

Interactive comment on “Neural Network Radiative Transfer for Imaging Spectroscopy” by Brian D. Bue et al.

Response: Many thanks to both reviewers for the helpful comments and for their time reviewing this work. They each identified several fundamental issues that were not clearly addressed in the original manuscript, and their input has substantially improved the quality of this work.

The primary changes we made to address the concerns of both reviewers are summarized below.

- We updated section 1 to more clearly state our novel contributions and their relation to prior work.
- We added additional content to section 2 that formally describes how we train and validate the channelwise subnetworks to construct the neural RTM.
- We added additional content to section 3 that states our modeling assumptions and the specific parameters used to train the channelwise subnetworks.
- We provide a commented LibRadTran config file associated with the PRISM flightline we considered in our case study as a supplement, citing the references therein in detail in the narrative.
- We reorganized the paper in a manner that makes interpreting the content in the provided figures and tables easier to understand.

Our responses to reviewer comments below are highlighted in red text. Our modifications to the text and figures are shown in blue text in the `nnrt_amt_r1_changes.pdf` file.

Anonymous Referee #1

Received and published: 2 February 2019

This paper is a compact presentation of a novel method of the emulation and use of radiative transfer models in compensating for the effects of illumination and the atmosphere to retrieve surface reflectance from measured radiances collected by an imaging spectrometer. It is quite dense and while it presents an adequate description, the results are just limited examples of the accuracy and a demonstration of its use.

It would be nice to see an expanded treatment of the first part of the article, the development of the forward emulation model, which is the main contribution. Section 2 presents the model with many details presented in a dense manner. It would be nice to expand and include more explicit equations for the various steps involved. More discussion on the selection and justification of the approach would also provide more insight.

Response: We added a new section that provides a comprehensive walkthrough of the entire neural RTM training procedure, and also describes the sampling/validation strategy we used to guide the optimization of the channelwise subnetworks.

We also provided more detailed description of the modeling assumptions made and parameters used to train each subnetwork used in our PRISM case study in sections 2 and 3.

In Section 3, the rationale behind limited the parameters of Table 1 would be good to see.

Response: We selected the suite of state parameters to act as a representative set of the types of states observed in operational atmospheric correction settings involving imaging spectroscopy data.

Our updates to the last paragraphs on page 7 clarify and motivate our parameter selection rationale.

In Figure 1, is the ρ in the middle of the figure supposed to ρ_{obs} ? If so please label as such. If not, how does it relate to terms defined in the text?

Response: We added a revised version of Figure 1 that corrects the notational errors in the earlier version, and provided a more descriptive caption for the content shown in the figure.

Figure 5 should have a color legend to better interpret the quantitative results. Also, why show mean squared error? Why not root mean square which can be better interpreted?

Response: The error shown in Figure 5 was actually RMSE, but was incorrectly labeled in the earlier draft. We modified the caption and corresponding text to remedy this issue.

The labels that overlap the contours give the RMSE error level associated with each contour, but we also added a color bar to make this figure a bit easier to parse.

Actually, is this the mean error across the wavelengths? That would mask errors that may be concentrated in the water vapor region. Please clarify.

Response: The RMSE error depicted in Figure 5 is indeed averaged across all wavelengths. However, as we mentioned in our analysis, the high error values associated with increasing water vapor concentrations shown in the H₂O row/column of the figure suggest that errors concentrated in the water absorption bands are not suppressed.

What was the source for the smoothed reflectance shown in Figure 6? Field spectrometer measurements? A library? Did you also use a dark target for the vicarious procedure?

Response: The radiance and the smoothed reflectance spectra shown in Figure 6 correspond to the same PRISM beach sand target. We adjusted the figure caption to clarify this detail.

This work presents an exciting development in the operational use of imaging spectrometers and deserves a more comprehensive presentation. Perhaps this could be the subject of future articles.

Response: thank you!

Anonymous Referee #2

Received and published: 22 February 2019

The paper presents a neural network developed to calculate radiative transfer for the solar spectral region in clear sky. The model is applied to retrieve the surface spectral reflectance function from PRISM (airborne imaging spectrometer) data. As examples the retrieved surface reflectance spectra for 5 different surface types are shown to demonstrate that the method works well. The neural network method highly accelerates the atmospheric correction methodology, thus it is certainly an interesting approach which might become standard since the data to be processed increases rapidly with improvements in spatial and spectral resolution of sensors. The topic of the paper fits well in the scope of AMT, however, it needs to be revised, because the methodology needs to be described more precisely. Further the authors need to point, what makes their approach novel compared to other neural networks based approaches. I recommend publication in AMT after these revisions.

Major comments:

1. The application of neural networks for remote sensing has become rather common in the last decade (besides the mentioned references e.g. Kox et al. 2014, Loyola et al. 2016, Eferemeko et al 2017, Strandgren et al. 2017, Loyola et. al 2018). The authors say that they for the first time apply a neural network based algorithm to imaging spectroscopy. I am not sure whether this is correct. At least, the operational cloud retrieval algorithm for TROPOMI, a spaceborne imaging spectrometer, also applies a neural network RTM (Loyola et al., 2018). So in my opinion the authors should point out, what makes their approach novel and what is the difference to other similar approaches.

Response: We do not claim that this work represents the first time a neural network has been applied to imaging spectroscopy. Rather, we believe that this work represents the first attempt to learn the RTM function $F(x) \rightarrow y$ such that the trained RTM emulator is capable of acting as a forward model in an atmospheric correction procedure, allowing us to retrieve surface reflectance over the entire VSWIR range in variable imaging conditions.

Also, thank you for pointing us to Loyola 2018, which represents the most similar approach to ours. They perform RTM emulation as part of an inversion (as we do), but their approach is designed for atmospheric measurements (i.e., cloud height and fraction) rather than surface retrievals spanning multiple wavelengths.

We rewrote the last two paragraphs in the introduction to clarify these details.

If I understood correctly, the physical knowledge that is incorporated is the "analytical decomposition of the radiance spectrum into quantities that are individually easier to model". Later I find, that this means to use reflectances ($\pi y / \phi_0 E_0$ with y -radiance, ϕ_0 solar

zenith angle and E0 -extraterrestrial irradiance; apart from E0 very uncommon nomenclature) instead of radiance. This is a very common approach, also in look-up-table based retrieval algorithms and not a new idea. I think that all publications mentioned below use reflectances and not radiances.

Response: Thank you for identifying this discrepancy. The quoted sentence was indeed poorly written and contained a rather egregious typo, and we did not intend to insinuate that using reflectances instead of radiances was a novel component of this work. Specifically:

“...analytical decomposition of the radiance spectrum into quantities that are individually easier to model...”

should be

“analytical decomposition of the radiative transfer function $F(x)$ into quantities that are individually easier to model...”

The above description refers to our approach to model the individual wavelengths of spectra generated by a first-principles RTM using separate neural network models. Using TOA reflectance rather than radiance to train a neural network is not (in itself) novel, the primary contribution is the channelwise decomposition of $F(x)$ that is used for predicting TOA responses at multiple wavelengths.

We believe that the updates we made to the last two paragraphs in the introduction address this ambiguity.

The second "novel" idea that is mentioned is the channel-wise training. This is as far as I know already applied in Kox et al. 2014. The only new idea, that I have not seen so far, is the channelwise weight propagation, which makes sense to me.

While it is indeed true that Kox et al. (and many others) train neural networks using multiple channels as input, those approaches are quite different from ours, both in terms of network architecture and derived outputs. Specifically:

- We train k separate neural networks, where each subnetwork computes the mapping from the state space to TOA reflectance for a single wavelength. The combined output of all k networks yields a full TOA reflectance spectrum, and allows us to directly apply the trained neural RTM to act as a forward model in an RTM-based atmospheric correction procedure.
- In contrast, Kox et al. construct a single neural network using a combination of k radiance channels + several auxiliary parameters to derive two atmospheric parameters (cloud optical thickness and top altitude).

To more clearly distinguish our approach from prior work, we added Algorithm 1 and several new paragraphs describing the end-to-end training + validation methodology we used.

2. The setup of the neural network needs a more comprehensive description:

- Most symbols in Figure 1 are not explained

Response: We added a revised version of Figure 1 that corrects the notational errors in the earlier version, and provided a more descriptive caption for the content shown in the figure.

- It is not clear to me, whether the "monochromatic subnetworks" (p4, l.4) are really monochromatic (for 1 wavelength) or for one channel (convolved with the spectral response function). E.g. on p5, l4 it reads "To ensure each subnetwork reliably models its corresponding channel ..."

Response: In principle, the neural RTM is designed to compute the mapping from the state space for each wavelength of the TOA reflectance spectra generated by the RTM (LibRadTran). However, the neural RTM we demonstrate in this work generates TOA reflectance spectra at (PRISM) instrument wavelengths, and the LibRadTran TOA spectra are indeed resampled with respect to the PRISM spectral response function to match those wavelengths. Consequently, they are not strictly monochromatic, but we decided that the dramatic reduction in computation made this an acceptable compromise for this proof of concept, and plan to construct a more general neural RTM at the same spectral resolution as the LibRadTran outputs in future work.

Our updates on p9, 1st paragraph describe these modeling assumptions + the associated benefits compromises of this approach in detail.

- Explain settings: "feed-forward architecture with two hidden layers of 50 units each ... Why is the neural network set up as shown in Figure 1? Why are these particular settings chosen for this particular application? Explain terms "rectifying" and "early stopping"

Response: Rectifying linear unit (RELU) activations were introduced by [Nair, 2010] and simply involve clipping the negative portion of the input (i.e., $f(x) = \max(0, x)$). RELU activations have been shown to be more robust than the conventional sigmoid / tanh activations used historically in neural networks. We updated the text to state this detail.

Early stopping the process of training each neural network until some criteria are met, for instance when test accuracy plateaus for several cycles. We removed any references to this item, as we explain our stopping conditions in more detail in section 3.

Our updates starting on p9, 2st paragraph describe the specific parameters we used to train each subnetwork, along with context motivating our design decisions.

3. The radiative transfer simulations for training and validation were performed using the libRadtran package, which is a comprehensive package including various parameterizations, atmospheric setups, aerosol types etc. The setup of libRadtran needs to be specified, i.e. which atmosphere is used, which absorption parameterization, how is the aerosol defined, which radiative transfer solver is used etc.

Response: We provided a commented libradtran configuration file (prm20151026t173148_libradtran_config) associated with the PRISM flightline described in our case study as a supplemental file, and describe the atmosphere used in the last paragraph of page 7.

The symbols in Table 1 (input parameters) are also not defined.

Response: the symbols in Table 1 were defined in the first paragraph of Section 2, but we agree that the table did not clearly reference the associated symbols. We moved this paragraph closer to the table, and added text labels for each symbol listed in the table for easier interpretation.

4. Figure 3: How is the linear regression done? Linear interpolation in a lookup-table for the parameters listed in Table 1? This is not explained in the text.

Response: We trained a single least squares linear regressor for each channel using the same set of training samples used to train the subnetwork associated with that channel. We then computed the test error per channel for each of trained regressor on the same set of test samples as the subnetwork. The channelwise test errors using the linear regressors provide an approximate upper bound of the error that would be incurred using piecewise / locally linear interpolation to infer TOA responses for intermediate states based on lookup tables.

These details are reflected in our updates to the first paragraph on page 10.

5. Why are only 5 examples of retrieved surface reflectivity spectra shown. Since the method is so fast, it could be easily be applied on a full image. The difference between the traditional atmospheric correction approach and the new neural network approach should be assessed in a statistical sense.

Response: Our focus in this work was to provide a detailed procedure to construct an RTM emulator that can be used as a forward model yielding high quality surface reflectance retrievals. To that end, we felt a qualitative analysis of our reflectance retrievals on spectra for which the surface materials were known was the most important component to demonstrate a functional proof of concept. We are currently applying the methodology described in this work in operational settings for entire images, and plan to describe those results in more detail in future publications.

Minor comments:

- Figure 5: I suggest to create a separate figure for the histograms shown in the diagonal with proper y-axis labels.

Response: Removing the histograms from the diagonal would leave large empty areas with no useful content, and we believe that associating each parameter on the diagonal with the error surfaces for the other parameters in each column provides informative context.

However, your point regarding the disconnect between the range of error values for the histograms and the error surfaces is well taken. To address this, we added text labels to the diagonal subplots specifying the minimum + maximum error values and placed those labels on the bar of the corresponding histogram bin in Figure 5.

- It is confusing that the term reflectance is used for the top-of-atmosphere reflectance and the surface-reflectance. It should always be specified, which reflectance is meant

Response: We updated each reference to “reflectance” to specify whether we refer to “surface reflectance” or “TOA reflectance.”

p11, l.5: traditional RTMs required over 10 minutes to run -> this depends very much on the settings, on the radiative transfer solution method etc. Specify, which kind of RTMs you refer to.

Response: Good point – The updates we added on page 9 specify the assumptions we made with respect to the RTM model, and we provided additional information in the last paragraph of page 11 regarding the runtime of our model versus the first-principles RTM. In particular, we noted the runtime of each approach with respect to the number of wavelengths in the output spectrum. Also, the libradtran config file we provide as a supplement should provide context regarding the RTM we used in this work.

References:

- Kox, S.; Bugliaro, L.; Ostler, A., Retrieval of cirrus cloud optical thickness and top altitude from geostationary remote sensing ATMOSPHERIC MEASUREMENT TECHNIQUES Volume: 7 Issue: 10 Pages: 3233-3246 Published: 2014
- Loyola, Diego G. R.; Pedergrana, Mattia; Garcia, Sebastian Gimeno Smart sampling and incremental function learning for very large high dimensional data NEURAL NETWORKS Volume: 78 Special Issue: SI Pages: 75-87 Published: JUN 2016
- Efremenko, Dmitry S.; Loyola, Diego G.; Hedelt, Pascal; et al. Volcanic SO2 plume height retrieval from UV sensors using a full-physics inverse learning machine algorithm INTERNATIONAL JOURNAL OF REMOTE SENSING Volume: 38 Supplement: 1 Pages: 1-27 Published: 2017
- Strandgren, Johan; Fricker, Jennifer; Bugliaro, Luca, Characterisation of the artificial neural network CiPS for cirrus cloud remote sensing with MSG/SEVIRI ATMOSPHERIC MEASUREMENT TECHNIQUES Volume: 10 Issue: 11 Pages: 4317-4339 Published: NOV 14 2017

Loyola, Diego G.; Garcia, Sebastian Gimeno; Lutz, Ronny; et al. The operational cloud retrieval algorithms from TROPOMI on board Sentinel-5 Precursor ATMOSPHERIC MEASUREMENT TECHNIQUES Volume: 11 Issue: 1 Pages: 409-427 Published: JAN 18 2018

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 and generate accurate radiance spectra at multiple wavelengths over a diverse range of surface/atmosphere state parameters. We also demonstrate such models can act as surrogate forward models for atmospheric correction procedures.](#) 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 (Schaeppman 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 (Schaeppman-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 image acquisition. 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 other methods (e.g., Jamet et al. (2012); Kox et al. (2014); Loyola et al. (2018)) permit retrieval of atmospheric or radiometric parameters based on models constructed using outputs generated by first-principles RTMs that span multiple wavelengths. However, to date, techniques designed to retrieve surface reflectance using learned RTM emulators have only 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), and not across the VSWIR range with state vector flexibilities that would permit a functionally-useful alternative for existing atmospheric correction routines (e.g., as a surrogate forward model). To our knowledge, this work represents the first demonstration on actual imaging spectroscopy data using nonparametric function approximation to emulate the RTM function $F(\mathbf{x}) \rightarrow y$ such that the RTM emulator is capable of acting as a forward model in an atmospheric correction procedure, thereby allowing us to retrieve surface reflectance over the entire VSWIR range for variable imaging conditions.

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 traditional 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 radiative transfer function $F(\mathbf{x})$ into quantities that are individually easier to model, *channelwise*, *monochromatic subnetworks* to simplify training and prediction, and *weight propagation* to account for correlation between adjacent channels and to reduce training time. We also describe an approach to partition the state space in a manner that guides each subnetwork to generate accurate predictions for states within the bounds of the state space. We evaluate our approach in a case study focusing on atmospheric correction for the PRISM imaging spectrometer, and demonstrate high quality surface reflectance retrievals using the optimal estimation approach of Thompson et al. (2018c) equipped with our neural RTM as the forward model. 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 RTM emulation technology.

2 Neural Networks for Radiative Transfer Modeling

Our goal is to construct a model that emulates a first-principles RTM using precalculated outputs generated by that RTM for a representative set of atmospheric, geometric, and surface states. More formally, we aim to model the RTM function $F(\mathbf{x}) \rightarrow \mathbf{y}$ that maps a set of m distinct state parameters $\{p_j\}_{j=1}^m$ with values captured in a state vector $\mathbf{x} \in \mathbb{R}^m$ to a vector $\mathbf{y} \in \mathbb{R}^n$ of observed at-sensor radiances for k channels centered at wavelengths $\{\lambda_1, \dots, \lambda_k\}$. We use boldface notation to signify vectors and matrices, with all matrices in capital letters.

We exploit two features of the problem to simplify $F(\mathbf{x})$. First, we leverage the fact that the observed radiance at any given channel is fully specified by the observation geometry, atmospheric state, and the surface reflectance *in that channel*. In statistical terms, absent any prior distribution that couples neighboring wavelengths, the channelwise radiances become conditionally independent of each other given the atmosphere and observation geometry. This permits an exact decomposition of $F(\mathbf{x})$ into *monochromatic functions* $F(\mathbf{x}) = f_i(\mathbf{x})_{i=1}^k$, where each $f_i(\mathbf{x}) \rightarrow y_i$ represents the RTM function for the channel centered at wavelength λ_i . Given this decomposition, we construct a *neural RTM emulator using a set of k channelwise subnetworks*, where each subnetwork is trained to model a single f_i . Figure 1 shows the topology of one of the channelwise subnetworks in the neural RTM. 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).

Second, we reduce the radiance spectrum analytically to the top of atmosphere reflectance, written ρ_{obs} , and solar illumination components. The top of atmosphere reflectance is defined as $\rho_{obs} = \mathbf{y}\pi/\phi_o\mathbf{e}_o$ where ϕ_o is the cosine of the solar zenith angle and $\mathbf{e}_o$ the extra-terrestrial solar irradiance. ρ_{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 infer the corresponding radiances.




Figure 1. Illustration of a single subnetwork in the neural RTM emulator. Each subnetwork predicts the top of atmosphere reflectance $\rho_{obs}(\lambda_i)$ for a single channel centered at wavelength λ_i provided state parameters $\mathbf{x}$ and surface reflectance $\rho_s(\lambda_i)$. Collecting the predictions generated by k subnetworks each modeling distinct channels and converting those predictions from ρ_{obs} to radiance emulates the RTM function $F(\mathbf{x}) \rightarrow \mathbf{y}$ for the selected channels.

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 in state is easier than modeling channels whose spectral responses vary substantially with respect to small changes in state. 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 test 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 without an informed sample stratification yields 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 describe this approach in more detail later in this work.

We can also improve model accuracy and reduce computational demands by exploiting characteristics of the state space in tandem with RTM modeling assumptions. 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 sufficiently low test 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. 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. In practice, requiring the first subnetwork to converge to a lower test error than the subsequent networks can help ensure that the propagated weights will be informative for subsequent channels. We can similarly apply different stopping criteria for subnetworks initialized with weight propagation representing poorly correlated neighboring channels to increase the likelihood that training converges to the appropriate channelwise responses.

Algorithm 1 describes the procedure to train the neural RTM emulator provided n samples from the m -dimensional state space and their corresponding ρ_{obs} outputs each spanning k channels. The output of the algorithm is a trained neural RTM that takes a state vector $\mathbf{x}$ of m parameters as input and outputs a k -dimensional prediction vector. The output radiance vector $\mathbf{y}$ is the concatenated output produced by the trained subnetworks $\{f_i\}_{i=1}^k$ and converted to radiance with respect to ϕ_o and $\mathbf{e}_o$ for each of the k channels.

Our goal is to train each subnetwork f_i to generate accurate predictions for states explicitly included in $\mathbf{X}$ and also (more importantly) for *intermediate* states not explicitly included in $\mathbf{X}$ but within the bounds of the state space. To achieve this, we use a sampling strategy that partitions the state space into training and test sets in a manner that helps optimize each subnetwork to accurately predict intermediate states. We first partition $\mathbf{X}$ and $\mathbf{Y}$ into disjoint training $(\mathbf{X}, \mathbf{Y})^{tr}$ and test $(\mathbf{X}, \mathbf{Y})^{te}$ sets such that $\mathbf{X}^{tr}$ contains all state vectors containing the boundary values $\{\min(p_j), \max(p_j)\}$ of all state parameters. This partitioning ensures the training set contains the convex hull of the Euclidean subspace of $\mathbb{R}^m$ defined by the state parameters, and also

Algorithm 1 Neural RTM Training

Input: $n \times m$ matrix $\mathbf{X}$ of n state vectors, each representing m parameters; $n \times k$ matrix $\mathbf{Y}$ of k -dimensional ρ_{obs} spectra associated with each state vector at wavelengths $\lambda = \{\lambda_i\}_{i=1}^k$; binary mask of training indices $\mathbf{v} = \{v_i\}_i^n$, $v_i \in \{0, 1\}$; number of neural network layers ℓ ; number of hidden neurons in each layer $\{h_l\}_{l=1}^\ell$; convergence tolerance ϵ ; maximum number of training epochs n_{epoch}

Output: Neural RTM model $F(\mathbf{x}) \rightarrow \mathbf{y}$ consisting of k trained neural network regressors f_i each mapping m dimensional state vector $\mathbf{x}$ to corresponding ρ_{obs} response $\mathbf{y}_i$ at wavelength λ_i

for $i = 1$ **to** k **do**

Let $\mathbf{y}_i = \mathbf{Y}_{:,i}$ be the ρ_{obs} responses at wavelength λ_i associated with the n state vectors.

Partition $(\mathbf{X}, \mathbf{y}_i)$ into $\{(\mathbf{X}, \mathbf{y}_i)^{tr}, (\mathbf{X}, \mathbf{y}_i)^{te}\}$ using training indices $\mathbf{v}$.

Let f_i be an L -layer neural network model with set of weight matrices $W_i = \{\mathbf{W}_l\}_{l=1}^L$ and corresponding bias vectors $b_i = \{\mathbf{b}_l\}_{l=1}^L$.

if $i = 1$ **then**

Initialize new model r_1 for first channel by populating W_1 , b_1 with random values (via (Glorot and Bengio, 2010)).

else

Propagate weight matrices and bias vectors from previous model r_{i-1} to current model f_i via $W_i = W_{i-1}$, $b_i = b_{i-1}$.

end if

for $e = 1$ **to** n_{epoch} **do**

Train f_i to minimize ρ_{obs} prediction error for channel centered at wavelength λ_i based on training set $(\mathbf{X}, \mathbf{y}_i)^{tr}$.

Compute average error e_{test} applying f_i to test set $(\mathbf{X}, \mathbf{y}_i)^{te}$.

end for

if e_{test} has converged or $e = n_{epoch}$ **then**

return Trained model f_i

end if

end for

return Trained neural RTM $F(\mathbf{x}) \rightarrow \mathbf{y} = \{f_i(\mathbf{x}) \rightarrow \mathbf{y}_i\}_{i=1}^k$

that all test states in $\mathbf{X}^{te}$ represent intermediate states within the hull. To capture the internal structure of the state space within the hull, the training set should also contain one or more intermediate state vectors for each p_j satisfying $\min(p_j) < \mathbf{x}_j^{tr} < \max(p_j)$. Given the training and test partitions, we train each subnetwork to model f_i by minimizing the L^2 -regularized mean squared error (MSE) between the predicted and the observed values of the n^{tr} training samples $(\mathbf{X}, \mathbf{y}_i)^{tr}$ representing the ρ_{obs}

5 responses for the i^{th} channel.

We use a feed-forward architecture with two hidden layers, rectifying linear activation functions in the hidden layers, which have been shown to be more robust than the conventional sigmoid / tanh activations used in traditional neural networks(Nair and Hinton, 2010), and a linear activation in the output layer for each subnetwork. We use the method proposed by Glorot et al. (Glorot and Bengio, 2010) to initialize subnetwork weights when necessary, and use the widely-used error back propagation

10 algorithm (Werbos, 1982) with adaptive moment estimation (Kingma and Ba, 2014) to optimize the weights via gradient descent. We train each subnetwork until the error converges to within 0.1% mean absolute error (MAE), or we reach the maximum number of epochs n_{epoch} . 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).

3 Neural RTM Emulation for PRISM

We define a case study demonstrating the capabilities of our RTM emulator for atmospheric correction on data acquired by 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.

State Parameter	State Values
Solar azimuth (ϕ_r)	$0, \frac{\pi}{8}, \dots, \frac{\pi}{2}, \dots, \frac{7\pi}{8}, \pi$
Observer zenith angle ($\cos(\theta_v)$)	0.94, 0.95, 0.96, 0.97 , 0.98, 0.99, 1.0
Aerosol optical depth (τ)	0.05, 0.1, 0.2 , 0.3
Water vapor (H_2O)	0, 0.5, 1.0, 1.5 , 2.0, 2.5
Surface reflectance (ρ_s)	0.05, 0.1, 0.25, 0.5, 1.0

Table 1. State parameters values used in LibRadTran model runs to generate ρ_{obs} spectra to train and validate the neural RTM. State vectors containing the median value of each auxiliary parameter (indicated by red text) are held out for testing, while the remaining state vectors are used for training the channelwise subnetworks.

Our state space consists of the surface reflectance ρ_s , represented by a single free parameter per instrument channel; along with $m = 4$ state parameters captured in state vector $\mathbf{x}$ represing a consise but representative suite of parameters used operational settings. These include: 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 $\cos(\theta_v)$; and the relative azimuth angle between the observer and the sun, written ϕ_r . Each of these free parameters 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 identified a set of values for each state parameter that covered the anticipated range of conditions that could occur during the flight campaign, and provided those values in Table 1. We generated RTM outputs using the LibRadTran radiative transfer code (Emde et al., 2016; Mayer and Kylling, 2005) with mid latitude summer atmosphere appropriate for the PRISM flightline we considered in this study. We provide the template LibRadTran config file in supplemental file (prm20151026t173148_libradtran_config) (Kurucz, 1994; Buehler et al., 2009; Bodhaine et al., 1999). Generating ρ_{obs} spectra for every combination of state parameter values yielded a total of $n = 9,072$ $k = 7101$ dimensional ρ_{obs} output spectra




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.

spanning the range of the PRISM instrument wavelengths with 0.1nm spacing. Our test data consists of the set of all state vectors containing the median value of each state parameter (shown in red text in Table 1) and the ρ_{obs} spectra associated with those states. The remaining states and their corresponding ρ_{obs} spectra form our training set.

Figure 2 depicts the changes in the ρ_{obs} spectra with respect to parameters ϕ_r (top left), $\cos(\theta_v)$ (top right), τ (bottom left) and H_2O (bottom right), while holding the other parameters fixed at their median values. Unsurprisingly, the most visibly dramatic changes occur as absorption features appear with increased H_2O vapor concentrations. Of the remaining auxiliary pa-

rameters, 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.

For this case study, we focused on modeling $F(\mathbf{x})$ based on LibRadTran outputs resampled to the PRISM instrument channels. This dramatically reduces the computation required to construct the neural RTM, as we only needed to train a total of 245 subnetworks representing each of the PRISM channels with 2.83nm spacing, rather than the 7101 channels at 0.1nm spacing generated by LibRadTran. As a consequence of convolving the LibRadTran spectra to the lower resolution PRISM Spectral Response Function (SRF) at each wavelength, the ρ_{obs} are no longer strictly monochromatic, but the instrument channels are well-separated so that channelwise coupling should not be a significant issue. We plan to construct a more general neural RTM that generates ρ_{obs} predictions at 0.1nm spacing, which are then convolved to the spectral response function associated with a particular sensor.

We observed experimentally that subnetworks consisting of at two hidden layers with 50 units each and a training cycle of at most 500 epochs (where one epoch consists of a full pass of gradient updates over the training set) with batch sizes ranging from 100 to 200 training samples was sufficient for each subnetwork to converge to our error requirements for the state space parametrized by values in Table 1. Notably, single layer networks were often insufficient to fit difficult to model channels (such as the water absorption bands) whose ρ_{obs} responses were unstable or highly nonlinear with respect to changes in the state parameters. We set the initial learning rate to 0.001 with the following adaptive moment estimation parameters $\{\beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 10^{-10}\}$ and set the L^2 regularization penalty term α to 10^{-4} for each subnetwork. A longer training cycle or additional hidden units can be used to match the RTM output more precisely, and would likely be necessary to model more complex state parameter spaces.




Figure 3. Top: ρ_{obs} test prediction error per channel using channelwise linear regressors (black line). Bottom: neural network test prediction error per channel using subnetworks trained from scratch (NN , red line) versus subnetworks trained weight propagation (NN_{WP} , blue line).

As a baseline comparison, we used the channelwise training samples $((\mathbf{X}, \mathbf{y}_i)^{tr}$ in Algorithm 1) to train a least squares linear regressor on the ρ_{obs} responses for each channel, and applied each regressor to generate predictions on the associated test samples $((\mathbf{X}, \mathbf{y}_i)^{te}$ in Algorithm 1)). The channelwise test errors using the linear regressors provide an approximate upper bound of the error that would be incurred using piecewise / locally linear interpolation to infer ρ_{obs} responses for intermediate states based on lookup tables. Figure 3 compares the ρ_{obs} test prediction error using the channelwise linear regressors (top figure, black line) to the error produced by the channelwise subnetworks trained from scratch (NN , blue line) versus channelwise subnetworks initialized with weight propagation (NN_{WP} , red line). The channelwise subnetworks yield an order of magnitude reduction in prediction error on all channels in comparison to the linear regressors, and demonstrates potentially significant issues with lookup table-based 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 the ρ_{obs} responses vary rapidly for adjacent channels.

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 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.




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.

Investigating further, we measured the average root mean square error (RMSE) on the test set with respect to the pairwise interactions between ρ_s and the four state parameters, and show the resulting error surfaces in Figure 5. Relatively small errors

for the majority of the parameter space indicate that the ρ_{obs} spectra vary smoothly with respect to most state parameter values, with the most significant variability emerging from a small range of values in the $\rho_s \in [0.4, 0.8]$ and $H_2O \in [1.0, 2.0]$ regions of the state space. 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

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 (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.

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 surface reflectance back into radiance space where it becomes a multiplicative correction factor applied independently to each channel. We generate a “standardized” surface reflectance target by performing a first-principles retrieval for a beach sand radiance spectrum manually selected from the target PRISM image. We smooth the resulting surface reflectance spectrum to suppress significant atmospheric features, and use the smoothed spectrum to generate radiometric correction factors appropriate to our flightline. Applying the resulting factors fine tunes the calibration for optimal results, and suppresses residual errors caused by uncertainty in spectral response or RTM inaccuracy. For reference, the beach sand radiance spectrum and the resulting smoothed surface reflectance spectrum are shown in Figure 6.

We applied the atmospheric correction procedure to a set of radiance spectra from the PRISM flightline representing a 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 extremely smooth and faithful estimate of a dark vegetation spectrum. Figure 8 shows additional radiance spectra (top), and their corresponding surface reflectance retrievals (bottom). The high quality surface reflectance estimates — evidenced by the lack of residual bumps caused by atmospheric absorption and the flat, low surface 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 PRISM spectrum (about 0.02 milliseconds per channel). This represents a several orders of magnitude reduction in runtime in comparison to analogous first-principles RTMs (i.e., monochromatic RTMs that solve the coupled scattering-absorption problem in a computationally exact manner, such as DISORT), which typically required over 10 minutes to generate a spectrum at 0.1nm spacing (about 0.15 seconds per channel).

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 provided by fast and accurate RTM emulators are particularly useful for 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 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 more “universal” neural RTM designed to generate ρ_{obs} predictions at much spectral resolution for a comprehensive set of states, and are also investigating Bayesian optimization or smart sampling approaches (e.g., Loyola R et al. (2016)) that may provide an informed sampling the state 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.

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.

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.
- Bodhaine, B. A., Wood, N. B., Dutton, E. G., and Slusser, J. R.: On Rayleigh optical depth calculations, *Journal of Atmospheric and Oceanic Technology*, 16, 1854–1861, 1999.
- 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.
- Buehler, S. A., John, V. O., Kottayil, A., Milz, M., and Eriksson, P.: ARTICLE IN PRESS, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 111, 602–615, <https://doi.org/10.1016/j.jqsrt.2009.10.018>, 2009.
- 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.
- 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.
- 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.
- 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.
- 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, 2015.
- Hochberg, E. J.: Remote sensing of coral reef processes, *Coral Reefs: An Ecosystem in Transition*, pp. 25–35, 2011.
- 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.
- 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.
- 5 Kingma, D. P. and Ba, J.: Adam: A method for stochastic optimization, arXiv preprint arXiv:1412.6980, 2014.
- Kox, S., Bugliaro, L., and Ostler, A.: Retrieval of cirrus cloud optical thickness and top altitude from geostationary remote sensing, *Atmospheric Measurement Techniques*, 7, 3233–3246, <https://doi.org/10.1175/1520-0442-12.1.170>, 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.
- 10 Kurucz, R. L.: Synthetic infrared spectra, in: *Infrared solar physics*, pp. 523–531, Springer, 1994.
- Loyola, D. G., Gimeno García, S., Lutz, R., Argyrouli, A., Romahn, F., Spurr, R. J. D., Pedernana, M., Doicu, A., Molina García, V., and Schüssler, O.: The operational cloud retrieval algorithms from TROPOMI on board Sentinel-5 Precursor, *Atmospheric Measurement Techniques*, 11, 409–427, <https://doi.org/10.5194/acp-8-6565-2008>, 2018.
- Loyola R, D. G., Pedernana, M., and García, S. G.: Smart sampling and incremental function learning for very large high dimensional data, *Neural Networks*, 78, 75–87, <https://doi.org/10.1016/j.neunet.2015.09.001>, 2016.
- 15 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.
- 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.
- 20 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, V., and Yee, K.: Portable Remote Imaging Spectrometer coastal ocean sensor: design, characteristics, and first flight results, *Appl. Opt.*, 53, 1363–1380, 2014.
- 25 Nair, V. and Hinton, G. E.: Rectified linear units improve restricted boltzmann machines, *International Conference on Machine Learning*, 2010.
- 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.
- 30 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.
- 35 Schaepman, 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.
- Schaepman-Strub, G., Schaepman, 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.
- 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.
- 5 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.
- 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 HypIRI preparatory campaign, *Remote Sensing of Environment*, 167, 64 – 77, 2015.
- 10 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.
- 15 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>,
- 20 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.
- 25 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.




Figure 5. Pairwise contour plots showing the ρ_{obs} test prediction error (RMSE) surfaces with respect to the state parameter values specified in Table 1. Contour labels on the off-diagonal subplots give the error levels associated with each contour. Diagonal subplots show the average RMSE in ten uniformly spaced bins spanning the (x-axis) range of each parameter. Vertical labels on the diagonal subplots indicate the minimum and maximum error values for each parameter and their corresponding bins.




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




Figure 7. Example surface 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 surface reflectance spectrum (black).




Figure 8. Selected radiance spectra (top) and corresponding surface 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).