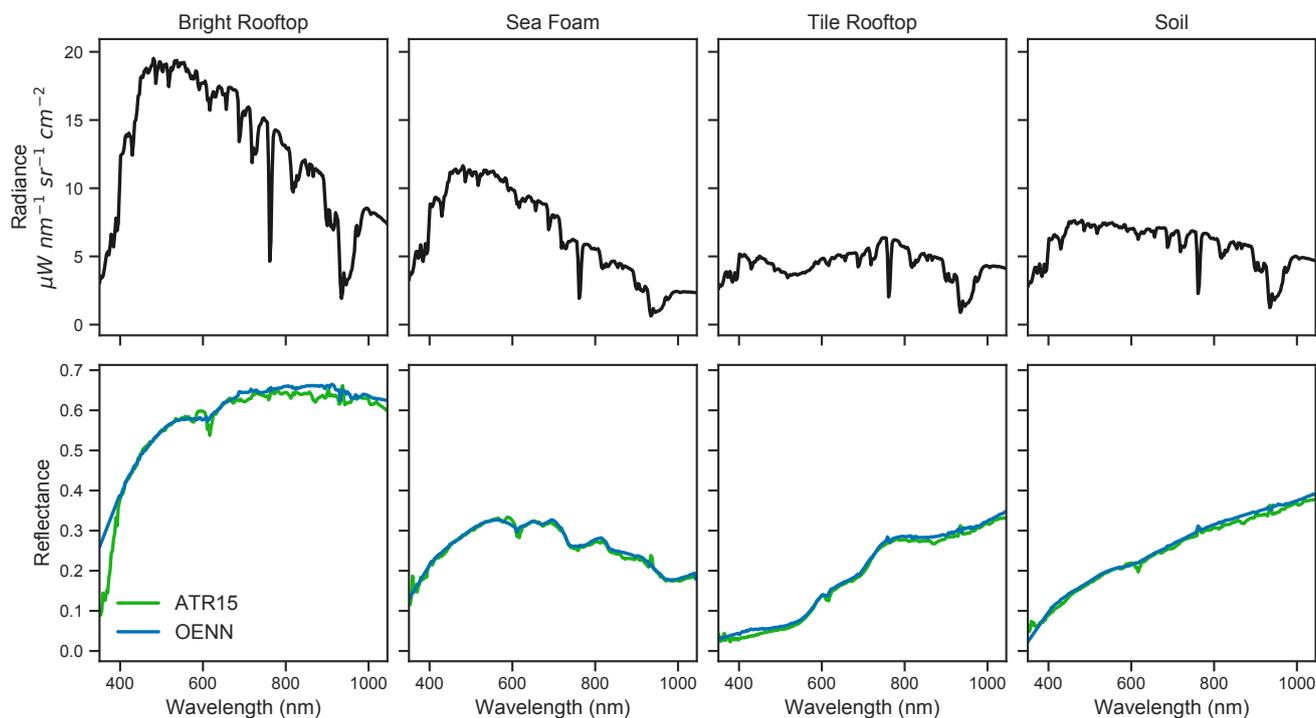


Figure 8. Selected radiance spectra (top) and corresponding reflectance retrievals (bottom) using the ATREM-based atmospheric correction approach of Thompson et al. (2015) (ATR15, green spectra) versus Optimal Estimation equipped with our neural network RTM emulator as the forward model (OENN, blue spectra).

with off-nadir views or haze. Finally, a combined model enables a rigorous, unified and quantitative treatment of uncertainty, respecting uncertainties in all measurement processes and modeled variables, and propagating posterior uncertainties for downstream analysis.

Our results also demonstrate the advantages of informed sampling of the state space. Finer grid sampling in rapidly varying regions of the state space is advantageous to capture complex and often nonlinear interactions among state parameters, while coarse sampling is beneficial in regions of the state space that vary smoothly to reduce redundancy and computational overhead. Uninformed sampling of the state space may not only lead to inaccurate models, but can also yield overly optimistic or inconsistent results when measuring test accuracy or convergence time during cross-validation. For example, as Figures 2 and 5 indicate, much of the state space is relatively smooth. Traditional cross-validation strategies that randomly partition the state space into training and test sets will indicate the subnetworks generalize well due to sampling bias in regions of the state space that are easy to model. Sample stratification approaches during cross-validation can help to ensure each subnetwork accurately captures the parameters that are more difficult to model. However, an informed sampling of the state space would not only eliminate the need for sample stratification during cross-validation, but would also ultimately yield more accurate models with reduced computational overhead.



Future work will train a “universal” neural RTM designed to generate higher spectral resolution predictions that captures a comprehensive range of atmospheric and viewing parameters using more well-informed sampling of the parameter space. We also aim to reduce approximation error still further, in order to keep the fractional contribution small for very dark and/or noisy targets, and are considering reparameterizing the model to retrieve additional aerosol optical properties.

- 5 *Code availability.* The python code used to train and apply the neural RTM for Optimal Estimation-based atmospheric correction are available at the following URL: <https://github.com/dsmbgu8/isofit/>.

Author contributions. Authors Brian Bue and David Thompson conceived, implemented and described the methodology and experiments described in this manuscript, and are the primary contributors to this work. Shubhankar Deshpande contributed data analysis and collection during a Caltech Summer Undergraduate Research Fellowship (SURF) internship at JPL in summer 2017. Terry Mullen contributed analysis and generated data during a Caltech SURF internship at JPL in summer 2018. Michael Eastwood and Robert Green provided expertise in applications and theory of imaging spectroscopy and the application of atmospheric correction techniques. Vijay Natraj provided radiative transfer modeling code and input regarding the theory and application radiative transfer models. Mario Parente provided support for Terry Mullen.

Competing interests. The authors declare that they have no conflict of interest.

15 *Acknowledgements.* We thank Sven Geier, Scott Nolte, and the AVIRIS-NG instrument and Science Data System teams for assistance in calibration and operations. Particular thanks go to Charles E. Miller and Michael Turmon whose feedback significantly improved the manuscript. We are also thankful for the counsel of colleagues including Phil Townsend, Phil Dennison, Dar Roberts, Steven Adler-Golden, Alexander Berk, Steven Massie, and Bruce Kindel. We acknowledge the support of the NASA Earth Science Division for the AVIRIS-NG instrument and the data analysis program “Utilization of Airborne Visible/Infrared Imaging Spectrometer Next Generation Data from an Airborne Campaign in India” NNH16ZDA001N-AVRSNG, managed by Woody Turner, for its support of the algorithm development. We are also thankful for the support of the Jet Propulsion Laboratory Research and Technology Development Program, the NASA Center Innovation Fund managed in conjunction with the Jet Propulsion Laboratory Office of the Chief Scientist and Technologist, and the Caltech SURF program. A portion of this research took place at the Jet Propulsion Laboratory, California Institute of Technology. US Government Support Acknowledged.



References

- Asner, G., Martin, R., Knapp, D., Tupayachi, R., Anderson, C., Sinca, F., Vaughn, N., and Llactayo, W.: Airborne laser-guided imaging spectroscopy to map forest trait diversity and guide conservation, *Science*, 355, 385–389, 2017.
- Brajard, J., Jamet, C., Moulin, C., and Thiria, S.: Use of a neuro-variational inversion for retrieving oceanic and atmospheric constituents from satellite ocean colour sensor: Application to absorbing aerosols, *Neural Networks*, 19, 178–185, <https://doi.org/10.1016/j.neunet.2006.01.015>, 2006.
- Emde, C., Buras-Schnell, R., Kylling, A., Mayer, B., Gasteiger, J., Hamann, U., Kylling, J., Richter, B., Pause, C., Dowling, T., and Bugliaro, L.: The libRadtran software package for radiative transfer calculations (version 2.0.1), *Geoscientific Model Development*, 9, 1647–1672, 2016.
- 10 ESAS: Thriving on Our Changing Planet: A Decadal Strategy for Earth Observation from Space. A Report by the Decadal Survey on Earth Science and Applications from Space, The National Academies Press, Washington, DC, available online at <http://sites.nationalacademies.org/DEPS/esas2017/index.htm> (accessed Jan. 2018), 2018.
- Fichot, C. G., Downing, B. D., Bergamaschi, B. A., Windham-Myers, L., Marvin-DiPasquale, M., Thompson, D. R., and Gierach, M. M.: High-Resolution Remote Sensing of Water Quality in the San Francisco Bay–Delta Estuary, *Envir. Sci. Tech.*, 50, 573–583, 2015.
- 15 Frankenberg, C., Thorpe, A. K., Thompson, D. R., Hulley, G., Kort, E. A., Vance, N., Borchardt, J., Krings, T., Gerilowski, K., Sweeney, C., Conley, S., Bue, B. D., Aubrey, A. D., Hook, S., and Green, R. O.: Airborne methane remote measurements reveal heavy-tail flux distribution in Four Corners region, *Proceedings of the National Academy of Sciences*, 113, 9734–9739, 2016.
- Gao, B. C., Heidebrecht, K. B., and Goetz, A. F.: Derivation of scaled surface reflectances from AVIRIS data, *Remote Sensing of Environment*, 44, 165–178, 1993.
- 20 Gao, B.-C., Montes, M. J., Ahmad, Z., and Davis, C. O.: Atmospheric correction algorithm for hyperspectral remote sensing of ocean color from space, *Applied Optics*, 39, 887–896, 2000.
- Gao, B. C., Montes, M. J., Li, R. R., Dierssen, H. M., and Davis, C. O.: An atmospheric correction algorithm for remote sensing of bright coastal waters using MODIS land and ocean channels in the solar spectral region, *IEEE Transactions on Geoscience and Remote Sensing*, 45, 1835–1843, 2007.
- 25 Glorot, X. and Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks., *AISTATS*, <https://dblp.org/rec/journals/jmlr/GlorotB10>, 2010.
- Gupta, M., Cotter, A., Pfeifer, J., Voevodski, K., Canini, K., Mangylov, A., Moczydlowski, W., and van Esbroeck, A.: Monotonic Calibrated Interpolated Look-Up Tables, *arXiv.org*, p. [arXiv:1505.06378](https://arxiv.org/abs/1505.06378), 2015.
- Hochberg, E. J.: Remote sensing of coral reef processes, *Coral Reefs: An Ecosystem in Transition*, pp. 25–35, 2011.
- 30 Jamet, C., Thiria, S., Moulin, C., and Crepon, M.: Use of a Neurovariational Inversion for Retrieving Oceanic and Atmospheric Constituents from Ocean Color Imagery: A Feasibility Study, *Journal of Atmospheric and Oceanic Technology*, 22, 460–475, <https://doi.org/10.1175/JTECH1688.1>, 2005.
- Jamet, C., Loisel, H., and Dessailly, D.: Retrieval of the spectral diffuse attenuation coefficient $K_d(\lambda)$ in open and coastal ocean waters using a neural network inversion, *Journal of Geophysical Research*, 117, <https://doi.org/10.1029/2012JC008076>, 2012.
- 35 Jetz, W., Cavender-Bares, J., Pavlick, R., Schimel, D., Davis, F. W., Asner, G. P., Guralnick, R., Kattge, J., Latimer, A. M., Moorcroft, P., et al.: Monitoring plant functional diversity from space, *Nature plants*, 2, 2016.
- Kingma, D. P. and Ba, J.: Adam: A method for stochastic optimization, *arXiv preprint arXiv:1412.6980*, 2014.



- Kruse, F.: Comparison of ATREM, ACORN, and FLAASH atmospheric corrections using low-altitude AVIRIS data of Boulder, CO., 13th JPL Airborne Geoscience Workshop, Jet Propulsion Laboratory Publication 05–3, 10, 2004.
- Martino, L., Vicent, J., and Camps-Valls, G.: Automatic emulator and optimized look-up table generation for radiative transfer models., IEEE International Geoscience And Remote Sensing Symposium, <https://dblp.org/rec/conf/igarss/MartinoVC17>, 2017.
- 5 Mayer, B. and Kylling, A.: Technical note: The libRadtran software package for radiative transfer calculations - description and examples of use, *Atmospheric Chemistry and Physics*, 5, 1855–1877, 2005.
- Mouroulis, P., Green, R. O., and Wilson, D. W.: Optical design of a coastal ocean imaging spectrometer, *Opt. Express*, 16, 9087–9096, 2008.
- Mouroulis, P., Gorp, B. V., Green, R. O., Dierssen, H., Wilson, D. W., Eastwood, M., Boardman, J., Gao, B.-C., Cohen, D., Franklin, B., Loya, F., Lundeen, S., Mazer, A., McCubbin, I., Randall, D., Richardson, B., Rodriguez, J. I., Sarture, C., Urquiza, E., Vargas, R., White,
- 10 V., and Yee, K.: Portable Remote Imaging Spectrometer coastal ocean sensor: design, characteristics, and first flight results, *Appl. Opt.*, 53, 1363–1380, 2014.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E.: Scikit-learn: Machine Learning in Python, *Journal of Machine Learning Research*, 12, 2825–2830, 2011.
- 15 Perkins, T., Adler-Golden, S., Matthew, M., Berk, A., Bernstein, L., Lee, J., et al.: Speed and accuracy improvements in FLAASH atmospheric correction of hyperspectral imagery, *Optical Engineering*, 51, 111 707–1–111 707–7, 2012.
- Richter, R. and Schlapfer, D.: Geo-atmospheric processing of airborne imaging spectrometry data, Part 2: atmospheric/topographic correction., *International Journal of Remote Sensing*, 23, 2631–2649, 2002.
- Rodgers, C. D.: *Inverse methods for atmospheric sounding: theory and practice*, vol. 2, World scientific, 2000.
- 20 Schaeppman, M. E., Ustin, S. L., Plaza, A. J., Painter, T. H., Verrelst, J., and Liang, S.: Earth system science related imaging spectroscopy—An assessment, *Remote Sensing of Environment*, 113, S123–S137, 2009.
- Schaeppman-Strub, G., Schaeppman, M. E., Painter, T. H., Dangel, S., and Martonchik, J. V.: Reflectance quantities in optical remote sensing—Definitions and case studies, *Remote sensing of environment*, 103, 27–42, 2006.
- Snoek, J., Larochelle, H., and Adams, R. P.: Practical Bayesian Optimization of Machine Learning Algorithms, *arXiv.org*, p. arXiv:1206.2944, <http://papers.nips.cc/paper/4522-practical>, 2012.
- 25 Stamnes, K., Tsay, S.-C., Wiscombe, W., and Jayaweera, K.: Numerically stable algorithm for discrete-ordinate-method radiative transfer in multiple scattering and emitting layered media, *Applied optics*, 27, 2502–2509, 1988.
- Stein, M.: Large sample properties of simulations using Latin hypercube sampling, *Technometrics*, 29, 143–151, 1987.
- Swayze, G. A., Clark, R. N., Goetz, A. F., Livo, K. E., Breit, G. N., Kruse, F. A., Sutley, S. J., Snee, L. W., Lowers, H. A., Post, J. L., Stoffregen, R. E., and Ashley, R. P.: Mapping Advanced Argillic Alteration at Cuprite, Nevada, Using Imaging Spectroscopy, *Economic Geology*, 109, 1179, 2014.
- 30 Thompson, D. R., Gao, B.-C., Green, R. O., Roberts, D. A., Dennison, P. E., and Lundeen, S. R.: Atmospheric correction for global mapping spectroscopy: ATREM advances for the HypSIRI preparatory campaign, *Remote Sensing of Environment*, 167, 64 – 77, 2015.
- Thompson, D. R., Boardman, J. W., Eastwood, M. L., Green, R. O., Haag, J. M., Mouroulis, P., and Gorp, B. V.: Imaging spectrometer stray spectral response: In-flight characterization, correction, and validation, *Remote Sensing of Environment*, 204, 850–860, <https://doi.org/10.1016/j.rse.2017.09.015>, 2018a.
- Thompson, D. R., Cawse-Nicholson, K., Erickson, Z., Fichot, C. G., Frankenberg, C., Gao, B.-C., Gierach, M. M., Green, R. O., Natraj, V., and Thompson, A.: Optimal Estimation for Coastal Ocean Imaging Spectroscopy, *Remote Sensing of Environment*, (in review), 2018b.



- Thompson, D. R., Natraj, V., Green, R. O., Helmlinger, M. C., Gao, B.-C., and Eastwood, M. L.: Optimal estimation for imaging spectrometer atmospheric correction, *Remote Sensing of Environment*, 216, 355 – 373, <https://doi.org/10.1016/j.rse.2018.07.003>, 2018c.
- Trinh, R. C., Fichot, C. G., Gierach, M. M., Holt, B., Malakar, N. K., Hulley, G., and Smith, J.: Application of Landsat 8 for Monitoring Impacts of Wastewater Discharge on Coastal Water Quality, *Frontiers in Marine Science*, 4, <https://doi.org/10.3389/fmars.2017.00329>, 5 2017.
- Ustin, S. L., Roberts, D. A., Gamon, J. A., Asner, G. P., and Green, R. O.: Using imaging spectroscopy to study ecosystem processes and properties, *BioScience*, 54, 523–534, 2004.
- Vermote, E. F., Tanré, D., Deuze, J. L., Herman, M., and Morcette, J. J.: Second simulation of the satellite signal in the solar spectrum, 6S: An overview, *Geoscience and Remote Sensing, IEEE Transactions on*, 35, 675–686, 1997.
- 10 Verrelst, J., Sabater, N., Rivera, J. P., Muñoz-Marí, J., Vicent, J., Camps-Valls, G., and Moreno, J.: Emulation of Leaf, Canopy and Atmosphere Radiative Transfer Models for Fast Global Sensitivity Analysis, *Remote Sensing*, 8, 673, 2016.
- Verrelst, J., Rivera Caicedo, J., Muñoz-Marí, J., Camps-Valls, G., and Moreno, J.: SCOPE-Based Emulators for Fast Generation of Synthetic Canopy Reflectance and Sun-Induced Fluorescence Spectra, *Remote Sensing*, 9, 927, 2017.
- Werbos, P. J.: Applications of advances in nonlinear sensitivity analysis, in: *System modeling and optimization*, pp. 762–770, Springer, 1982.