

Evaluating the accuracy of Gaussian approximations in VSWIR imaging spectroscopy retrievals

Kelvin M. Leung, David R. Thompson, *Senior Member, IEEE*, Jouni Susiluoto, Jayanth Jagalur, Amy Braverman, Youssef Marzouk

Abstract—The joint retrieval of surface reflectances and atmospheric parameters in VSWIR imaging spectroscopy is a computationally challenging high-dimensional problem. Using NASA’s Surface Biology and Geology mission as the motivational context, the uncertainty associated with the retrievals is crucial for further application of the retrieved results for environmental applications. Although Markov chain Monte Carlo (MCMC) is a Bayesian method ideal for uncertainty quantification, the full-dimensional implementation of MCMC for the retrieval is computationally intractable.

In this work, we developed a block Metropolis MCMC algorithm for the high-dimensional VSWIR surface reflectance retrieval that leverages conditional linearity structure in the forward radiative transfer model to enable tractable fully Bayesian computation. We use the posterior distribution from this MCMC algorithm to assess the limitations of optimal estimation, the state-of-the-art Bayesian algorithm in operational retrievals which is more computationally efficient but uses a Gaussian approximation to characterize the posterior. Analyzing the differences in the posterior computed by each method, the MCMC algorithm was shown to give more physically sensible results and reveals the non-Gaussian structure of the posterior, specifically in the atmospheric aerosol optical depth parameter and the low-wavelength surface reflectances.

Index Terms—Remote sensing, imaging spectroscopy, Markov chain Monte Carlo, Bayesian computation.

I. INTRODUCTION

WHETHER airborne or orbital, all remote sensing missions face a common challenge of characterizing distant objects using only measurements made at the sensor. In the Earth sciences, investigators will often solve this problem with physics-based models that use the state of the surface or atmosphere to predict the remote measurement. Investigators can then *retrieve*, or determine, the state most consistent with the remote data. Model inversion methods are used for diverse sensors ranging from infrared or microwave sounders, to multiangle imagers, to radiometers. Perhaps one of the most challenging applications from a computational perspective is remote Visible/Short-Wave Infrared (VSWIR) imaging spectroscopy [1].

VSWIR imaging spectrometers acquire a data cube with two spatial dimensions and one spectral dimension. In other words, they produce images in which each pixel contains a radiance spectrum covering the entire solar reflected interval from 380 to 2500 nm. This interval is sensitive to diverse surface and atmospheric processes, making these sensors useful for a wide range of studies from terrestrial and aquatic

ecology, to geology, to hydrology and cryosphere studies. These Earth surface studies aim to measure properties of the surface that create characteristic features in the reflectance spectra. Roughly speaking, surface reflectance is the fraction of incident illumination at the surface which is reflected back in the direction of the sensor. However, imaging spectrometers observe radiance at the top of the atmosphere, so inference to remove the effects of the atmosphere is required to estimate surface reflectance at each pixel. The reflectance can then be used to further estimate properties of the Earth surface. Because of the high data volume of these sensors and their broad spectral range encompassing a wide range of physical phenomena, VSWIR imaging spectrometers present a particularly challenging test case for efficient inference algorithms.

Our motivating context for this problem is NASA’s Surface Biology and Geology mission (SBG) [2], [3]. The objective of SBG is to track changes in surface properties pertaining to ecosystems, coastal zones, agriculture, and snow and ice accumulations over time, for the entire planet, by first retrieving the surface reflectances. For these types of scientific applications, the uncertainty associated with the retrieval is particularly important. This motivates the need for a Bayesian method to determine the posterior distribution of the surface reflectances and related atmospheric parameters.

Markov chain Monte Carlo (MCMC) is a Bayesian sampling method that was introduced in the context of remote sensing retrievals in [4]–[6]. Recent retrieval problems are generally high-dimensional and require methods of dimension reduction to lower the computational complexity. For example, [7] breaks up the high-dimensional parameter space into low-dimensional blocks that can be sampled in parallel. [6] and [8] implement MCMC in a low-dimensional parameter subspace obtained from principal component analysis (PCA), which defines the subspace based on the leading directions of an eigenvalue problem.

More recent methods of dimension reduction such as likelihood informed subspace (LIS) have also been considered in the retrieval context, where MCMC is performed in a specific low-dimensional subspace, defined by a different eigenvalue problem, that is determined by the data [9]. [10] uses LIS in atmospheric methane retrievals, and [11] uses LIS in retrievals for atmospheric concentrations of carbon dioxide in NASA’s Orbiting Carbon Observatory-2 (OCO-2) mission. Contrary to most retrievals where the number of parameters is under 100, there are over 300 parameters in the SBG retrieval problem, which makes the problem much more difficult in terms of computational tractability. Furthermore, since dimension reduction methods such as PCA or LIS are low-dimensional

K. Leung, J. Jagalur, and Y. Marzouk are with the Massachusetts Institute of Technology. D. R. Thompson, J. Susiluoto, and A. Braverman are with the Jet Propulsion Laboratory, California Institute of Technology.

approximations, significantly reducing the dimension leads to problems in convergence to the posterior distribution.

Optimal estimation (OE) [12] is the current state-of-the-art algorithm for an operational setting, such as for NASA's EMIT mission [13], [14]. OE is a Bayesian retrieval algorithm that computes a maximum a posteriori (MAP) estimate of the parameters and characterizes the posterior distribution using the Laplace approximation. Although this leads to fast Gaussian approximations of the posterior, the posterior is not Gaussian in the general case.

Our objective is to explore how well this approximation holds. There are two main contributions of this work.

- 1) We developed a computationally tractable block Metropolis MCMC algorithm for the VSWIR retrieval problem. This fully Bayesian algorithm allows for the characterization of a non-Gaussian posterior and performs exact inference in the limit of infinite samples.
- 2) We use this algorithm to evaluate the limitations of the OE method and identify the scenarios in which it is sufficiently accurate.

In contrast to the PCA and LIS methods mentioned above, our method that has two conceptual advantages over dimension reduction. The block Metropolis algorithm that we develop is asymptotically exact and does not rely on truncation of the eigenspectrum. In addition, PCA and LIS are methods used to try and identify underlying problem structure, whereas we directly enforce structure in the algorithm. We therefore do not make comparisons with dimension reduction methods in our results.

We emphasize that we are comparing the fully Bayesian MCMC algorithm to OE, focusing on the differences in the posterior distribution characterized by each method. Absolute performance of the algorithm relative to the ground truth depends on factors outside the scope of this paper. While we do not expect MCMC to be used in an operational setting due to its longer runtime than OE, it can be viewed as a method to validate and identify shortcomings in OE.

II. SETUP OF THE REMOTE SENSING PROBLEM

The remote sensing retrieval considered in this paper is modelled as an inverse problem. For each pixel of the image captured by the imaging spectrometer, the quantities of interest are the surface and atmospheric parameters that are retrieved given the radiance at the same pixel. This type of retrieval can be thought of as a statistical inference problem for one set of multidimensional parameters with one set of multidimensional data.

We use the notation \mathbf{y} to denote the set of radiance observations from the imaging spectrometer. The radiances are used to infer the state \mathbf{x} , which consists of the surface and atmospheric parameters. Incoming solar radiation is reflected off the Earth surface, and the transfer of radiation through the atmosphere is modelled by a vector-valued forward function $f(\cdot)$. The full expression is known and is written in (4). The observations are represented by the output of the forward model with additive noise, $\mathbf{y} = f(\mathbf{x}) + \epsilon$.

The setup of this inference problem and the desire for uncertainty quantification leads to a Bayesian formulation. Given

the prior and likelihood distributions, the posterior distribution of the surface and atmospheric parameters conditioned on the observed radiance, $\pi(\mathbf{x}|\mathbf{y})$, is obtained using Bayes rule:

$$\pi(\mathbf{x}|\mathbf{y}) = \frac{\pi(\mathbf{y}|\mathbf{x})\pi(\mathbf{x})}{\pi(\mathbf{y})} \propto \pi(\mathbf{y}|\mathbf{x})\pi(\mathbf{x}). \quad (1)$$

This section provides an overview of the elements associated with the prior and likelihood distributions, including parameters and data. Formulations to determine the posterior distribution are described in Section III.

A. Surface and atmospheric parameters

The inversion problem estimates the surface and atmospheric parameters, which are concatenated into one state vector:

$$\mathbf{x} = [\mathbf{x}_{\text{refl}}, \mathbf{x}_{\text{atm}}]^\top, \quad (2)$$

where $\mathbf{x}_{\text{refl}} \in \mathbb{R}^n$ and $\mathbf{x}_{\text{atm}} \in \mathbb{R}^2$. There are $n = 432$ surface parameters and two atmospheric parameters.

- The surface parameters \mathbf{x}_{refl} are surface reflectances that describe the proportion of solar radiation that is reflected off the surface at each of the n wavelengths. The wavelengths range from 380 to 2500 nm.
- The two atmospheric parameters are $\mathbf{x}_{\text{atm}} = [x_{\text{AOD}}, x_{\text{H}_2\text{O}}]^\top$, which consist of Aerosol Optical Depth (AOD) at 550 nm and column precipitable water vapour (cm).

Aerosol optical depth (AOD) is a measure of the atmospheric concentration of aerosols. Specifically, it is the proportion of radiation that is absorbed by aerosols at wavelength 550 nm. The second atmospheric parameter is the column precipitable water vapour (cm), which is the volume of water per vertical column of atmosphere. Other gas concentrations that affect light transmittance and scattering are accounted for within the radiative transfer model.

Prior. The prior on the parameters is modelled as a normal distribution. The surface and atmospheric parameters are treated independently, giving a block diagonal structure to the prior covariance:

$$\boldsymbol{\mu}_{\text{pr}} = [\boldsymbol{\mu}_{\text{refl}}^0, \boldsymbol{\mu}_{\text{atm}}^0]^\top, \quad \boldsymbol{\Gamma}_{\text{pr}} = \begin{bmatrix} \boldsymbol{\Gamma}_{\text{refl}}^0 & 0 \\ 0 & \boldsymbol{\Gamma}_{\text{atm}}^0 \end{bmatrix} \quad (3)$$

The surface prior is created using a set of over 1400 historical reflectance spectra from the EcoSIS spectral library [15]. These spectra are fitted to a Gaussian mixture model with eight components, each corresponding to a different type of terrain on the Earth surface that have similar characteristics, such as vegetation or aquatic environments. For a particular inversion, the component with shortest Mahalanobis distance to the initial state estimate is used as the Gaussian surface prior.

For the atmospheric parameters, the priors are chosen to have large variances to allow relatively unconstrained exploration of the posterior.

B. Forward model

The forward model, $f(\cdot) = [f_1(\cdot), \dots, f_n(\cdot)]^\top$, approximates the propagation of photons through the Earth atmosphere from the surface to the imaging spectrometer. The model used in this work consists of two stages to map the state \mathbf{x} to the radiance \mathbf{y} .

The first step is the computation of intermediate parameters using MODTRAN 6.0, a high-fidelity radiative transfer model [16]. The outputs of MODTRAN given \mathbf{x}_{atm} are three n -dimensional parameters that describe light propagation through the atmosphere: path reflectance ρ_a , spherical albedo s , and atmospheric transmission t . For computational efficiency, a lookup table is generated for a set of reference atmospheric conditions.

The second step is to use the intermediate parameters to calculate the radiance. After linearly interpolating these parameters in the lookup table, the mechanics of the forward model for $i = 1, \dots, n$ are given by:

$$f_i(\mathbf{x}) = \frac{\phi_0}{\pi} e_{0,i} \left[\rho_{a,i}(\mathbf{x}_{\text{atm}}) + \frac{t_i(\mathbf{x}_{\text{atm}}) \cdot x_{\text{refl},i}}{1 - s_i(\mathbf{x}_{\text{atm}}) \cdot x_{\text{refl},i}} \right], \quad (4)$$

where ϕ_0 is the cosine of the solar zenith angle and e_0 is the solar downward irradiance at the top of the atmosphere [17].

C. Noise model

We model the observation uncertainty matrix by combining covariance matrices from different independent error sources as in [18]. First, we determine the intrinsic sensor noise, which includes uncertainty due to discrete photon counts, electronic uncertainty in the analog-to-digital conversion, and thermal noise from the instrument itself. These sources are all independent and uncorrelated in each channel leading to a diagonal covariance structure. The photon noise contribution depends on the magnitude of the radiance itself, so we use the measurement itself to predict its own noise level. Any error induced by this circularity is acceptable since noise is small relative to the total magnitude, with a signal-to-noise ratio of 500 or 1000 for typical spectra.

In addition to the instrument noise, we also model several systematic error sources following [12]. We use a diagonal matrix to represent a 1% uncertainty in calibration. Another diagonal matrix represents systematic radiative transfer model errors due to spectral calibration uncertainty and the intrinsic uncertainty in the unretrieved components of the atmospheric model. The covariance matrices from these independent error sources combine additively to form a final observation error matrix Γ_{obs} . For this work, we set Γ_{obs} to be a diagonal matrix with a signal-to-noise ratio of 50 to account for these error sources.

D. Retrieval

The prior, forward, and noise models are used to solve the inversion problem of retrieving the surface and atmospheric parameters given the observations. From a Bayesian perspec-

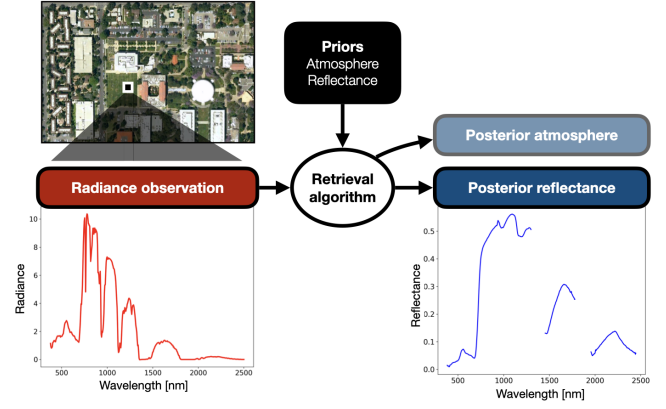


Fig. 1. Retrieval flowchart for an example observation over a grassy field.

tive, the retrieval is equivalent to characterizing the posterior that is given by:

$$\pi(\mathbf{x}|\mathbf{y}) \propto \exp \left(-\frac{1}{2} \|\mathbf{x} - \boldsymbol{\mu}_{\text{pr}}\|_{\Gamma_{\text{pr}}}^2 - \frac{1}{2} \|\mathbf{y} - f(\mathbf{x})\|_{\Gamma_{\text{obs}}}^2 \right) \quad (5)$$

Methods for inference are described in Section III. An example of the retrieval process from a pixel in the satellite image to the surface reflectance is shown in Figure 1. Possible retrieval algorithms are discussed in the next section. The gaps in the reflectance represent the wavelengths for which most of the radiation is absorbed by water droplets in the atmosphere, rendering the retrieval meaningless in those regions. Ignoring these wavelengths, there are 324 reflectance parameters of interest.

III. ESTIMATION AND INFERENCE FORMULATIONS

The existing state-of-the-art Bayesian method used in VSWIR remote sensing problems is optimal estimation (OE) [12], which is a computationally efficient way to obtain estimates of the surface and atmospheric parameters. In this work we place the emphasis on the uncertainty quantification of the retrieval of these parameters. We consider a sampling-based Markov chain Monte Carlo (MCMC) approach to characterize the full posterior distribution.

A. Optimal estimation

Given a Gaussian prior and the forward and noise models, the OE method solves an optimization problem to estimate the surface and atmospheric parameters. The negative log posterior is used as the objective function, and the resulting parameter estimate is denoted as $\mathbf{x}_{\text{MAP}} = \arg \max_{\mathbf{x}} c(\mathbf{x})$, where:

$$c(\mathbf{x}) = \frac{1}{2} (\mathbf{x} - \boldsymbol{\mu}_{\text{pr}})^\top \Gamma_{\text{pr}}^{-1} (\mathbf{x} - \boldsymbol{\mu}_{\text{pr}}) + \frac{1}{2} (\mathbf{y} - f(\mathbf{x}))^\top \Gamma_{\text{obs}}^{-1} (\mathbf{y} - f(\mathbf{x})). \quad (6)$$

The covariance estimate is a Laplace approximation [19] derived from linear Bayesian inversion theory [20] using the local linearization of the forward model at the MAP estimate [21].

$$\Gamma_{\text{L}} = \left(\nabla f(\mathbf{x}_{\text{MAP}})^\top \Gamma_{\text{obs}}^{-1} \nabla f(\mathbf{x}_{\text{MAP}}) + \Gamma_{\text{pr}}^{-1} \right)^{-1} \quad (7)$$

Since the OE posterior is characterized using a local Gaussian approximation, we sometimes refer to OE as the *approximate Bayesian* method to contrast with the *fully Bayesian* MCMC approach.

B. Fully Bayesian approach

Although the Laplace approximation in optimal estimation would be accurate if the posterior is approximately Gaussian, this is not the case in general. Since the forward model is nonlinear, the posterior shape cannot be determined a priori, making it impossible to determine whether a normal approximation is sufficiently accurate. Therefore, a method of characterizing the full posterior distribution is needed to obtain an accurate measure of uncertainty associated with the retrieval.

Markov chain Monte Carlo (MCMC) [22] is a probabilistic sampling method that addresses the issue of characterizing the posterior, but is computationally intractable in the high-dimensional VSWIR retrieval problem. Methods for dimension reduction such as [9] [23] [24] designed for Bayesian inverse problems can limit the sampling to a low-dimensional subspace. However, these methods are considered *approximate* inference because they involve truncation of information that are deemed less important based on the eigenvalues. This paper presents a technique for *exact* inference specifically for the VSWIR retrieval that takes advantage of the conditional linearity structure of the surface and atmospheric parameters in the state vector. The next subsection outlines the structure that we then use to develop the sampling methodology in Section IV.

C. Linear approximations to the forward model

The structure in the forward model arises from \mathbf{x}_{refl} and \mathbf{x}_{atm} being treated as independent in (4). When the atmospheric parameters are held fixed, the forward model is approximately linear. That is, for \mathbf{x}_{atm} held constant, we can define a submodel conditioned on the atmospheric parameters:

$$f_{\text{refl}}(\mathbf{x}_{\text{refl}}) \approx \mathbf{A}\mathbf{x}_{\text{refl}} + \mathbf{b}, \quad (8)$$

where \mathbf{A} , a diagonal matrix with $A_{ii} = \frac{\phi_0}{\pi} e_{0,i} t_i$, and $b_i = \frac{\phi_0}{\pi} e_{0,i} \rho_{a,i}$, $i = 1 \dots n$, are the deterministic constants computed from the intermediate parameters computed using MODTRAN. The denominator of the second term in (4) is approximately equal to one. The approximately linear structure in the surface parameters can be exploited to accelerate the sampling process.

IV. SAMPLING METHODOLOGY

A computationally tractable fully Bayesian algorithm was developed to obtain samples from the posterior distribution of the surface and atmospheric parameters. The algorithm, based on a block Metropolis MCMC algorithm [22] [25], generates alternating samples of the reflectance and atmospheric parameter blocks. Contrary to algorithms involving dimension reduction, this algorithm performs exact inference, meaning that it converges to the true posterior distribution in

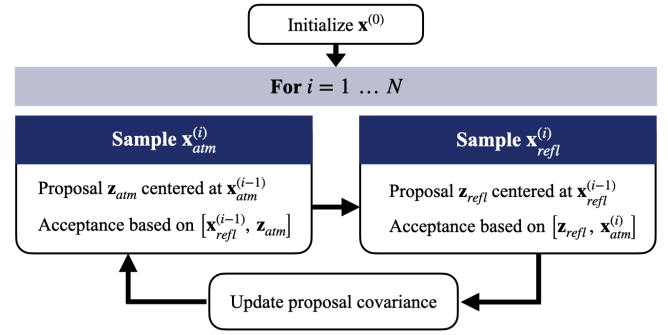


Fig. 2. Flowchart for the block Metropolis MCMC retrieval algorithm.

the limit of infinite samples. The algorithm is described in this section, including the overall structure and the parameter tuning process.

A. Exploiting structure in the forward model

The forward model is known to be approximately linear in the reflectances conditioned on fixed atmospheric parameters. Since the objective is to develop a fully Bayesian algorithm, the linear model described in Section III-C is not used explicitly in the inversion. It is instead used to provide structure to the sampling algorithm.

The motivation behind the block Metropolis algorithm is to restrict the “difficult” sampling to the atmospheric parameters. After obtaining a sample from the atmospheric block, the sampling within the surface block can converge much faster thanks to the approximate conditional linearity given a fixed atmosphere. Without this structure, the algorithm would have to blindly explore the n -parameter space, which is computationally infeasible in practice.

The sampling procedure is depicted in Figure 2. Different initializations for the Markov chain were attempted and resulted in the same stationary distribution with varying burn-in times. Initializing at the MAP estimate obtained using optimal estimation led to sufficiently low burn-in for all cases and was therefore used throughout the experiments. Each subsequent sample is split into the atmospheric and reflectance blocks, each with a proposal and acceptance step. The proposal is a sample from the normal distribution centered at the previous sample with some proposal covariance. This proposal covariance is discussed in the next subsection. The proposal is then accepted or rejected with some acceptance probability computed based on the posterior density. The new samples from both blocks are then concatenated and added to the chain. The acceptance step ensures asymptotic convergence of the chain. The full algorithm is outlined in Section IV-C.

B. Choice of proposal covariance

Different approaches were taken for the two blocks since the structure is known for the reflectances conditioned on the atmospheric parameters but not vice versa. For the atmospheric block, the proposal covariance follows the update procedure of the Adaptive Metropolis algorithm [26], in which the proposal attempts to adapt to the shape of the posterior based on the

previous samples in attempt to explore the parameter space more efficiently. For samples $i \leq 1000$, the proposal covariance is a diagonal matrix with elements ϵ_0 . In subsequent samples $i > 1000$, the proposal is the sum of the covariance of all previous samples scaled with parameter γ and an additive term to prevent singularity. This adaptive scheme is given as:

$$\Gamma_{\text{atm}}^{(i)} = \begin{cases} \epsilon_0 \mathbf{I}_2 & i \leq 1000 \\ \gamma \text{cov}(\mathbf{x}_{\text{atm}}^{(0)}, \dots, \mathbf{x}_{\text{atm}}^{(i-1)}) + \gamma \epsilon_{\text{AM}} \mathbf{I}_2 & i > 1000 \end{cases}, \quad (9)$$

where \mathbf{I}_2 is the 2×2 identity matrix, $\gamma = \frac{2.38^2}{2}$ and $\epsilon_0 = \epsilon_{\text{AM}} = 10^{-3}$.

Two methods of obtaining the proposal covariance for the reflectance block were compared. Both involve computing some approximation of the posterior covariance of the reflectance parameters.

1) **Linear inversion theory.** Modelling the forward sub-model as linear, i.e. making (8) an equality, closed form expressions of the posterior covariance can be derived from linear Bayesian inversion theory [20]. For the linear model in (8) and using the same prior and noise model, the posterior covariance of the reflectances can be expressed as:

$$\Gamma_{\text{lin}}^{(i)} = (\mathbf{I} - \Gamma_{\text{pr}} \mathbf{A}^{(i)T} \Gamma_y^{-1(i)} \mathbf{A}^{(i)}) \Gamma_{\text{pr}}, \quad (10)$$

where $\Gamma_y^{(i)} = \mathbf{A}^{(i)} \Gamma_{\text{pr}} \mathbf{A}^{(i)T} + \Gamma_{\text{obs}}$ is the marginal covariance of the data. A scaled version of this posterior approximation, $\epsilon_1 \Gamma_{\text{lin}}$, where $\epsilon_1 < 1$ is used as the proposal covariance.

2) **Laplace approximation.** Another method is to directly use the Laplace approximation, Γ_L , obtained from OE [19] [21]. This can be done as a preprocessing step to avoid having to compute an inversion problem for every iteration of the chain. The proposal covariance is the scaled Laplace approximation, $\epsilon_2 \Gamma_L$, for some $\epsilon_2 < 1$.

The block Metropolis algorithm was implemented to compare the quality of mixing in the chain for each method of obtaining the proposal covariance. Two million samples were generated for each chain, which were then thinned by taking every 10 samples for a total of 2×10^5 samples. The mixing characteristics were compared by analyzing trace plots and the effective sample sizes. Both scaling parameters ϵ_1 and ϵ_2 were tuned to achieve a near-optimal acceptance rate of approximately 23% [27]. The subsequent results use the scaling parameters $\epsilon_1 = 0.14$ and $\epsilon_2 = 0.11$.

The samples from the reflectance block affect the acceptance in the atmospheric block due to the alternating nature of the sampling process, so we must evaluate the effect of this proposal covariance on the parameters in both blocks. The trace plots for the atmospheric parameters are shown in Figure 3. While the H_2O parameter trace plots are similar across the two methods, the AOD plots show that the chain explores the parameter space more efficiently when using the Laplace approximation. In the proposal from linear inversion theory, the chain requires more samples to move toward a different region of the space.

Since MCMC leads to dependent samples, another useful metric is the autocorrelation time τ_a , which is the time it

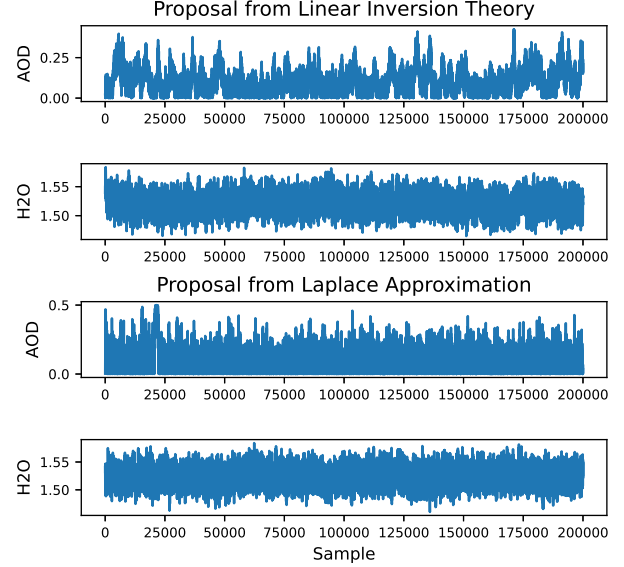


Fig. 3. Trace plots of the atmospheric parameters using the two methods of obtaining proposal covariance for the reflectance block.

takes for the samples to become effectively independent. The effective sample size (ESS) can then be defined as:

$$\text{ESS} = \frac{N}{\tau_a}, \quad (11)$$

where N was taken as 1.8×10^5 after removing the first 2×10^4 samples in the chain as burn-in.

The ESS was computed for each of the $n + 2$ parameters and summarized in Table I. The median value is around 1000, indicating that one effectively independent sample is generated every 180 samples. The proposal covariance obtained from the Laplace approximation generates greater sample sizes throughout, suggesting better mixing and greater sampling efficiency. This method performs particularly better for the atmospheric AOD parameter, where the ESS is roughly a factor of 4 greater than the linear inversion method. For the reflectances, it induces a much better performance for wavelengths less than 1000 nm, as shown in Figure 4. While the ESS from the Laplace method remain fairly consistent throughout all wavelengths, the ESS in low wavelength regions drop significantly for the linear inversion method.

Another advantage of the Laplace method is that the Laplace approximation is constant and only needs to be computed once as a preprocessing step. The linear inversion method requires the computation of (10) between the two blocks in the algorithm for each sample. Although this does not make a noticeable difference in computational time, it simplifies the algorithm. In the final implementation of the algorithm, the proposal covariance for the reflectance block is equal to $\epsilon_2 \Gamma_L$, where the scaling parameter was tuned to be $\epsilon_2 = 0.11$.

C. Block Metropolis Algorithm

Algorithm 1 outlines the final algorithm, which takes advantage of the forward model structure in which $f_{\text{refl}}(\cdot)$ is approximately linear. By using this property along with the scaled Laplace approximation as the proposal covariance, we

TABLE I
EFFECTIVE SAMPLE SIZES FOR MCMC ON BUILDING 177

	Proposal from linear inversion	Proposal from Laplace
Refl Min	108	120
Refl Med	1278	1375
Refl Max	3294	4399
AOD	166	633
H ₂ O	527	784

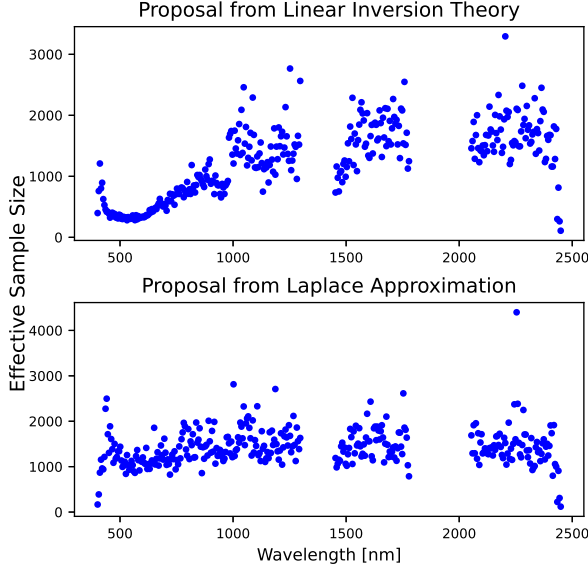


Fig. 4. Effective sample sizes of the reflectance chains.

are able to obtain samples that efficiently explore the parameter space of both the surface and atmospheric parameters. Note that in the atmospheric block, \mathbf{z}_{atm} is drawn from a truncated normal distribution with a lower bound of zero.

Algorithm 1 Block Metropolis

```

1: Initialize  $\mathbf{x}^{(0)} = \mathbf{x}_{\text{MAP}}$ 
2: for  $i = 1 \dots N$  do
3:   Sample  $\mathbf{x}_{\text{atm}}^{(i)}$ 
4:   Proposal  $\mathbf{z}_{\text{atm}} \sim \mathcal{N}(\mathbf{x}_{\text{atm}}^{(i-1)}, \Gamma_{\text{atm}}^{(i)})$  such that  $z_{\text{atm}} \geq 0$ 
5:   Metropolis accept/reject for  $[\mathbf{x}_{\text{refl}}^{(i-1)}, \mathbf{z}_{\text{atm}}]$ 
6:   Sample  $\mathbf{x}_{\text{refl}}^{(i)}$ 
7:   Proposal  $\mathbf{z}_{\text{refl}} \sim \mathcal{N}(\mathbf{x}_{\text{refl}}^{(i-1)}, \epsilon_2 \Gamma_L)$ 
8:   Metropolis accept/reject for  $[\mathbf{z}_{\text{refl}}, \mathbf{x}_{\text{atm}}^{(i)}]$ 
9:   Compute  $\Gamma_{\text{atm}}^{(i+1)}$ 
10: end for

```

V. RESULTS

The results focus on comparing the posterior distribution characterized by the fully Bayesian MCMC method with the posterior approximated by optimal estimation. The MCMC algorithm was executed for four radiance datasets. The datasets, collected over the JPL campus using the airborne AVIRIS-NG instrument [28] [29], attempt to include a variety of terrain types and are named Building 177, Building 306, Mars Yard, and Parking Lot. The radiance spectra are shown in Figure 5.

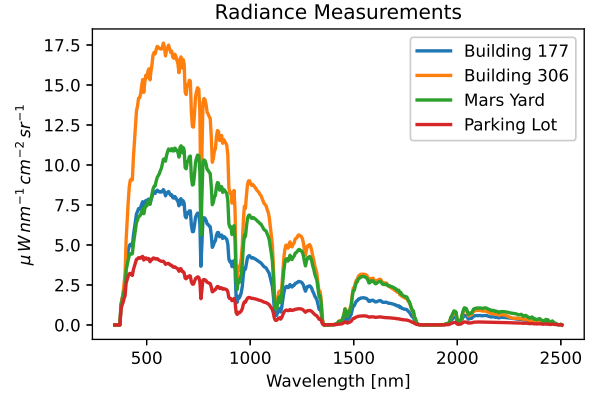


Fig. 5. Radiance measurements for test cases collected over JPL campus.

We reiterate that the scope of this paper is comparing the posterior distributions characterized by MCMC and optimal estimation. Comparisons with the ground truth are omitted since the absolute retrieval accuracy depends on factors outside this scope such as the accuracy of the physics-based radiative transfer model. In this section, we compare the surface posterior obtained from both methods using several metrics, followed by the posterior on the atmospheric parameters. Then, we evaluate the Gaussianity of the full posterior distribution.

In our implementation of the block Metropolis MCMC algorithm, two million samples were obtained for each chain, with the first 2×10^5 discarded as burn-in since the chain stabilizes by then for most cases, as seen in the trace plots in Figure 3. The chain was thinned by taking every tenth sample to reduce storage. The overall acceptance rate for all cases ranged from 0.2 to 0.3.

A. Surface posterior comparison

We first compare the mean and covariance of the posterior distribution. Figure 6 plots the posterior mean reflectance for the four test cases and the MAP estimate from optimal estimation. Figure 7 plots the relative difference of the two methods normalized over the values from the MCMC method. In the first three cases, the greatest differences occur in the low wavelength regions and all peak around 0.02. The Parking Lot case begins with a high relative difference at low wavelength but generally remain below 0.02. The relative difference between the posterior mean and MAP estimates are below 6% throughout the spectrum, which gives confidence in both methods that they converge to a sensible location in the parameter space. Parking Lot generally follows different patterns from the other cases in our analysis since it is difficult to retrieve due to low signal levels and heightened light scattering from neighbouring buildings [12].

Figure 8 shows the marginal variance of the posterior distribution in the reflectances obtained using Laplace approximation in optimal estimation and MCMC. For the first three cases, the Laplace approximation predicts a higher marginal variance than the MCMC posterior variance, particularly in the low wavelength regions below 1000 nm. For the Parking Lot case, the posterior marginal variance is slightly higher than

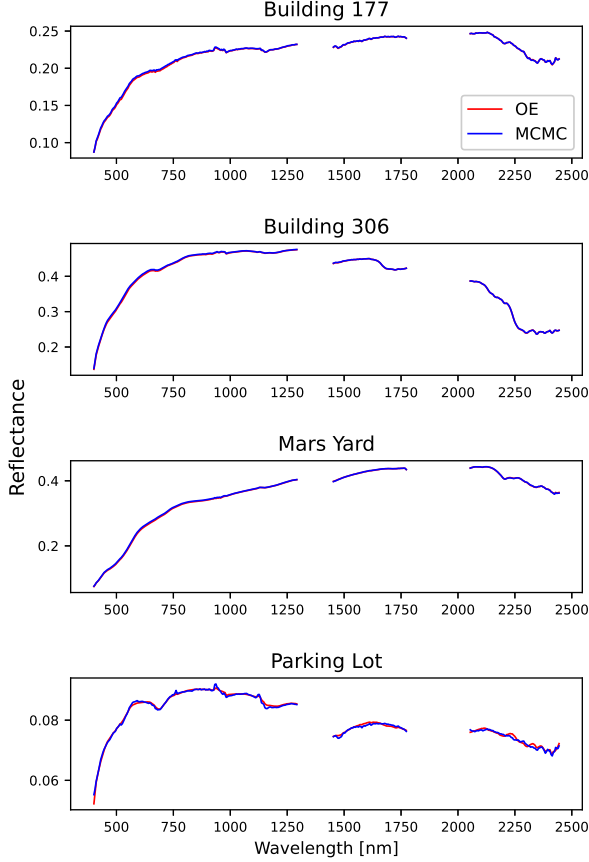


Fig. 6. Posterior mean and MAP estimates for reflectances.

the Laplace approximation except for the regions around 380 nm and 550 nm.

In addition to the marginal variances, it is also necessary to look into the differences in cross-correlations. We compare the MCMC and OE covariances (Γ_M and Γ_L) using three metrics involving trace, Frobenius norm, and Förstner distance, which are defined below. The first metric is the relative difference in trace normalized by the trace of the MCMC posterior covariance:

$$d_{tr} = \left| \frac{\text{tr}(\Gamma_M) - \text{tr}(\Gamma_L)}{\text{tr}(\Gamma_M)} \right|. \quad (12)$$

The second metric is the relative difference in Frobenius norm normalized by the Frobenius norm of the MCMC posterior covariance:

$$d_{norm} = \frac{\|\Gamma_M - \Gamma_L\|_F}{\|\Gamma_M\|_F}. \quad (13)$$

The third metric involves the Förstner distance, which is a metric that measures the distance between two symmetric positive definite matrices [30]. The Förstner distance between two SPD matrices Γ_A and Γ_B is defined by:

$$d_f = \sqrt{\sum_i \ln^2(\sigma_i)}, \quad (14)$$

where σ_i are the generalized eigenvalues of the eigenpencil (Γ_A, Γ_B) . The relative difference defined using this metric is

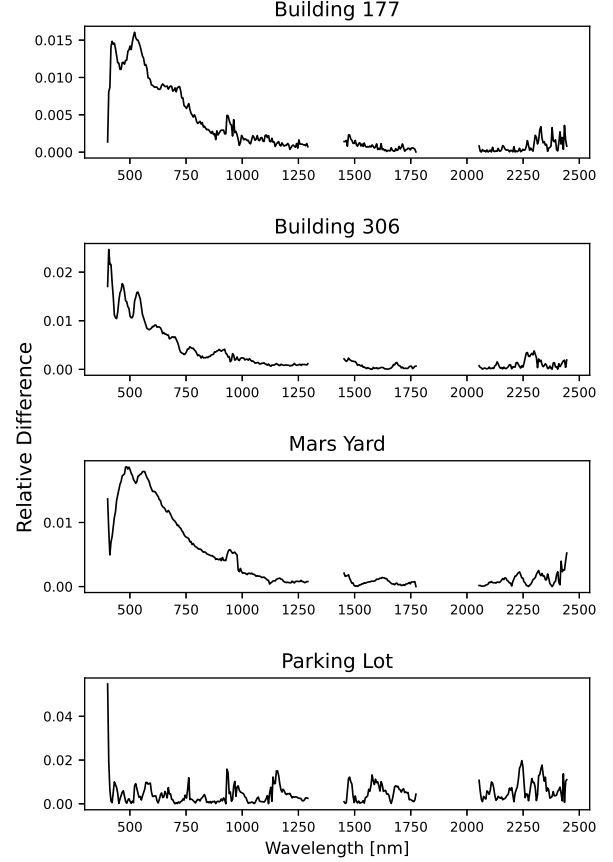


Fig. 7. Relative difference between posterior mean and MAP estimates.

TABLE II
RELATIVE DIFFERENCE BETWEEN Γ_M AND Γ_L

	Trace	Frobenius Norm	Förstner Distance
Building 177	0.331	1.820	0.319
Building 306	0.509	2.601	0.282
Mars Yard	0.289	0.758	0.356
Parking Lot	0.063	0.929	0.539

normalized by the distance between the MCMC covariance and the prior covariance:

$$d_F = \frac{d_f(\Gamma_M, \Gamma_L)}{d_f(\Gamma_M, \Gamma_{pr})}. \quad (15)$$

The results of these comparisons are shown in Table II for the four test cases. The relative differences in trace are greater than 25% with the exception of Parking Lot. While Parking Lot has one of the lowest differences in trace, it has the highest relative difference in Förstner distance as defined in (15). By including the cross-covariances, the relative differences in Frobenius norm are over 75%. This numerical comparison using multiple metrics establishes that there is a nontrivial deviation in the covariances obtained using the approximate Bayesian and fully Bayesian algorithms, motivating the exploration of the differences in more detail.

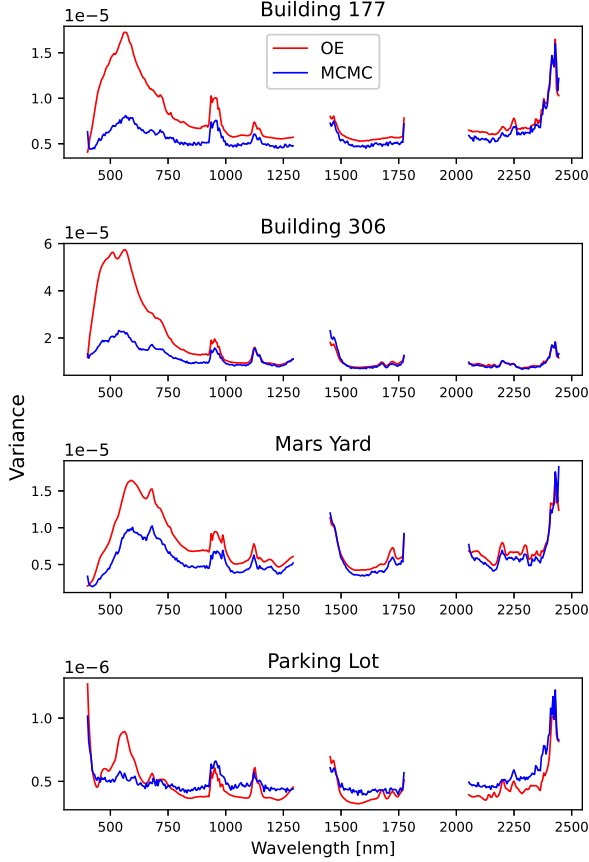


Fig. 8. Marginal variance in reflectances predicted by MCMC and OE methods.

B. Eigenanalysis of the surface posterior

Expanding on the eigenvalue problem used in the Förstner distance metric, we explore the interpretability of eigenproblems to reveal structure in the difference between the two covariance matrices. We compare them with respect to the eigendirections of one of the matrices. Specifically, we focus on the following eigenvalue problem involving the sample covariance of the Building 177 posterior:

$$\Gamma_M v_M = \lambda_M v_M, \quad (16)$$

where λ_M and v_M are the eigenvalues and eigenvectors of Γ_M . The eigenvalue spectrum is shown in Figure 9.

Then, the variance of Γ_L in the direction $v_{M,i}$ can be expressed as $v_{M,i}^\top \Gamma_L v_{M,i}$. This directional variance can be normalized using the corresponding eigenvalue as follows:

$$\sigma_{M,i}^L = \frac{v_{M,i}^\top \Gamma_L v_{M,i}}{v_{M,i}^\top \Gamma_M v_{M,i}} = \frac{v_{M,i}^\top \Gamma_L v_{M,i}}{\lambda_{M,i}} \quad (17)$$

Figure 10 plots this quotient, ranked in order of highest to lowest eigenvalue $\lambda_{M,i}$. A value greater than 1 can be interpreted as the Laplace approximation having greater variance in the $v_{M,i}$ direction, and vice versa. Consistent with Figure 8, the Laplace approximation overestimates the variance of the MCMC variance in most directions. The overall pattern is that the variances are similar along most of the top 25

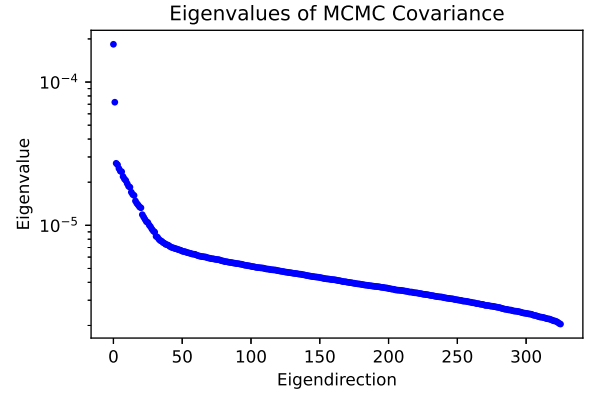


Fig. 9. Eigenvalues of the MCMC covariance matrix.

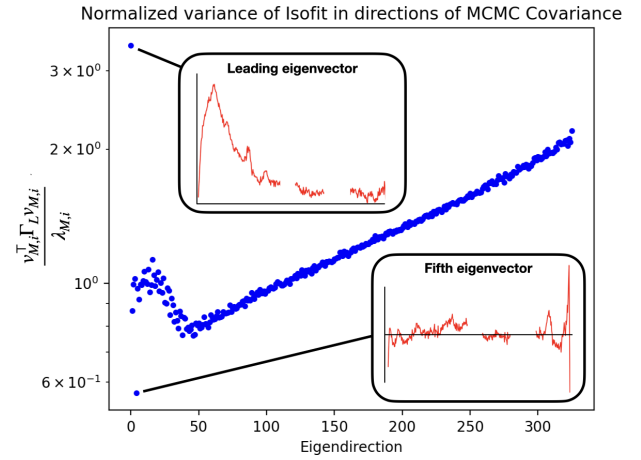


Fig. 10. Ratio of Laplace approximation variance and MCMC variance in the eigendirections of the MCMC covariance.

eigendirections, but the MCMC variance becomes smaller than the Laplace approximation variance as we move into the less important directions toward the right of the plot.

There are two outliers among the leading eigendirections, and the corresponding eigenvectors are shown in red in Figure 10. Comparing the shape to the posterior variance plot in Figure 8, the first eigenvector (top outlier) resembles the main reflectance feature in the lower wavelengths. The Laplace approximation predicts variance three times higher in this direction. In Section V-D we show how this may be related to the non-Gaussianity in the low wavelength region, which makes the Laplace approximation less accurate. The fifth eigenvector (bottom outlier) describes some of the noisy features, particularly the spike near 2500 nm, and the MCMC result predicts a variance around 70% higher than the Laplace approximation in this direction.

C. Atmospheric posterior comparison

While the reflectances are the quantities of interest ultimately used in subsequent analysis of the Earth surface, their behaviour is conditioned on the atmospheric parameters.

Figure 11 is a 2D marginal density plot of the posterior for the two atmospheric parameters. The MAP estimate

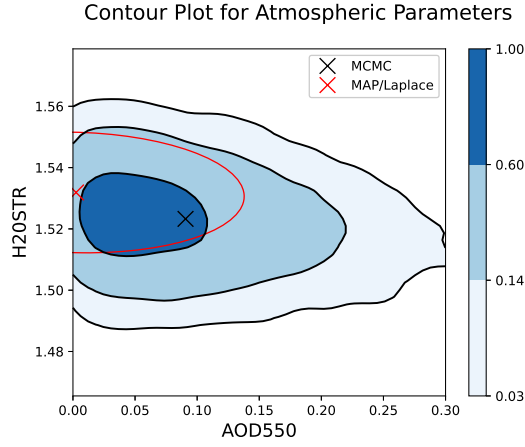


Fig. 11. 2D marginal density plot of the atmospheric posterior distribution for Building 177.

from optimal estimation is plotted in red along with an ellipse representing one standard deviation obtained using the Laplace approximation. There are two visible improvements from characterizing the posterior distribution using MCMC. First, optimal estimation has no way of ensuring positivity of the parameters, so the probabilistic interpretation is that the probability is obtaining a negative AOD parameter is almost 0.5, for example. The MCMC implementation fixes the samples to be positive and therefore leads to results that are more representative of the physical quantities. The second improvement is that MCMC sampling reveals a non-elliptical shape to the posterior, suggesting that it is not Gaussian. The Gaussianity of the posterior for both surface and atmospheric parameters is further explored in the next subsection.

D. Evaluating Gaussianity

The motivation for turning to a fully Bayesian approach is that the posterior is non-Gaussian in general. Here, we first demonstrate the non-Gaussianity of the posterior distribution qualitatively using normal Q-Q plots, and then quantitatively using hypothesis testing for individual parameters in one dimension.

Figure 12 shows the Q-Q plots for the two atmospheric parameters across all four cases. The red line is the reference for a truncated normal distribution and the MCMC samples are plotted in blue. The truncated normal was used since the samples for the atmospheric block were constrained to positive values in the algorithm. While the H_2O parameters closely follow the truncated normal, the right tail of the AOD plots deviate from the red line, especially for the Building 177 and Parking Lot cases. Qualitatively, these two cases look the least Gaussian, which is likely related to the absolute value of the reflectances and radiance being lower. Comparing the overall brightness levels in Figure 5, Building 177 and Parking Lot are the darkest. Building 306, the brightest, has a Q-Q plot for AOE that looks the most Gaussian in Figure 12. This may be the case because the additive contribution of aerosol scattering to the observed radiance is smaller in proportion to the large surface reflected signal.

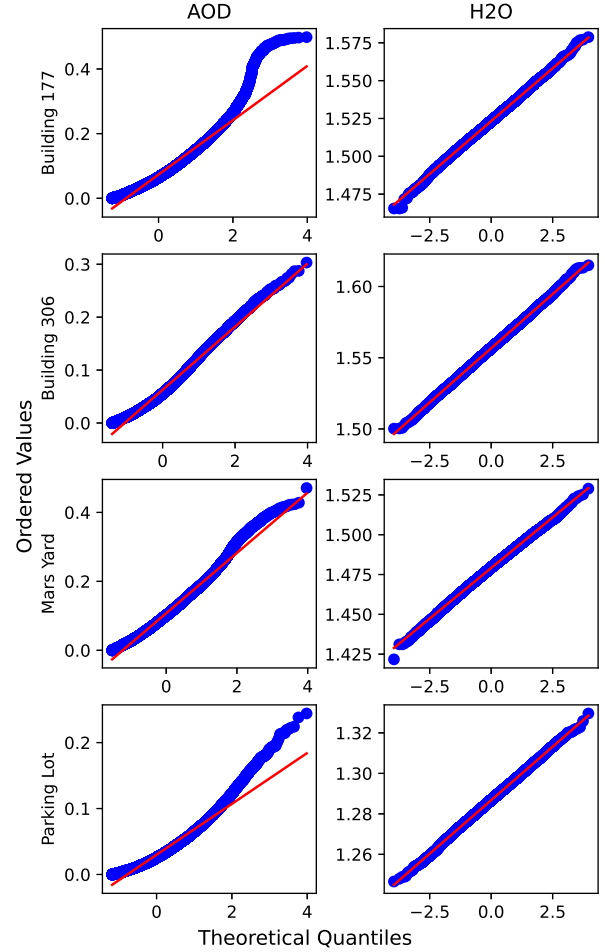


Fig. 12. Q-Q plots of atmospheric parameters across all four cases. The red line is a reference indicating the truncated normal distribution.

Figure 13 shows the Q-Q plots for select reflectance parameters across the spectrum for the Building 177 case. Although the MCMC samples closely follow a normal distribution, the two plots for 596 nm and 746 nm have tails that deviate from the reference normal. This is likely due to the effect of the non-Gaussian AOD parameter. We present a more comprehensive analysis of the reflectances at all wavelengths using a hypothesis testing approach.

Treating the reflectances individually, we use the Kolmogorov-Smirnov test [31] on the empirical marginal distribution of the MCMC samples with the null hypothesis being that the reflectances are normally distributed. The p -values for each reflectance parameter are shown in Figure 14, with the red line representing $p = 0.05$. In three of the cases, the non-Gaussian phenomenon observed in certain wavelengths in Figure 13 are present in the entire low wavelength regime, with $p \approx 0$. However, the extent of this regime varies for all three cases, with the low p -value region in the Mars Yard case extending to nearly 1000 nm.

This is consistent with the findings in Figure 10, which showed that the largest difference between the OE and MCMC posterior was in the direction that represents the lower wavelength regime. The departure from Gaussianity for the reflectances in this regime may be the reason why the Laplace

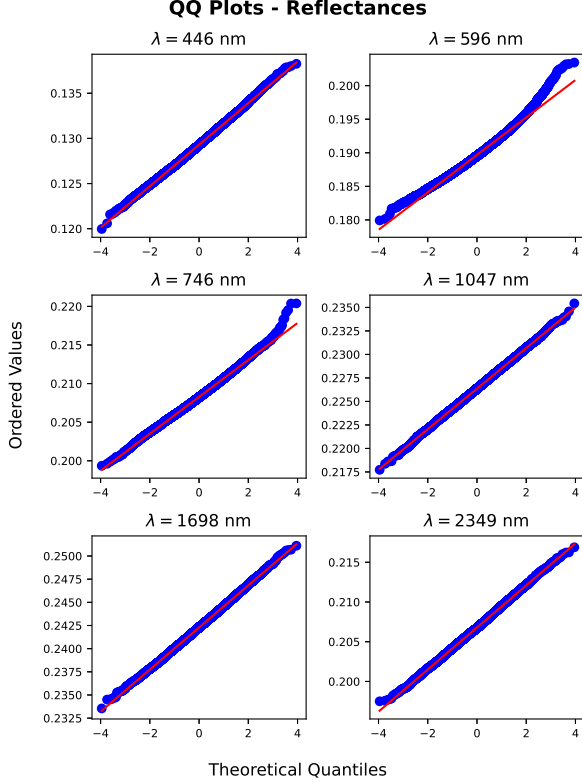


Fig. 13. Q-Q plots of select reflectance parameters for the Building 177 case.

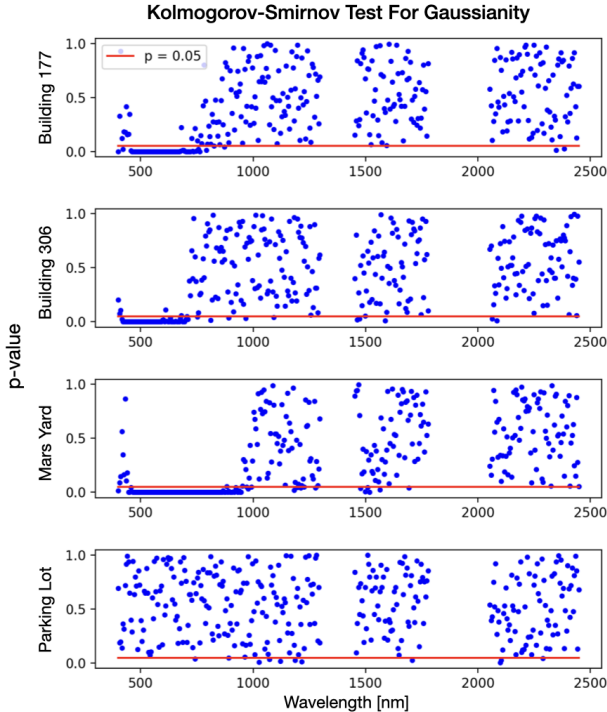


Fig. 14. Kolmogorov-Smirnov test for reflectances across all four cases.

approximation also departs from the posterior characterized by MCMC.

VI. DISCUSSION

We presented a fully Bayesian MCMC algorithm for the remote sensing problem that characterizes the posterior distribution of the surface reflectances and atmospheric parameters. This posterior was used to identify and understand the limitations of optimal estimation, the current state-of-the-art approximate Bayesian approach. There are three main takeaways from the results presented in this paper.

- The fully Bayesian solution and Laplace approximation produce posterior mean reflectance estimates that are virtually identical for the range of conditions studied.
- The fully Bayesian solution and Laplace approximation yield very different covariances. We analyzed the differences in terms of three different metrics and the eigendirections of the MCMC covariance matrix.
- The posterior distribution of the atmospheric parameters is more physically sensible than the Laplace approximation.
- Non-Gaussianity in the posterior is revealed by the fully Bayesian solution.

We identified regions of the spectrum and in the atmospheric parameters for which the Laplace approximation would not be able to sufficiently represent the non-Gaussian distribution. From the eigenanalysis, the OE posterior covariance was shown to be the most different from the MCMC posterior covariance in the low-wavelength region, which is the same region that was shown to depart from Gaussianity. Any further work on non-Gaussian posterior characterizations could focus on only this regime of the reflectances and the AOD atmospheric parameter. There is potential to develop a new combined method that uses MCMC or another non-Gaussian method for these parts, and OE for the rest of the parameters.

While we based our analysis on four test cases, we expect MCMC to be just as appropriate in other cases when we compare using the same forward radiative transfer model. Regions with heavier aerosol loads are expected to exhibit more non-Gaussianity, making MCMC even more appealing than OE in these cases.

In terms of NASA's Surface Geology and Biology mission, characterizing the posterior distribution is important for subsequent analysis. The surface reflectances are ultimately used to further infer properties of the Earth surface pertaining to themes such as terrestrial and aquatic ecosystems, geology, and water resources. Accurately quantifying the uncertainty of the surface properties is especially important for these scientific applications.

A. Limitations

Although we have created a computationally tractable fully Bayesian algorithm by exploiting structure in the problem, it is not computationally feasible in an operational setting. Generating 2×10^6 posterior samples takes on the order of hours, whereas the approximate Bayesian approach for one retrieval is on the order of seconds. However, the two methods

could be combined for the operational setting. For example, a fully Bayesian retrieval can be used as a validation step to verify and potentially correct the Laplace approximation. Or, since the Laplace approximation is sufficient to approximate many of the parameters, the algorithm in this work can be narrowed to solve a smaller subproblem on a subset of parameters using a fully Bayesian approach while maintaining the original approximate approach for the other parameters.

The performance of the MCMC algorithm is contingent on the prior, noise model, and forward model. Since the noise model dictates the likelihood distribution, the posterior would change along with changes in the prior distribution and noise model. The forward model is subject to modelling error, which affects the interactions between the atmospheric and surface parameters. This would be one of the main improvements to the current algorithm when considering operational use, since ultimately we are interested in matching the ground truth reflectances.

B. Future work

The retrievals performed in this work are pixel-by-pixel, meaning that for each radiance vector, there is one corresponding state vector to be inferred. The main efforts for future work are focused on extending fully Bayesian algorithms to spatial and temporal fields. Including spatial and temporal correlations can increase the retrieval accuracy and reduce the number of retrievals required, which would reduce computational time in an operational setting.

VII. CONCLUSION

In this work, we developed a computationally tractable fully Bayesian retrieval method for the high-dimensional VSWIR retrieval. Leveraging the conditional linearity structure in the forward radiative transfer model, which is approximately linear with respect to the surface for a fixed atmosphere, we implemented a block Metropolis MCMC algorithm that alternates samples between the atmospheric and surface parameter blocks. Unlike other algorithms that involve dimension reduction, the block Metropolis algorithm is asymptotically exact.

We compared this fully Bayesian algorithm to the current state-of-the-art optimal estimation method in several ways to identify limitations of the approximate Bayesian approach. For the surface parameters, both methods had a similar posterior mean but the posterior variance was shown to be significantly different. The eigendirections in which they are different were analyzed and interpreted. Through hypothesis testing, MCMC revealed non-Gaussianity in the aerosol parameter and the low-wavelength regime of the surface reflectances. This work is the first step in combining the block Metropolis algorithm with optimal estimation to allow non-Gaussian retrievals to be performed in an operational setting.

ACKNOWLEDGMENTS

A portion of this research was performed at the Jet Propulsion Laboratory, California Institute of Technology, under

contract with the National Aeronautics and Space Administration. This research was funded in part by a Jet Propulsion Laboratory Strategic University Partnership grant.

REFERENCES

- [1] M. National Academies of Sciences, Engineering, *Thriving on Our Changing Planet: A Decadal Strategy for Earth Observation from Space*. Washington, DC: The National Academies Press, 2018. [Online]. Available: <https://nap.nationalacademies.org/catalog/24938/thriving-on-our-changing-planet-a-decadal-strategy-for-earth>
- [2] K. Cawse-Nicholson, P. A. Townsend, D. Schimel, A. M. Assiri, P. L. Blake, M. F. Buongiorno, P. Campbell, N. Carnon, K. A. Casey, R. E. Correa-Pabón, K. M. Dahlin, H. Dashti, P. E. Dennison, H. Dierssen, A. Erickson, J. B. Fisher, R. Frouin, C. K. Gatebe, H. Gholizadeh, M. Gierach, N. F. Glenn, J. A. Goodman, D. M. Griffith, L. Guild, C. R. Hakkenberg, E. J. Hochberg, T. R. Holmes, C. Hu, G. Hulley, K. F. Huemmrich, R. M. Kudela, R. F. Kokaly, C. M. Lee, R. Martin, C. E. Miller, W. J. Moses, F. E. Muller-Karger, J. D. Ortiz, D. B. Otis, N. Pahlevan, T. H. Painter, R. Pavlick, B. Poulter, Y. Qi, V. J. Realmuto, D. Roberts, M. E. Schaepman, F. D. Schneider, F. M. Schwandner, S. P. Serbin, A. N. Shiklomanov, E. N. Stavros, D. R. Thompson, J. L. Torres-Perez, K. R. Turpie, M. Tzortziou, S. Ustin, Q. Yu, Y. Yusup, and Q. Zhang, "NASA's surface biology and geology designated observable: A perspective on surface imaging algorithms," *Remote Sensing of Environment*, vol. 257, p. 112349, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0034425721000675>
- [3] K. Cawse-Nicholson, A. Raiho, D. Thompson, G. Hulley, C. Miller, K. Miner, B. Poulter, D. Schimel, F. Schneider, P. Townsend, and S. Zareh, "Surface Biology and Geology imaging spectrometer: A case study to optimize the mission design using intrinsic dimensionality," *Remote Sensing of Environment*, vol. 290, p. 113534, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0034425723000858>
- [4] J. Tamminen and E. Kyrölä, "Bayesian solution for nonlinear and non-Gaussian inverse problems by Markov chain Monte Carlo method," *Journal of Geophysical Research: Atmospheres*, vol. 106, no. D13, pp. 14 377–14 390, 2001.
- [5] J. Tamminen, "Validation of nonlinear inverse algorithms with markov chain monte carlo method," *Journal of Geophysical Research: Atmospheres*, vol. 109, no. D19, 2004. [Online]. Available: <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2004JD004927>
- [6] H. Haario, M. Laine, M. Lehtinen, E. Saksman, and J. Tamminen, "Markov Chain Monte Carlo Methods for High Dimensional Inversion in Remote Sensing," *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, vol. 66, no. 3, pp. 591–607, 2004. [Online]. Available: <http://www.jstor.org/stable/3647497>
- [7] Y. Wang, X. Jiang, B. Yu, and M. Jiang, "A Hierarchical Bayesian Approach for Aerosol Retrieval Using MISR Data," *Journal of the American Statistical Association*, vol. 108, no. 502, pp. 483–493, 2013. [Online]. Available: <https://doi.org/10.1080/01621459.2013.796834>
- [8] S. Tukiainen, J. Railo, M. Laine, J. Hakkarainen, R. Kivi, P. Heikkinen, H. Chen, and J. Tamminen, "Retrieval of atmospheric CH₄ profiles from Fourier transform infrared data using dimension reduction and MCMC," *Journal of Geophysical Research: Atmospheres*, vol. 121, no. 17, pp. 10–312, 2016.
- [9] T. Cui and O. Zahm, "Data-free likelihood-informed dimension reduction of Bayesian inverse problems," *Inverse Problems*, vol. 37, 2021.
- [10] O. Lamminpää, M. Laine, S. Tukiainen, and J. Tamminen, "Likelihood informed dimension reduction for remote sensing of atmospheric constituent profiles," *2017 MATRIX Annals*, 2017. [Online]. Available: <https://api.semanticscholar.org/CorpusID:88515698>
- [11] O. Lamminpää, J. Hobbs, J. Brynjarsdóttir, M. Laine, A. Braverman, H. Lindqvist, and J. Tamminen, "Accelerated MCMC for Satellite-Based Measurements of Atmospheric CO₂," *Remote Sensing*, vol. 11, no. 17, 2019. [Online]. Available: <https://www.mdpi.com/2072-4292/11/17/2061>
- [12] D. R. Thompson, V. Natraj, R. O. Green, M. C. Helmlinger, B.-C. Gao, and M. L. Eastwood, "Optimal estimation for imaging spectrometer atmospheric correction," *Remote Sensing of Environment*, vol. 216, pp. 355–373, 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0034425718303304>
- [13] R. O. Green, N. Mahowald, D. R. Thompson, C. Ung, P. Brodrick, R. Pollock, M. Bennett, S. Lundeen, M. Joyce, W. Olson-Duvall, B. Oaida, C. Bradley, E. Diaz, R. Clark, S. Vannan, G. Swayze, R. Kokaly, P. Ginoux, R. Miller, G. Okin, C. P. Garcia-Pando,

- B. Ehlmann, O. Kalashnikova, T. H. Painter, V. Realmuto, D. Chadwick, E. Ben-Dor, D. H. Pearlshien, L. Gunter, B. Phillips, K. Reath, A. Thorpe, L. Shaw, A. Keebler, F. Ochoa, K. Grant, A. Sen, R. Duren, V. Obiso, M. Gonçalves-Ageitos, and Y. Huang, "Performance and Early Results from the Earth Surface Mineral Dust Source Investigation (EMIT) Imaging Spectroscopy Mission," in *2023 IEEE Aerospace Conference*, 2023, pp. 1–10.
- [14] D. R. Thompson, R. O. Green, C. Bradley, P. G. Brodrick, N. Mahowald, E. B. Dor, M. Bennett, M. Bernas, N. Carmon, K. D. Chadwick *et al.*, "On-orbit calibration and performance of the emit imaging spectrometer," *Remote Sensing of Environment*, vol. 303, p. 113986, 2024.
- [15] E. E. Team, "Ecological spectral information system (Eco-SIS)," 2020.
- [16] A. Berk, P. Conforti, R. Kennett, T. Perkins, F. Hawes, and J. van den Bosch, "MODTRAN6: a major upgrade of the MODTRAN radiative transfer code," in *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XX*, M. Velez-Reyes and F. A. Kruse, Eds., vol. 9088, International Society for Optics and Photonics. SPIE, 2014, pp. 113 – 119. [Online]. Available: <https://doi.org/10.1117/12.2050433>
- [17] D. R. Thompson, L. Gunter, A. Berk, B.-C. Gao, R. Richter, D. Schläpfer, and K. J. Thome, "Retrieval of atmospheric parameters and surface reflectance from visible and shortwave infrared imaging spectroscopy data," *Surveys in Geophysics*, vol. 40, no. 3, pp. 333–360, May 2019. [Online]. Available: <https://doi.org/10.1007/s10712-018-9488-9>
- [18] D. R. Thompson, A. Braverman, P. G. Brodrick, A. Candela, N. Carmon, R. N. Clark, D. Connelly, R. O. Green, R. F. Kokaly, L. Li *et al.*, "Quantifying uncertainty for remote spectroscopy of surface composition," *Remote Sensing of Environment*, vol. 247, p. 111898, 2020.
- [19] W. Fulks and J. O. Sather, "Asymptotics. II. Laplace's method for multiple integrals," *Pacific Journal of Mathematics*, vol. 11, no. 1, pp. 185 – 192, 1961. [Online]. Available: <https://doi.org/10.1137/140977308>
- [20] A. Spantini, A. Solonen, T. Cui, J. Martin, L. Tenorio, and Y. Marzouk, "Optimal Low-rank Approximations of Bayesian Linear Inverse Problems," *SIAM Journal on Scientific Computing*, vol. 37, no. 6, pp. A2451–A2487, 2015. [Online]. Available: <https://doi.org/10.1137/140977308>
- [21] L. Tierney and J. B. Kadane, "Accurate approximations for posterior moments and marginal densities," *Journal of the American Statistical Association*, vol. 81, no. 393, pp. 82–86, 1986. [Online]. Available: <https://www.tandfonline.com/doi/abs/10.1080/01621459.1986.10478240>
- [22] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, "Equation of State Calculations by Fast Computing Machines," *The Journal of Chemical Physics*, vol. 21, no. 6, pp. 1087–1092, 1953. [Online]. Available: <https://doi.org/10.1063/1.1699114>
- [23] O. Zahm, T. Cui, K. J. H. Law, A. Spantini, and Y. M. Marzouk, "Certified dimension reduction in nonlinear Bayesian inverse problems," *Math. Comput.*, vol. 91, pp. 1789–1835, 2018. [Online]. Available: <https://api.semanticscholar.org/CorpusID:53696711>
- [24] P. G. Constantine, E. Dow, and Q. Wang, "Active subspace methods in theory and practice: Applications to kriging surfaces," *SIAM Journal on Scientific Computing*, vol. 36, no. 4, pp. A1500–A1524, 2014. [Online]. Available: <https://doi.org/10.1137/130916138>
- [25] S. Chib and E. Greenberg, "Understanding the Metropolis-Hastings Algorithm," *The American Statistician*, vol. 49, no. 4, pp. 327–335, 1995. [Online]. Available: <https://www.tandfonline.com/doi/abs/10.1080/00031305.1995.10476177>
- [26] H. Haario, E. Saksman, and J. Tamminen, "An adaptive Metropolis algorithm," *Bernoulli*, vol. 7, no. 2, pp. 223 – 242, 2001. [Online]. Available: <https://doi.org/10.1080/00031305.1995.10476177>
- [27] A. Gelman, W. R. Gilks, and G. O. Roberts, "Weak convergence and optimal scaling of random walk Metropolis algorithms," *The annals of applied probability*, vol. 7, no. 1, pp. 110–120, 1997.
- [28] L. Hamlin, R. O. Green, P. Mouroulis, M. Eastwood, D. Wilson, M. Dudik, and C. Paine, "Imaging spectrometer science measurements for Terrestrial Ecology: AVIRIS and new developments," in *2011 Aerospace Conference*, 2011, pp. 1–7.
- [29] H. A. Bender, P. Z. Mouroulis, R. O. Green, and D. W. Wilson, "Optical design, performance, and tolerancing of next-generation airborne imaging spectrometers," in *Imaging Spectrometry XV*, S. S. Shen and P. E. Lewis, Eds., vol. 7812, International Society for Optics and Photonics. SPIE, 2010, p. 78120P. [Online]. Available: <https://doi.org/10.1117/12.861331>
- [30] W. Förstner and B. Moonen, *A Metric for Covariance Matrices*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, pp. 299–309. [Online]. Available: https://doi.org/10.1007/978-3-662-05296-9_31

- [31] D. A. Darling, "The Kolmogorov-Smirnov, Cramer-von Mises Tests," *The Annals of Mathematical Statistics*, vol. 28, no. 4, pp. 823–838, 1957. [Online]. Available: <http://www.jstor.org/stable/2237048>



approaches to Bayesian inference. With a background in aerospace engineering from the University of Toronto, Kelvin also received his S.M. degree from MIT in Aeronautics and Astronautics.



David R. Thompson is a Senior Research Scientist at the Jet Propulsion Laboratory, California Institute of Technology. He serves named roles on NASA imaging spectrometer space missions including EMIT, SBG-VSWIR, Lunar Trailblazer, and Carbon-I. He is recipient of the NASA Medals for Exceptional Engineering Achievement, Exceptional Technology Achievement, and Early Career Achievement. His code has flown on three Earth-orbiting spacecraft and twice to the surface of Mars.



Jouni Susiluoto is a researcher at the Jet Propulsion Laboratory (JPL) with expertise in data science and modeling applications in geosciences and remote sensing. He has worked on retrieval algorithms for imaging spectroscopy missions and collaborated with Caltech on developing fast emulators for radiative transfer models. Jouni's previous work includes a wide range of data science and modeling applications in geosciences and remote sensing, such as spatial data fusion and Bayesian model selection. Prior to JPL, he spent two years in the Uncertainty Quantification group of Department of Aeronautics and Astronautics at MIT, studying spatio-temporal statistical modeling of massive remote sensing data. Susiluoto has a doctorate in mathematics from University of Helsinki, Finland, where he was a member of the Inverse Problems group at the Department of Mathematics and Statistics. His research work and interests are in uncertainty quantification, inverse problems, Bayesian methods, MCMC, parameter estimation, climate models, data assimilation, and carbon cycle modeling.

Kelvin M. Leung is a Ph.D. candidate in the Uncertainty Quantification (UQ) group at the Massachusetts Institute of Technology (MIT). His research is broadly focused on Bayesian methods and inverse problems in the context of Earth remote sensing, including dimension reduction techniques and optimal transport-based

David R. Thompson is a Senior Research Scientist at the Jet Propulsion Laboratory, California Institute of Technology. He serves named roles on NASA imaging spectrometer space missions including EMIT, SBG-VSWIR, Lunar Trailblazer, and Carbon-I. He is recipient of the NASA Medals for Exceptional Engineering Achievement, Exceptional Technology Achievement, and Early Career Achievement. His code has flown on three Earth-orbiting spacecraft and twice to the surface of Mars.

Jouni Susiluoto is a researcher at the Jet Propulsion Laboratory (JPL) with expertise in data science and modeling applications in geosciences and remote sensing. He has worked on retrieval algorithms for imaging spectroscopy missions and collaborated with Caltech on developing fast emulators for radiative transfer models. Jouni's previous work includes



Dr. **Jayanth Jagalur** is a research scientist at the Center for Applied Scientific Computing (CASC) in the Lawrence Livermore National Laboratory (LLNL). His work and research interests include Data Assimilation & Bayesian Statistical Methods, Uncertainty Quantification, Design of experiments.



Dr. **Amy Braverman** is a Senior Research Scientist at NASA's Jet Propulsion Laboratory. She is the Technical Lead for Statistical Methods and Applications, and a member of the Uncertainty Quantification and Statistical Analysis Group. She graduated from Swarthmore College in 1982, and, after working in industry, returned to graduate school at UCLA in 1991. She earned a Ph.D. in Statistics, and began her second career as a post-doc at JPL in 1999. She has been at the Lab ever since. Dr. Braverman is at the forefront of JPL's efforts to bring rigorous statistical methods and UQ to remote sensing science.



Youssef Marzouk is a professor in the Department of Aeronautics and Astronautics at the Massachusetts Institute of Technology (MIT), and co-director of the Center for Computational Science and Engineering within the MIT Schwarzman College of Computing.

He is also a core member of MIT's Statistics and Data Science Center and a PI in the MIT Laboratory for Information and Decision Systems (LIDS).

His research interests lie at the intersection of statistical inference, computational mathematics, and physical modeling. He develops new methodologies for uncertainty quantification, Bayesian computation, and machine learning in complex physical systems, motivated by a broad range of engineering and science applications. Some of his main research contributions have centered on algorithms for inference, with applications to data assimilation and inverse problems; dimension reduction methodologies for high-dimensional learning and surrogate modeling; optimal experimental design; and transportation of measure as a tool for inference and stochastic modeling.

He received his SB, SM, and PhD degrees from MIT and spent four years at Sandia National Laboratories before joining the MIT faculty in 2009.