

Simulation-Based Uncertainty Quantification for Estimating Atmospheric CO₂ from Satellite Data*

Jonathan Hobbs[†], Amy Braverman[†], Noel Cressie[‡], Robert Granat[†], and Michael Gunson[†]

Abstract. Remote sensing of the atmosphere has provided a wealth of data for analyses and inferences in earth science. Satellite observations can provide information on the atmospheric state at fine spatial and temporal resolution while providing substantial coverage across the globe. For example, this capability can greatly enhance the understanding of the space-time variation of the greenhouse gas, carbon dioxide (CO₂), since ground-based measurements are limited. NASA's Orbiting Carbon Observatory-2 (OCO-2) collects tens of thousands of observations of reflected sunlight daily, and the mission's retrieval algorithm processes these indirect measurements into estimates of atmospheric CO₂. The retrieval is an inverse problem and consists of a physical forward model for the transfer of radiation through the atmosphere that includes absorption and scattering by gases, aerosols, and the surface. The model and other algorithm inputs introduce key sources of uncertainty into the retrieval problem. This article develops a computationally efficient surrogate model that is embedded in a simulation experiment for studying the impact of uncertain inputs on the distribution of the retrieval error.

Key words. Bayesian inference, inverse problem, surrogate model, radiative transfer, simulation experiment, optimal estimation, nonlinear model

AMS subject classifications. 62F15, 62P12

DOI. 10.1137/16M1060765

1. Introduction. In recent decades, atmospheric remote sensing has provided a wealth of data for understanding the Earth system. Remote sensing instruments, particularly Earth-orbiting satellites, exploit characteristics of electromagnetic radiation to make inferences about the state of the atmosphere. The retrieval problem, namely, estimating the atmospheric state from a satellite's observed radiation, is a primary scientific inference objective for remote sensing data. Each instrument has one or more associated *retrieval algorithms* that estimate a quantity of interest (QOI) from the instrument's observed radiances. Retrieval algorithms use a variety of approaches for estimating the atmospheric state. Some examples include construction of lookup tables, statistical modeling in combination with likelihood inference, and Bayesian inverse inference. Formal uncertainty quantification (UQ) can be a valuable tool

*Received by the editors February 9, 2016; accepted for publication (in revised form) June 2, 2017; published electronically September 6, 2017.

<http://www.siam.org/journals/juq/5/M106076.html>

Funding: The research described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with NASA. The work of the authors was supported by the Orbiting Carbon Observatory-2 (OCO-2) mission. The work of the third author was partially supported by NASA grant NNN11ZDA001N-OCO2 and by a 2015–2019 Australian Research Council Discovery Grant DP 150104576.

[†]Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109 (Jonathan.M.Hobbs@jpl.nasa.gov, Amy.Braverman@jpl.nasa.gov, Robert.Granat@jpl.nasa.gov, Michael.Gunson@jpl.nasa.gov).

[‡]School of Mathematics and Applied Statistics, University of Wollongong, Wollongong, NSW, Australia, and Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109 (ncressie@uow.edu.au).

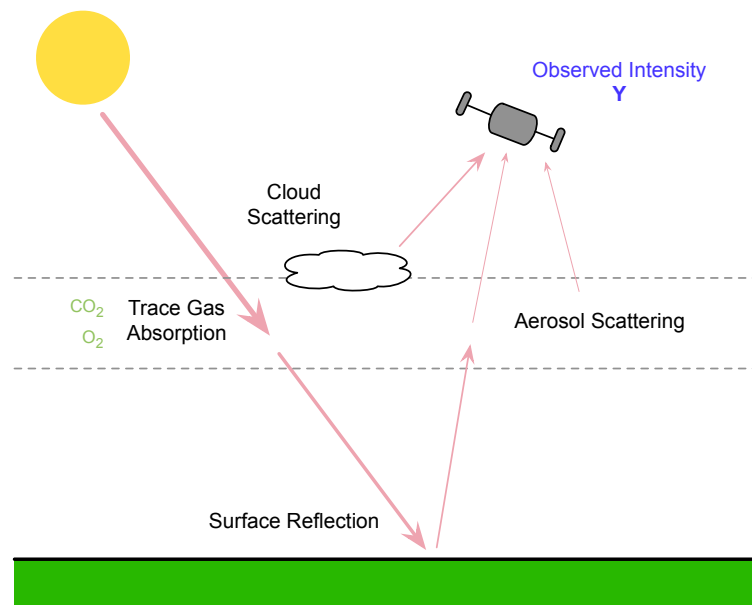


Figure 1. Summary of key sources and sinks of radiation along a path through the atmosphere to the satellite.

in any of these situations by providing a framework for propagating the impact of algorithm choices, including the sources of uncertainty that accompany them, through the retrieval process.

In satellite remote sensing, the QOI (e.g., the atmospheric state) is inferred from observable radiance spectra (Figure 1), making inference an example of an inverse problem. Inverse problems present a number of challenges, including a tendency to be ill-posed and highly sensitive, particularly when the relationship between the state and the observation is nonlinear [17, 7]. Bayesian inference is an appealing option in this situation because additional information about the state or other model parameters can be introduced. In remote sensing, this approach has been developed into the so-called optimal estimation (OE) retrieval [22]. In OE retrieval, the distribution of the observed spectra given the state and the (marginal) distribution of the state are modeled probabilistically. From these distributions, a posterior distribution of the state given the observed spectra can be used to infer the unknown state. Because of the inherently nonlinear relationship between the state and the observed spectra, in practice this posterior distribution is rarely available in closed form.

There are a number of strategies for interrogating the resulting posterior distribution, and practical considerations, such as the volume of data to be processed and the computational expense of the nonlinear forward model relating the radiances to the state, often take priority. Markov chain Monte Carlo (MCMC) sampling from the posterior distribution has been implemented in remote sensing retrieval problems [25, 12], but this approach requires a large number of forward model evaluations. The recently launched Orbiting Carbon Observatory-2 (OCO-2) provides tens of thousands of retrievals per day, requiring the retrieval process to be computationally fast [9, 19]. The data volume means that the information extracted from the posterior distribution is minimal, being restricted to a point estimate and an approximate

covariance matrix. As detailed in section 2.2, a typical approach is to search for the posterior mode, the maximum a posteriori estimate, with numerical approaches and to obtain the covariance matrix through linearization. Some theoretical aspects of this retrieval process have been demonstrated [7, 8], and linear error analysis has identified key sensitivities for this OE retrieval [23, 4].

The present paper develops a simulation-based framework for the OE retrieval applied to atmospheric CO₂ retrievals that addresses several objectives. First, the approach samples the retrieval error distribution under standard conditions without assuming linearity. Second, it characterizes the impact of key OE-algorithm choices on the distribution of the retrieval error. Finally, it is contrasted with the linear error analysis that is commonly used in remote sensing retrievals through a retrieval error budget that separates contributions from linear and nonlinear sources. In the process, the true bias and covariance of the retrieval errors can be determined. This approach and the underlying statistical model resemble simulation studies of nonlinear mixed effects (NLME) models [13, 14]. In the remote sensing application, the inference objective focuses on the state, which would be considered the random effect in the NLME context. A simulation framework allows an extension of the linear approximation in traditional OE retrieval error analysis [23]. This simulation-based strategy requires an OE retrieval that is computationally fast in order to facilitate large Monte Carlo sample sizes in the simulation experiment. In fact, the OCO-2 operational algorithm is not fast enough, so we develop a surrogate forward model and retrieval.

This article is organized as follows. Section 2 describes OCO-2 and its role in carbon cycle science, along with the mathematical details for the OE retrieval. Section 3 outlines a UQ simulation framework and an associated surrogate model. Section 4 describes a simulation experiment that examines dominant sources of uncertainty for OCO-2, with the results discussed in section 5. Section 6 offers some concluding remarks and future research directions.

2. Remote sensing and OCO-2. Later sections summarize simulation experiments using a nonlinear radiative transfer model and OE retrieval. Figure 2 provides a schematic overview of this framework, which could be applied to retrievals from a general remote sensing instrument. A particular instance requires an appropriate forward model for simulating synthetic radiances from specified atmospheric states, plus a retrieval algorithm for estimating the state given the observed radiances. The experiment developed in section 4 specifically targets the OE retrieval and radiative transfer model for estimating atmospheric CO₂ concentration. As motivation, we provide background on this measurement and the mathematical framework for the OE retrieval.

The OCO-2 launched in July 2014 with an objective of providing global estimates of atmospheric carbon dioxide at fine spatial resolution. OCO-2's primary QOI is the column-averaged dry air mole fraction of CO₂, a quantity termed X_{CO_2} . The estimation of X_{CO_2} is discussed further in section 2.2. The OCO-2 instrument's global coverage and data volume are providing a more comprehensive picture of atmospheric carbon dioxide (CO₂) concentration, especially regional spatial patterns, seasonal cycles, and interannual variability. Remote sensing data are an important data source for CO₂, since in situ measurements are sparse and concentrated in midlatitude land regions. A comprehensive picture of the CO₂ field can aid the understanding of the global carbon cycle. In particular, X_{CO_2} estimates are combined

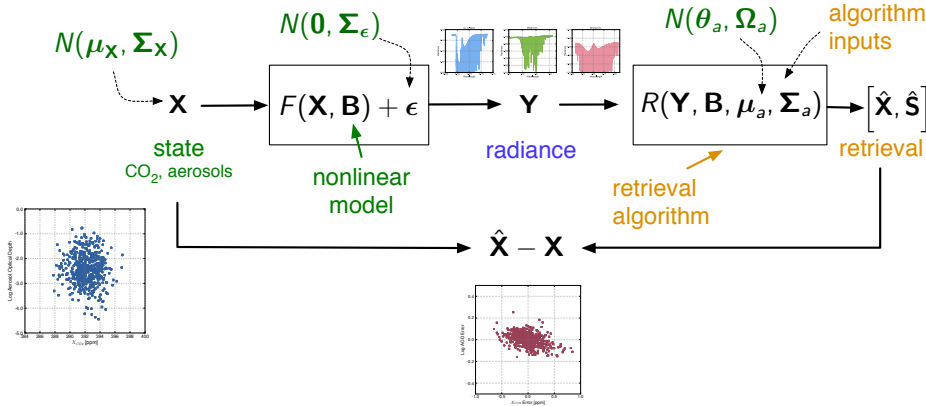


Figure 2. Schematic diagram of the Monte Carlo framework using the OCO-2 surrogate model.

with transport models to infer carbon fluxes between the surface and the lower atmosphere. Fluxes vary substantially across the globe, with source regions often located in close proximity to sink regions, such as in the tropics where substantial deforestation has occurred [1].

Emissions from human activities such as fossil-fuel burning and land-use change are key components of the global carbon budget. The combined land and ocean sinks remove approximately half of anthropogenic carbon emissions, but there is pronounced year-to-year variability in this proportion [3]. The mechanisms behind this variability are largely unknown, and substantial uncertainty exists as to the relative impact of tropical forests and boreal forests of the northern hemisphere as land carbon sinks. Continuous monitoring across the globe from remote sensing instruments has the potential to more precisely identify sources and sinks and their evolution over time. At the same time, appropriate uncertainties must be attached to the remote sensing retrievals so that they can be propagated through the flux-inversion process. A comprehensive understanding of the OCO-2 OE retrieval and associated sources of uncertainty is a critical component of this end-to-end inference problem.

2.1. Measurement. The OCO-2 instrument includes three imaging grating spectrometers measuring solar radiation reflected from the Earth's surface in the infrared (IR) portion of the spectrum. Each spectrometer corresponds to an IR band with a resolution of approximately 1000 wavelengths (colors) over a narrow wavelength range of less than 50 nm. Molecular oxygen (O_2) absorbs strongly in one of the bands, termed the O_2 -A band, and the other two bands are known as the weak CO_2 band and the strong CO_2 band. The collection of observed radiances from the three bands at a particular time make up a *sounding*. The satellite is in sun-synchronous polar orbit in a formation of satellites called the A-train at 700 km above the Earth's surface. The orbit track crosses the equator on the daytime side in the early afternoon local time, and about 15 orbits are completed each day [9].

Let the random vector \mathbf{Y} represent the set of radiances for a single OCO-2 sounding. Figure 3 gives an example of a radiance vector from the surrogate forward model outlined in section 3. The observed radiances are a result of the interaction between the radiation and the composition of the atmosphere and of the Earth's surface along the path from the top of the

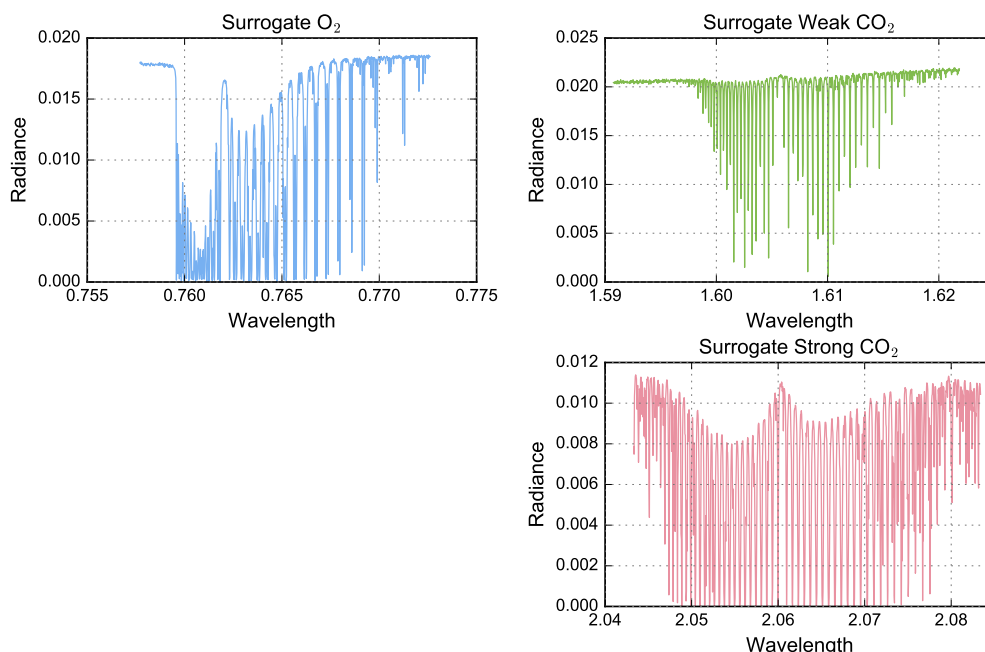


Figure 3. Example of a radiance vector \mathbf{Y} .

atmosphere to the surface and back to the satellite. The general goal is to estimate the atmospheric state, which we denote as \mathbf{X} , from the observed radiances, along with characterizing the uncertainty of the estimate. In particular, certain atmospheric constituents will absorb and/or scatter radiation. The extent of absorption and scattering depends on the wavelength as well as the amount and type of the constituent, as shown in Figure 1.

The mathematical relationship between the atmospheric state \mathbf{X} and the radiances \mathbf{Y} is captured through a *forward model*, $\mathbf{F}(\mathbf{X}, \mathbf{B})$. The inputs of the forward model include the state as well as a set of forward-model parameters \mathbf{B} that are characteristics of the instrument and any other quantities not included in the state \mathbf{X} . In general the parameters are not perfectly known, and their treatment in the retrieval problem is discussed in the next subsection.

For many remote sensing applications, including OCO-2, the forward model discretizes the atmospheric vertical profile into a set of layers. The composition of different layers can be different, but the atmosphere is assumed homogenous within a layer. This discretization allows for a numerical solution to the equation of radiative transfer (RT), and this numerical solution is the resulting value of $\mathbf{F}(\mathbf{X}, \mathbf{B})$. For the OCO-2 surrogate model defined in section 3, the elements of the state vector can be grouped into the following general categories:

- **CO₂ vertical profile.** The dry-air mole fraction, or the number of moles of CO₂ per mole of dry air, varies vertically in the atmospheric column. For OCO-2, it is defined at 20 fixed pressure levels in the atmosphere, corresponding to the upper and lower boundaries of each of the discrete layers. Absorption of CO₂ occurs at numerous wavelengths, often called *absorption lines*, in both the strong and weak CO₂ bands. Therefore, the amount of CO₂ present is strongly related to the radiances at many wavelengths in these bands. This relationship reflects the total number of molecules

of CO₂ present and, hence, additional information about the total amount of dry air is required.

- **Surface pressure.** The surface pressure is a single component of the state vector that helps identify the total number of molecules of air in the atmospheric column. Since molecular O₂ has a nearly constant dry air mole fraction anywhere in the atmosphere, the absorption of O₂ can accurately reflect the total amount of dry air. Surface pressure is identified with this information and a representation of the presence of water vapor in the atmosphere. Many O₂ absorption lines are present in the O₂ A-band.
- **Surface albedo.** Earth's surface acts as a boundary condition in the RT problem. Some radiation is extinguished and some is reflected at the surface. Surface albedo is the fraction of reflected radiation to total incoming radiation at the surface. This behavior varies as a function of wavelength. The state vector includes two albedo coefficients for each of the three bands. The first is the albedo at a reference wavelength at the center of the band (intercept), and the second is a slope that defines the linear change in albedo across the band.
- **Aerosols.** Small particles in the atmosphere interact with incoming radiation in complex ways. Some radiation is extinguished, and the extent of this extinction is often summarized by aerosol optical depth (AOD), which is defined as the natural logarithm of the ratio of incoming to transmitted radiation. Since the ratio is larger than unity, AOD is strictly positive, and larger values correspond to more opaque conditions due to radiation extinction by aerosols. In addition, some radiation is scattered in different directions, represented as different angles with respect to the direct path from the sun. The forward model accounts for the angular dependence of scattering through a phase function. The OCO-2 state vector includes three coefficients to describe the aerosol vertical profile for up to four different aerosol types. For a given aerosol type, one coefficient is the natural logarithm of the total AOD in the O₂ A-band. The second coefficient represents the vertical height where the aerosol concentration is a maximum. The third coefficient represents the depth of the aerosol profile; a small value indicates a “thin” aerosol layer. The state vector can include these coefficients for an arbitrary number of different aerosol types, which are characterized by different scattering properties in the forward-model parameters **B**.

These components represent the key state variables in our investigation and are included in the state vector for the surrogate model in section 3. These components, particularly aerosols and albedo, critically impact the uncertainty in retrieved CO₂ [19]. Their actual implementation in the RT model is outlined in Appendix B, and additional components of the full physics state vector are noted in Table 1. The OCO-2 mission's primary QOI is the CO₂ mole fraction, but it is important to include other components in the state vector because they play important roles in the forward model. Since they are not perfectly known, they are estimated as part of the retrieval. These additional quantities are often termed *nuisance parameters* in statistics and have been termed *interferences* in the remote sensing retrieval literature [23]. The CO₂ retrieval problem is particularly challenging due to the nonlinear nature of the forward model and the heterogeneous makeup of the state vector. Further, the sensitivity of the measured radiance to these interferences is often larger than to changes in CO₂.

Table 1

Composition of the state vector \mathbf{X} in the OCO-2 full physics (FP) forward model and in the surrogate forward model.

Component	FP	Surrogate
20-level CO ₂ profile	✓	✓
Surface air pressure	✓	✓
Surface albedo	✓	✓
Aerosol profile	✓	✓
Temperature scaling	✓	
Humidity scaling	✓	
Wavelength offset, scaling	✓	
Fluorescence	✓	

2.2. Optimal estimation. The relationship between the n -dimensional vector of satellite radiances \mathbf{Y} and the r -dimensional state vector \mathbf{X} , where typically $n \gg r$, can be represented through a simple statistical model,

$$(1) \quad \mathbf{Y} = \mathbf{F}(\mathbf{X}, \mathbf{B}) + \boldsymbol{\epsilon}.$$

The random errors $\boldsymbol{\epsilon}$ can represent measurement error along with model discrepancy. Here we assume

$$\boldsymbol{\epsilon} \sim \text{Gaussian}(\mathbf{0}, \boldsymbol{\Sigma}_{\boldsymbol{\epsilon}}).$$

The state vector can also be treated as a random vector with a marginal distribution,

$$\mathbf{X} \sim \text{Gaussian}(\boldsymbol{\mu}_{\mathbf{X}}, \boldsymbol{\Sigma}_{\mathbf{X}}).$$

Inference for the state can be carried out through its conditional (or posterior) distribution given the radiances and relevant parameters $\boldsymbol{\phi} = (\boldsymbol{\mu}_{\mathbf{X}}, \boldsymbol{\Sigma}_{\mathbf{X}}, \boldsymbol{\Sigma}_{\boldsymbol{\epsilon}}, \mathbf{B})$,

$$\begin{aligned} [\mathbf{X}|\mathbf{Y}, \boldsymbol{\phi}] &= \frac{[\mathbf{X}, \mathbf{Y}|\boldsymbol{\phi}]}{\int [\mathbf{X}, \mathbf{Y}|\boldsymbol{\phi}] d\mathbf{X}} \\ &= \frac{[\mathbf{X}|\boldsymbol{\phi}][\mathbf{Y}|\mathbf{X}, \boldsymbol{\phi}]}{\int [\mathbf{X}|\boldsymbol{\phi}][\mathbf{Y}|\mathbf{X}, \boldsymbol{\phi}] d\mathbf{X}}, \end{aligned}$$

where the notation $[\mathbf{A}|\mathbf{B}]$ denotes the conditional probability distribution of \mathbf{A} given \mathbf{B} . The conditional mean $E(\mathbf{X}|\mathbf{Y}, \boldsymbol{\phi})$ can serve as an estimate of the state, and the conditional variance $\text{Var}(\mathbf{X}|\mathbf{Y}, \boldsymbol{\phi})$ can characterize the uncertainty of the estimate. This inference framework is known as OE in the remote sensing literature [22]. OE retrievals for atmospheric constituents such as carbon monoxide, CO₂, and ozone have been implemented for a number of recent Earth-observing satellites [19, 27]. Despite the multivariate Gaussian assumption for the random errors and the atmospheric state, the posterior distribution is not Gaussian if the forward model is nonlinear. Generally, an analytical form for the posterior distribution is unavailable. However, sampling from the posterior distribution is possible with MCMC [24,

12], but can be prohibitively expensive for the number of soundings processed for a mission like OCO-2. Evaluation of the forward model $\mathbf{F}(\mathbf{X}, \mathbf{B})$ is time consuming, so the full posterior distribution must be summarized in an efficient manner that limits the number of evaluations of the forward model.

A strategy commonly advocated in remote sensing and used in the OCO-2 FP retrieval algorithm is to search for the posterior mode. This is equivalent to minimizing a “cost function” of the form,

$$(2) \quad -2 \ln[\mathbf{X}|\mathbf{Y}, \phi] = (\mathbf{Y} - \mathbf{F}(\mathbf{X}, \mathbf{B}))^T \Sigma_e^{-1} (\mathbf{Y} - \mathbf{F}(\mathbf{X}, \mathbf{B})) \\ + (\mathbf{X} - \mu_{\mathbf{X}})^T \Sigma_{\mathbf{X}}^{-1} (\mathbf{X} - \mu_{\mathbf{X}}) + \text{constant}.$$

A variety of optimization algorithms can be used for solving this nonlinear least squares problem. The Levenberg–Marquardt (LM) algorithm, which is a tunable generalization of gradient descent and the Gauss–Newton algorithm, is often used in remote sensing applications [22]. The actual implementation of the algorithm includes nontrivial choices such as the starting value, convergence criterion, and initial value for the LM regularization parameter. The algorithm determines step size and direction, in part, based on the gradient of the cost function (2), which requires the forward-model Jacobian,

$$\mathbf{K}(\mathbf{X}) = \frac{\partial \mathbf{F}(\mathbf{X}, \mathbf{B})}{\partial \mathbf{X}} \equiv \left(\frac{\partial F_i(\mathbf{X}, \mathbf{B})}{\partial X_k} \right).$$

Notice that the Jacobian is generally a function of the atmospheric state.

In an operational setting such as the OCO-2 FP retrieval, other algorithm choices must be made as well. In particular, values for key parameters are set at fixed values. Since their true values are not generally known, we distinguish these retrieval parameters from their true counterparts.

- The retrieval forward-model parameters are set at $\hat{\mathbf{B}}$, and the true forward-model parameters are \mathbf{B} .
- The retrieval prior mean vector is set at μ_a , and the true marginal mean for the state is $\mu_{\mathbf{X}}$.
- The retrieval prior covariance matrix is set at Σ_a , and the true marginal covariance for the state is $\Sigma_{\mathbf{X}}$.
- The retrieval radiance error covariance matrix is set at Σ_e , and the true radiance error covariance is Σ_{ϵ} .

The value of the state vector at the last step of a nominally converged LM algorithm is declared the retrieved state and denoted $\hat{\mathbf{X}}$. It is a function of the data \mathbf{Y} . An expression for the posterior covariance [22] can be computed through a linear approximation,

$$\mathbf{S}(\mathbf{X}) \equiv [\mathbf{K}(\mathbf{X})^T \Sigma_e^{-1} \mathbf{K}(\mathbf{X}) + \Sigma_a^{-1}]^{-1}.$$

This approximation involves the Jacobian, which must be evaluated at a chosen value of the state vector. This choice of \mathbf{X} , or linearization point, can impact the overall uncertainty if, for example, the retrieval $\hat{\mathbf{X}}$ is used as the linearization point. The OCO-2 operational retrieval

uses this convention, so this choice is used throughout the rest of this paper. Henceforth, we define

$$\hat{\mathbf{S}} \equiv \mathbf{S}(\hat{\mathbf{X}}) = \left[\mathbf{K}(\hat{\mathbf{X}})^T \boldsymbol{\Sigma}_e^{-1} \mathbf{K}(\hat{\mathbf{X}}) + \boldsymbol{\Sigma}_a^{-1} \right]^{-1}.$$

The combination of $\hat{\mathbf{X}}$ and $\hat{\mathbf{S}}$ are typically used for inference in operational retrievals that implement OE. More computationally intensive approaches such as randomized likelihood use conditional simulation to characterize the uncertainty in the estimate of the state [15].

The primary QOI for OCO-2 is X_{CO_2} , the column-averaged dry-air mole fraction of CO_2 . Fundamentally, this is the ratio of the number of CO_2 molecules in a column to the total number of molecules of dry air in the column. We decompose the state vector,

$$\mathbf{X} = \begin{bmatrix} \mathbf{X}_\alpha \\ \mathbf{X}_\beta \end{bmatrix},$$

where \mathbf{X}_α is the vertical profile of CO_2 and \mathbf{X}_β is the rest of the state vector. The prior mean vector of the state,

$$\boldsymbol{\mu}_a = \begin{bmatrix} \boldsymbol{\mu}_{a,\alpha} \\ \boldsymbol{\mu}_{a,\beta} \end{bmatrix},$$

can be similarly decomposed, and the covariance matrix can be written as

$$\hat{\mathbf{S}} = \begin{bmatrix} \hat{\mathbf{S}}_{\alpha\alpha} & \hat{\mathbf{S}}_{\alpha\beta} \\ \hat{\mathbf{S}}_{\beta\alpha} & \hat{\mathbf{S}}_{\beta\beta} \end{bmatrix},$$

where $\hat{\mathbf{S}}_{\alpha\alpha}$ is the block of the covariance matrix corresponding to the vertical profile of CO_2 .

Given the configuration of the state vector, X_{CO_2} can be constructed as a weighted average of the vertical profile of CO_2 [19]. The vector of weights $\mathbf{h}(\mathbf{X}_\beta)$ has the same dimension as \mathbf{X}_α , and the weights are generally a function of the other state vector elements. However, the weights are fixed for the surrogate model defined in section 3, and we drop the dependence of \mathbf{h} on the state vector,

$$X_{CO_2} = \mathbf{h}^T \mathbf{X}_\alpha.$$

In a similar fashion, the retrieved X_{CO_2} and a variance estimate can be computed from the retrieval,

$$\begin{aligned} \hat{X}_{CO_2} &\equiv \mathbf{h}^T \hat{\mathbf{X}}_\alpha, \\ \widehat{Var}_{X_{CO_2}} &\equiv \mathbf{h}^T \hat{\mathbf{S}}_{\alpha\alpha} \mathbf{h}. \end{aligned}$$

2.3. Error analysis. Linear error analysis is a standard framework for diagnosing error characteristics in OE retrievals [22]. Through an analytic formulation, the technique quantifies the *linear* propagation of uncertainty for particular sources, including the inherent variability of the state, the noisy measurements, and systematic errors in parameters and the forward model, into the variability in the retrieval errors. In this article, we compare and contrast this approach with simulation-based UQ, which can additionally characterize nonlinearity and uncertainty propagation from any other retrieval algorithm choices, specifically uncertainty in the prior mean, that are not handled in the OE framework. The linear error analysis technique

in OE uses a linearization of the retrieval error, $\Delta = \hat{\mathbf{X}} - \mathbf{X}$, to decompose the contribution from the sources noted previously. The linearization process relies on the Jacobian and two additional operators.

1. The *gain matrix* \mathbf{G} has dimension $r \times n$ and characterizes the linear response of the retrieval to the measurements,

$$\mathbf{G}(\hat{\mathbf{X}}) = \left[\left(\mathbf{K}(\hat{\mathbf{X}}) \right)^T \Sigma_e^{-1} \mathbf{K}(\hat{\mathbf{X}}) + \Sigma_a^{-1} \right]^{-1} \left(\mathbf{K}(\hat{\mathbf{X}}) \right)^T \Sigma_e^{-1}.$$

2. The *averaging kernel* \mathbf{A} has dimension $r \times r$ and characterizes the linear response of the retrieval to the state vector,

$$\mathbf{A}(\hat{\mathbf{X}}) = \mathbf{G}(\hat{\mathbf{X}}) \mathbf{K}(\hat{\mathbf{X}}).$$

In this framework, the retrieval error can be decomposed into several contributions [22],

$$\begin{aligned} \Delta &= \hat{\mathbf{X}} - \mathbf{X} \\ &= \left(\mathbf{A}(\hat{\mathbf{X}}) - \mathbf{I} \right) (\mathbf{X} - \boldsymbol{\mu}_a) && \text{(smoothing)} \\ &\quad + \mathbf{G}(\hat{\mathbf{X}}) \boldsymbol{\epsilon} && \text{(noise)} \\ &\quad + \boldsymbol{\gamma} && \text{(nonlinearity)}. \end{aligned}$$

The nonlinearity term $\boldsymbol{\gamma}$ is zero for a linear forward model, as outlined in Appendix A. Additional contributions arise if the forward model used in the retrieval is different from the true forward function. Parameter error is also introduced if the retrieval model parameters $\hat{\mathbf{B}}$ are different from the true model parameters \mathbf{B} . For an operational retrieval such as OCO-2, these are important contributions to the retrieval error; these other contributions will not be addressed in the current work.

The analogous error budget has been developed for X_{CO2} [5]:

$$\begin{aligned} (3) \quad \Delta_{XCO2} &= \hat{X}_{CO2} - X_{CO2} \\ &= \mathbf{h}^T \left(\mathbf{A}_{\alpha\alpha}(\hat{\mathbf{X}}) - \mathbf{I}_{\alpha\alpha} \right) (\mathbf{X}_\alpha - \boldsymbol{\mu}_{a,\alpha}) && \text{(smoothing)} \\ &\quad + \mathbf{h}^T \mathbf{A}_{\alpha\beta}(\hat{\mathbf{X}}) (\mathbf{X}_\beta - \boldsymbol{\mu}_{a,\beta}) && \text{(interference)} \\ &\quad + \mathbf{h}^T \mathbf{G}_\alpha(\hat{\mathbf{X}}) \boldsymbol{\epsilon} && \text{(noise)} \\ &\quad + \gamma_{XCO2} && \text{(nonlinearity)}. \end{aligned}$$

Here, the averaging kernel matrix is partitioned in a similar fashion to the covariance matrix with $\mathbf{A}_{\alpha\alpha}(\hat{\mathbf{X}})$ and $\mathbf{A}_{\alpha\beta}(\hat{\mathbf{X}})$ representing the CO_2 -profile rows of the averaging kernel. Further, $\mathbf{G}_\alpha(\hat{\mathbf{X}})$ represents the first 20 rows, corresponding to the CO_2 profile, of the gain matrix.

In this budget, the smoothing error for the full state vector is further divided for X_{CO2} into smoothing error for the CO_2 profile and interference error due to the correlation between retrieval errors in the CO_2 profile and retrieval errors in other state-vector elements [5, 23]. The final term, γ_{XCO2} , is a catch-all that arises from the nonlinearity of the forward model,

the role of this nonlinearity in the behavior of the retrieval algorithm, and the choice of linearization point. In an operational setting, the true state \mathbf{X} and random error ϵ are not known, so the OE error analysis focuses on characterizing the plausible marginal variability of each contributor to the budget based on the assumed probability distribution of the true state and random error [5]. Correlations between contributions are ignored. Through our simulation experiment (section 4), components of the error budget can be computed directly from the known true state and model discrepancy. Error budget components can be evaluated jointly.

3. Surrogate model. The previous section highlighted some of the critical choices in the practical implementation of the OCO-2 remote sensing retrieval. Parameters that are in reality uncertain are fixed, and the LM algorithm is configured in a specified fashion. These choices can impact the distribution of the retrieval $\hat{\mathbf{X}}$ and the adequacy of $\hat{\mathbf{S}}$ as a measure of the variability of the distribution of the retrieval error,

$$\Delta = \hat{\mathbf{X}} - \mathbf{X}.$$

Particular attention is focused on the retrieval error for X_{CO_2} , namely,

$$\Delta_{X_{CO_2}} = \hat{X}_{CO_2} - X_{CO_2}.$$

We wish to study this distribution by simulation experiments through extensive Monte Carlo draws under different combinations of geophysical conditions and algorithm choices. However, the computational cost of the OCO-2 FP forward model limits the scope of any experiments involving this model.

Consequently, we have developed a computationally efficient surrogate model and retrieval based on the physical principles in the OCO-2 FP forward model and measurement approach. There are multiple strategies for surrogate-model development in the literature. Statistical models, which are usually Gaussian process models, are often developed as emulators of complex computer models [17, 20]. Another approach involves developing a surrogate of reduced order or complexity based on the original parent model, which is the approach implemented in this article. The surrogate model makes some simplifications for interpretability and computational efficiency while attempting to maintain the key components of the state vector and RT that contribute substantially to uncertainty in X_{CO_2} . Scattering of radiation in the atmosphere by aerosols has been shown to contribute to errors in retrieved X_{CO_2} for other remote sensing instruments [19], so aerosols are a primary focus for investigation with the surrogate model. After some initial investigation with even simpler surrogate models that did not include aerosol scattering, we found that the surrogate model presented here exhibits a satisfactory level of nonlinear behavior for the experiments desired. As implemented, the surrogate model achieves computational efficiency over the FP model through a reduced state vector, fixed absorption coefficients, a simplified instrument model, and reduced-accuracy numerics for RT. Further details on the surrogate model can be found in Appendix B.

The surrogate-model state vector includes the same configuration as the FP state vector for the CO_2 profile, surface pressure, surface albedo, and aerosols, as defined in section 2.1. Other elements of the FP state vector are not included in the surrogate state vector. Table 1

highlights the makeup of the two models' state vectors. In their most extensive formulation, the surrogate state vector includes 39 elements and the FP state vector includes 49 elements. A more detailed description of the representation of the state vector and the RT included in the surrogate model can be found in Appendix B.

Evaluation of the surrogate forward model provides a substantial computational speedup; a five-iteration retrieval takes approximately 200 seconds with the FP model and approximately 10 seconds for the surrogate model. This speed improvement allows extensive Monte Carlo experiments with the surrogate model. Figure 2 provides an overview of the general experimental setup. An experiment requires specification of the true marginal distribution for the state \mathbf{X} , through $(\boldsymbol{\mu}_{\mathbf{X}}, \boldsymbol{\Sigma}_{\mathbf{X}})$, the random error characteristics through $\boldsymbol{\Sigma}_{\epsilon}$, and the forward model parameters \mathbf{B} . Similar choices are made for the surrogate retrieval inputs such as $\boldsymbol{\mu}_a, \boldsymbol{\Sigma}_a, \boldsymbol{\Sigma}_e, \hat{\mathbf{B}}$. We distinguish two key approaches for choosing these inputs. One option is to fix these inputs at specified values, which we call the *sensitivity mode*. Another option, as illustrated in Figure 2, is to generate random inputs to reflect uncertainty in retrieval inputs. This option is termed *stochastic mode*.

The experiment proceeds by simulating a large random sample of state vectors \mathbf{X} , each of which is used to evaluate the forward model. Random errors are added to yield synthetic radiance vectors \mathbf{Y} . A surrogate retrieval is then performed to yield retrievals $\hat{\mathbf{X}}$ and covariances $\hat{\mathbf{S}}$.

4. UQ simulation experiment. In this section we develop a surrogate-model experiment to investigate the impact of systematic misspecification of and uncertainty in the retrieval prior mean $\boldsymbol{\mu}_a$ on the retrieval error distribution. These experiments focus on the impact of the prior mean choices for surface albedo and aerosols. Representing the surface and aerosols is an ongoing challenge in remote sensing retrievals like OCO-2, since they appear to contribute a substantial portion of the variability in retrieval errors [19].

4.1. Marginal distribution. The geophysical states are constructed from available data sources, which include remote sensing and reanalysis datasets. These sources provide geophysically plausible mean states and intraseasonal variability, which is adequate for studying the error distribution under a range of geophysical conditions and algorithm choices. The experiment considers a marginal distribution based on typical conditions near Izaña, Tenerife, Spain in July. Influenced by atmospheric transport from northern Africa, this location is characterized by moderate CO_2 variability and high mean AOD, particularly from dust.

A few key data sources provide the basis for the marginal distribution. In each case, daily “data” from June–August 2013 near the location of interest are extracted. Daily values for the necessary components of the state vector are treated as replicates, and their empirical means and covariances are assembled to produce a marginal mean vector $\boldsymbol{\mu}_{\mathbf{X}}$ and a marginal covariance matrix $\boldsymbol{\Sigma}_{\mathbf{X}}$. Daily data on vertical profiles for CO_2 come from a simulation of NASA’s Goddard Earth Observing System model, version 5 (GEOS-5) [21]. Daily data on surface pressure and aerosols come from the modern era retrospective analysis for research and applications aerosol reanalysis (MERRAero) [2]. Finally, daily data on surface albedo data come from the Moderate Resolution Imaging Spectrometer (MODIS) albedo product [26].

4.2. Simulation of the radiances. The (marginal) distribution of \mathbf{X} , with mean $\boldsymbol{\mu}_{\mathbf{X}}$ and covariance matrix $\boldsymbol{\Sigma}_{\mathbf{X}}$, is used to simulate synthetic state vectors. For each simulated state \mathbf{X} , the surrogate model $\mathbf{F}(\mathbf{X}, \mathbf{B})$ is evaluated at each wavelength in each band, and random errors $\boldsymbol{\epsilon}$ are added to yield synthetic radiance vectors \mathbf{Y} . The error covariance matrix $\boldsymbol{\Sigma}_{\boldsymbol{\epsilon}}$ is a diagonal matrix. The individual variances are defined to be proportional to the expected signal. Specifically, let $\mathbf{Y} \equiv \{Y_{i,j} : i = 1, \dots, n_j; j = 1, 2, 3\}$, where j indexes the spectral band (O_2 , weak CO_2 , strong CO_2) and i indexes wavelength within a band. Hence, $n_1 + n_2 + n_3 = n$. Then the variance for each radiance $Y_{i,j}$ is related to its expectation, as follows:

$$\begin{aligned} Y_{i,j} &= F_{i,j}(\mathbf{X}, \mathbf{B}) + \epsilon_{i,j}, \\ \text{Var}(Y_{i,j}) &= c_j F_{i,j}(\mathbf{X}, \mathbf{B}). \end{aligned}$$

The band-specific constant c_j is specified to yield signal-to-noise ratios that are comparable to those characteristic of the OCO-2 instrument. This model for the error variance follows the general behavior of the instrument with a slightly simplified structure. The OCO-2 operational algorithm develops wavelength-specific variances based on known instrument characteristics [11]. These distributional assumptions for generating synthetic states \mathbf{X} and radiances \mathbf{Y} are applied for all treatments in the experiment.

4.3. Treatments in the simulation experiment. The experiment explores the impact of uncertainty in the retrieval prior mean $\boldsymbol{\mu}_a$, as depicted on the right side of Figure 2; the prior covariance $\boldsymbol{\Sigma}_a$ is fixed at $\boldsymbol{\Sigma}_{\mathbf{X}}$. In particular, each retrieval uses a prior mean that is generated from a hyperdistribution,

$$\boldsymbol{\mu}_a \sim \text{Gaussian}(\boldsymbol{\theta}_a, \boldsymbol{\Omega}_a).$$

The experiment includes two factors with levels that reflect different choices for the hyperparameters $\boldsymbol{\theta}_a$ and $\boldsymbol{\Omega}_a$. The two factors described below included five and three levels, respectively, and the experiment was run in a full two-way factorial design, yielding 15 treatments.

The first factor is the systematic error present in the prior mean $\boldsymbol{\mu}_a$, reflected by the choice of the hyperparameter $\boldsymbol{\theta}_a$. In general, this parameter is defined as an offset from the true marginal mean,

$$\boldsymbol{\theta}_a = \boldsymbol{\mu}_{\mathbf{X}} + \boldsymbol{\delta}.$$

The five levels of this factor reflect varying amounts of misspecification:

- MA: $\boldsymbol{\delta} = -2\sqrt{\text{diag}(\boldsymbol{\Sigma}_{\mathbf{X}})}$,
- MB: $\boldsymbol{\delta} = -\sqrt{\text{diag}(\boldsymbol{\Sigma}_{\mathbf{X}})}$,
- MC: $\boldsymbol{\delta} = \mathbf{0}$,
- MD: $\boldsymbol{\delta} = \sqrt{\text{diag}(\boldsymbol{\Sigma}_{\mathbf{X}})}$,
- ME: $\boldsymbol{\delta} = 2\sqrt{\text{diag}(\boldsymbol{\Sigma}_{\mathbf{X}})}$.

Here, $\sqrt{\text{diag}(\boldsymbol{\Sigma}_{\mathbf{X}})}$ represents a vector with a single nonzero element given by the marginal standard deviation for the natural logarithm of the AOD (log AOD) for the dominant aerosol type, which is dust for the location of interest. The element is in its appropriate place in the state vector, and all other elements are set to 0 for all treatments. We know from the physics behind the retrieval and preliminary surrogate-model experiments that uncertainty in the AOD component of the prior mean is among the most problematic.

Table 2

Treatments for the uncertain prior mean (μ_a) experiment. Each treatment is named as a combination of the magnitude of systematic error (MA, MB, MC, MD, ME) in the prior mean and the level of uncertainty (V0, V1, V2) in the prior mean.

		Covariance Ω_a		
		0	$\frac{1}{10}\Sigma_X$	$\text{diag}(\Sigma_X)$
Mean offset δ	$-2\sqrt{\text{diag}(\Sigma_X)}$	MAV0	MAV1	MAV2
	$-\sqrt{\text{diag}(\Sigma_X)}$	MBV0	MBV1	MBV2
	0	MCV0	MCV1	MCV2
	$\sqrt{\text{diag}(\Sigma_X)}$	MDV0	MDV1	MDV2
	$2\sqrt{\text{diag}(\Sigma_X)}$	MEV0	MEV1	MEV2

The second factor is the degree of uncertainty present in the specification of the prior mean, reflected by the choice of the hyperparameter Ω_a . The three levels of this factor reflect no uncertainty, small uncertainty, and moderate uncertainty, respectively:

- V0: $\Omega_a = 0$,
- V1: $\frac{1}{10}\Sigma_X$,
- V2: $\Omega_a = \text{diag}(\Sigma_X)$.

The treatments are summarized in Table 2.

For the treatments that include some degree of uncertainty in the retrieval's prior mean μ_a , it is possible to estimate components of the variance in X_{CO_2} through the use of the conditional-variance formula,

$$\text{Var}(\Delta_{XCO_2}) = E(\text{Var}(\Delta_{XCO_2}|\mu_a)) + \text{Var}(E(\Delta_{XCO_2}|\mu_a)).$$

The first contribution, $E(\text{Var}(\Delta_{XCO_2}|\mu_a))$, is the variability in the retrieval errors given the prior mean, averaged across the distribution of prior means. This variability results from the inherent variability in the state \mathbf{X} as well as the random errors in the radiances \mathbf{Y} , and the posterior covariance $\hat{\mathbf{S}}$ accounts for these, at least to the extent that the linear approximation is adequate. The second contribution, $\text{Var}(E(\Delta_{XCO_2}|\mu_a))$, is variability in the retrieval bias for a given prior mean across the distribution of prior means. The posterior covariance $\hat{\mathbf{S}}$ conditions on the prior mean μ_a and does not capture this second contribution to the variability. These components can both be computed in the Monte Carlo framework if a hierarchical sampling strategy is used. Specifically,

- Generate $p = 1, \dots, 50$ random prior mean vectors

$$\mu_{a,p} \sim N(\theta_a, \Omega_a).$$

- For each prior mean vector $\mu_{a,p}$, generate $q = 1, \dots, 400$ simulated states and radiances $\mathbf{X}_{p,q}$, $\mathbf{Y}_{p,q}$ and perform retrievals.

The sample size of 400 for each prior mean represents a compromise that achieves a satisfactory Monte Carlo precision while allowing a reasonable outer loop sample size (50). The treatments representing no uncertainty in the prior mean (V0) do not require hierarchical sampling. For these treatments, a total of 5000 independent state and radiance vectors were simulated.

5. Results. This section summarizes the results of the experiment in several ways. Since X_{CO_2} is the primary QOI, it receives additional focus, both in terms of the components of

variance relative to variability in the retrieval prior mean and in terms of the components of the error budget. In addition, the bias and covariance of the retrieval errors for the full state vector \mathbf{X} are summarized using a small set of summary figures of merit. These diagnostics reveal key properties of the CO_2 retrieval and represent a suite of tools that could additionally be used in summarizing simulation experiments for other remote sensing retrievals and similar nonlinear Bayesian inverse problems.

5.1. X_{CO_2} components of variance. Figure 4 summarizes the error distributions for X_{CO_2} for each of the treatments in the experiment. The error distribution for each prior mean μ_a , which is fixed for the V0 treatments (left column) and randomly generated (center and right columns), is summarized with its mean and two extreme quantiles. The impact of the increasing level of uncertainty in the retrieval prior mean is evident both in the V1 treatments, where a modest amount of additional variability is present in the overall error distribution, and in the V2 treatments, where there is especially noticeable variability in the conditional means (points) of the X_{CO_2} errors for the randomly selected prior means. In addition, there is a weak relationship between this conditional bias and the prior mean log AOD, which is particularly evident in the MAV2 and MEV2 treatments. As the log AOD prior mean increases, the mean X_{CO_2} retrieval error decreases. This relationship clearly does not explain all of the variability in the conditional bias, so other elements of the prior mean vector play a role as well.

Table 3 summarizes the bias and variance in the X_{CO_2} retrieval error for each treatment in the experiment. For the V1 and V2 treatments, the variance is separated into the contributions from the average error variance within each prior mean $E(\text{Var}(\Delta_{X_{\text{CO}_2}}|\mu_a))$ and from the variance of average errors across prior means $\text{Var}(E(\Delta_{X_{\text{CO}_2}}|\mu_a))$. In addition, the average of the estimated posterior variances $E(\widehat{\text{Var}}_{X_{\text{CO}_2}})$, is reported for comparison.

From a practical standpoint, the retrieval bias is small (less than 0.1 ppm) for all except the extreme MA and ME treatments. There is a trend from negative to positive bias moving from MA to ME. This suggests that the prior-mean specification may reflect the importance of nonlinearity in the presence of parameter error, a topic that is studied further in section 5.2. The volatility is also reflected in the variance of the retrieval errors. Both components of the error variance are largest for the MAV2 and MEV2 treatments. The between-prior variance is largest for the V2 treatments and is relatively modest in the V1 treatments.

The average of the estimated posterior variances, $E(\widehat{\text{Var}}_{X_{\text{CO}_2}})$, compares well to the empirical error variance computed from the Monte Carlo simulations for the V0 treatments, although the empirical error variance is at least slightly larger for every treatment. The posterior variance attempts to capture the inherent variability in the atmospheric state and the noise present in the radiances, and the inflation in the V0 treatments may be due in part to nonlinearity. In addition, the posterior-variance calculation assumes a fixed (known) prior mean μ_a , so the V1 and V2 treatments will exhibit additional variability in the retrieval errors that would not be captured in the calculation of $\widehat{\text{Var}}_{X_{\text{CO}_2}}$. This mismatch is noticeable, around 20%, in the small-uncertainty (V1) treatments and becomes more substantial, as large as 50%, for the moderate-uncertainty (V2) treatments. This result underscores the impact of uncertainty propagation for a particular algorithm input, μ_a , through uncertainty in the primary QOI.

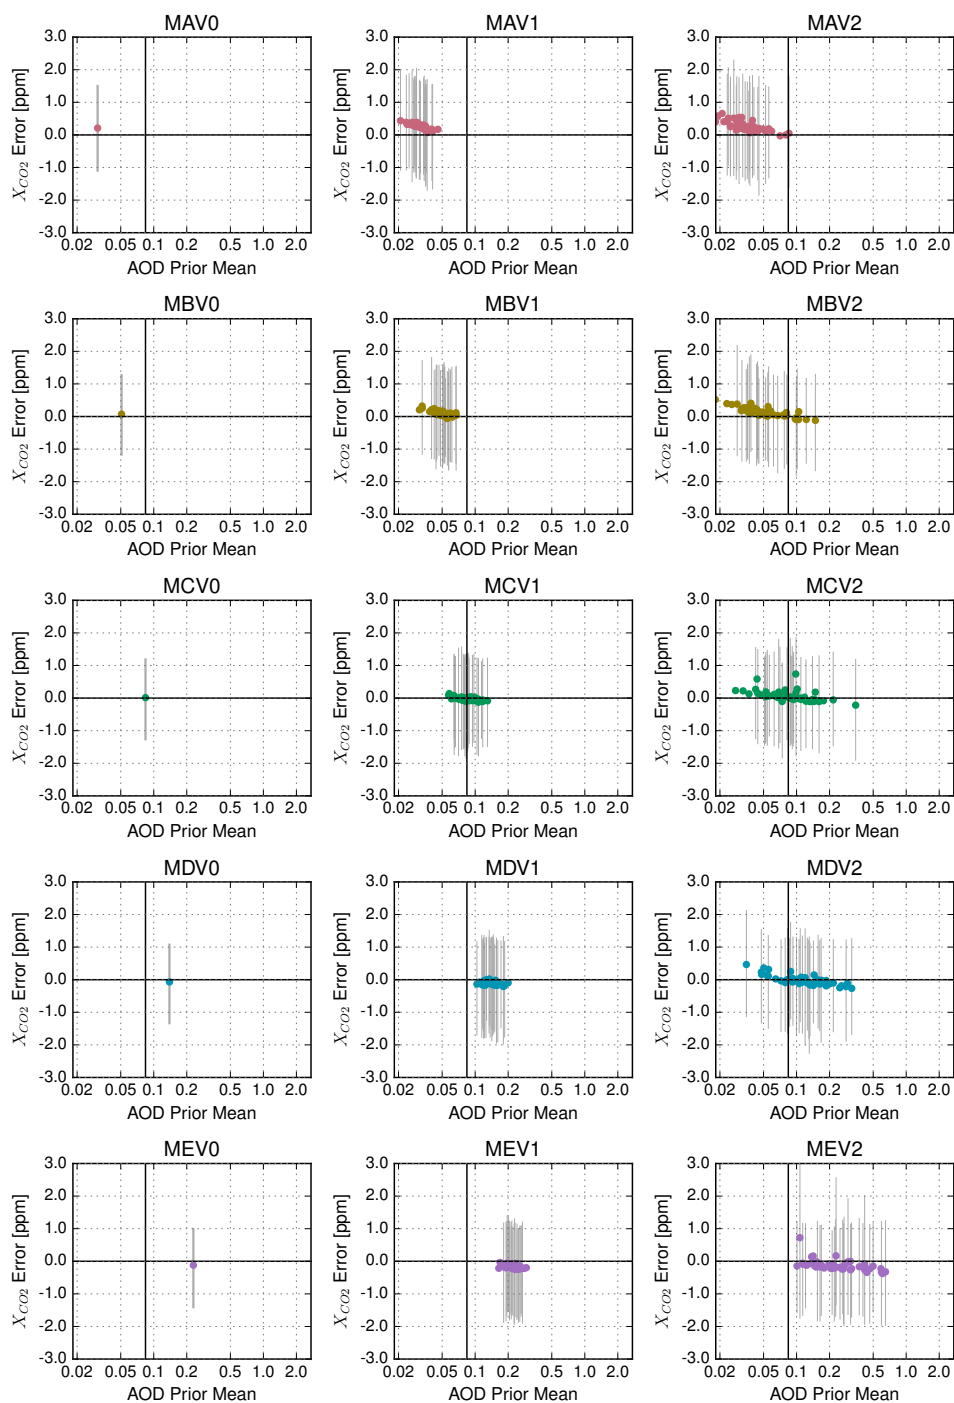


Figure 4. Distribution of retrieval errors for X_{CO_2} , under the hierarchical sampling strategy and different experiment conditions, plotted against the log AOD component of the prior mean. The solid vertical line depicts the true marginal mean of log AOD. Solid circles depict the distribution's mean and error bars cover the center 95% of the retrieval-error distribution.

Table 3

Summary of X_{CO_2} bias and variance for the uncertain prior mean experiment. Bias is reported in units of ppm and variance is reported in units of ppm^2 . The total variance of the retrieval errors is $\text{Var}(\Delta_{XCO_2}) = E(\text{Var}(\Delta_{XCO_2}|\mu_a)) + \text{Var}(E(\Delta_{XCO_2}|\mu_a))$, which is the sum of the two components above it in the table. This total can be contrasted with the retrieval's mean estimated variance $E(\widehat{\text{Var}}_{XCO_2})$.

$E(\Delta_{XCO_2})$	MAV0	MAV1	MAV2
$E(\text{Var}(\Delta_{XCO_2} \mu_a))$	0.210	0.264	0.312
$\text{Var}(E(\Delta_{XCO_2} \mu_a))$	0.436	0.588	0.663
$\text{Var}(\Delta_{XCO_2})$	0.006	0.006	0.034
$E(\widehat{\text{Var}}_{XCO_2})$	0.436	0.594	0.697
	0.344	0.482	0.483
$E(\Delta_{XCO_2})$	MBV0	MBV1	MBV2
$E(\text{Var}(\Delta_{XCO_2} \mu_a))$	0.073	0.097	0.144
$\text{Var}(E(\Delta_{XCO_2} \mu_a))$	0.382	0.553	0.588
$\text{Var}(\Delta_{XCO_2})$	0.006	0.006	0.022
$E(\widehat{\text{Var}}_{XCO_2})$	0.382	0.559	0.610
	0.331	0.466	0.471
$E(\Delta_{XCO_2})$	MCV0	MCV1	MCV2
$E(\text{Var}(\Delta_{XCO_2} \mu_a))$	0.015	-0.023	0.067
$\text{Var}(E(\Delta_{XCO_2} \mu_a))$	0.388	0.545	0.661
$\text{Var}(\Delta_{XCO_2})$	0.003	0.003	0.027
$E(\widehat{\text{Var}}_{XCO_2})$	0.388	0.548	0.688
	0.324	0.456	0.461
$E(\Delta_{XCO_2})$	MDV0	MDV1	MDV2
$E(\text{Var}(\Delta_{XCO_2} \mu_a))$	-0.069	-0.110	-0.021
$\text{Var}(E(\Delta_{XCO_2} \mu_a))$	0.386	0.543	0.582
$\text{Var}(\Delta_{XCO_2})$	0.003	0.003	0.023
$E(\widehat{\text{Var}}_{XCO_2})$	0.386	0.546	0.605
	0.318	0.444	0.456
$E(\Delta_{XCO_2})$	MEV0	MEV1	MEV2
$E(\text{Var}(\Delta_{XCO_2} \mu_a))$	-0.120	-0.166	-0.127
$\text{Var}(E(\Delta_{XCO_2} \mu_a))$	0.371	0.533	0.658
$\text{Var}(\Delta_{XCO_2})$	0.003	0.003	0.027
$E(\widehat{\text{Var}}_{XCO_2})$	0.371	0.536	0.685
	0.313	0.438	0.437

Since each retrieval, \hat{X}_{CO_2} , has a corresponding reported variance, $\widehat{\text{Var}}_{XCO_2}$, the distribution of retrieval errors can also be diagnosed by normalizing the retrieval error by the square root of this reported variance. The distribution of this unitless quantity,

$$Z_{p,q} = \frac{\Delta_{XCO_2,p,q}}{\sqrt{\widehat{\text{Var}}_{XCO_2,p,q}}}, \quad p = 1, \dots, 50, \quad q = 1, \dots, 400,$$

is summarized in Figure 5 for each treatment in the experiment. The standardized errors $\{Z_{p,q}\}$ are sorted and plotted against standard Gaussian quantiles, yielding a quantile-quantile plot. The slope of the resulting regression line yields a scaling of the standard deviation of the true retrieval errors relative to $\sqrt{\widehat{\text{Var}}_{XCO_2,p,q}}$, which is based on the linear approximation. This slope is closest to unity for the V0 and V1 treatments but deviates more substantially

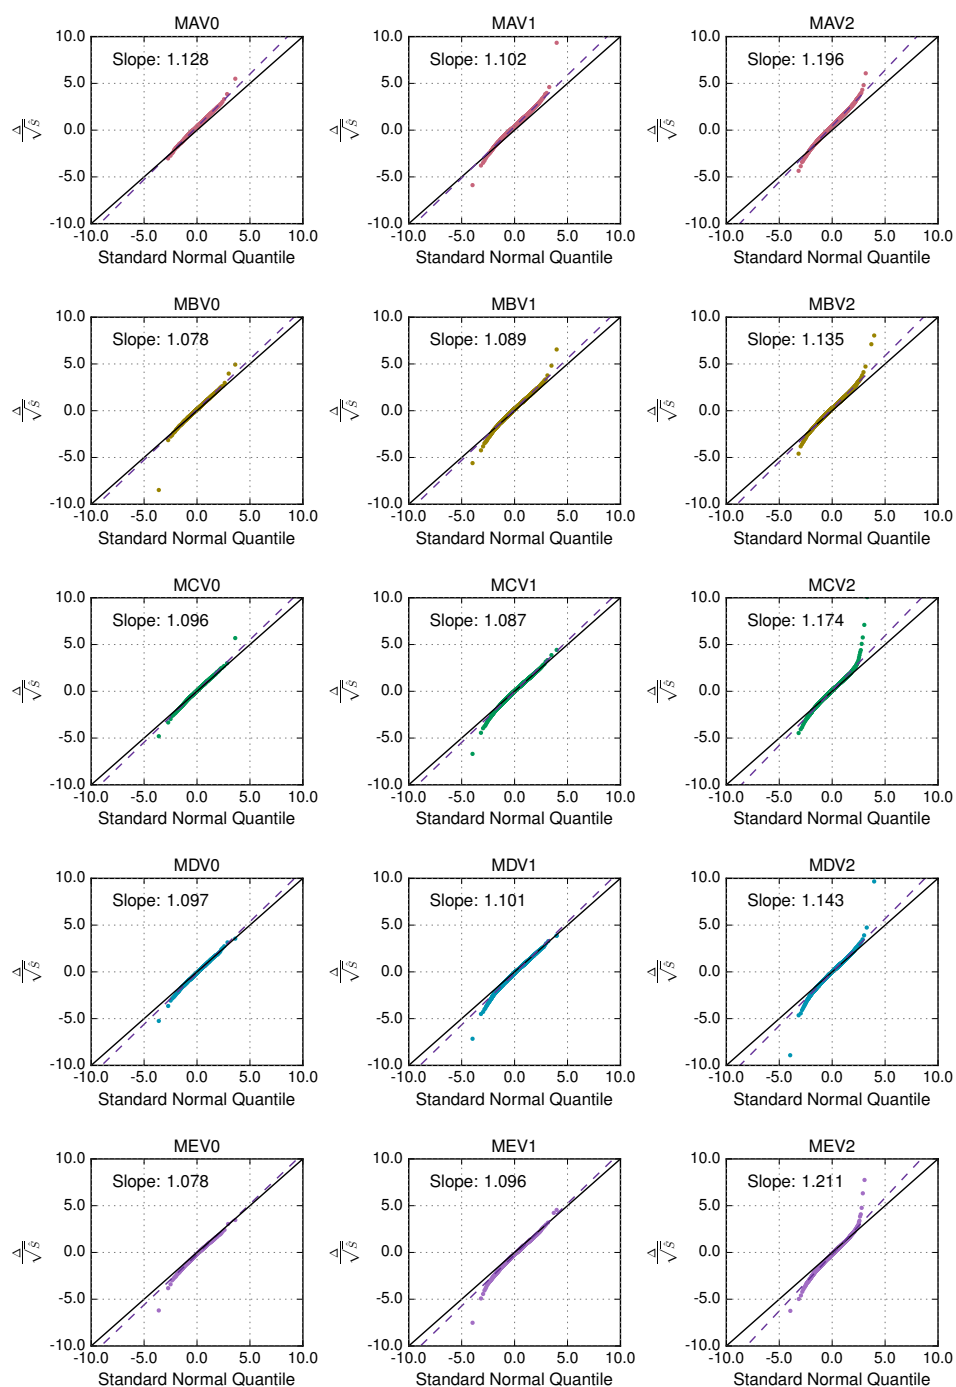


Figure 5. Distribution of normalized X_{CO_2} retrieval errors under different experimental conditions, plotted against quantiles from a standard normal distribution.

in the V2 treatments. In particular, the V2 treatments show a tendency toward skewed and heavy-tailed error distributions.

5.2. X_{CO_2} error budget. Section 2.3 outlined an error budget (3) that is often used in diagnosing remote sensing retrievals. Three of the four error terms, namely, smoothing, interference, and noise, can be computed directly for each Monte Carlo draw and corresponding retrieval. Since the total X_{CO_2} error is available as well, the error due to nonlinearity can be computed as a difference between the total and the sum of the other three components. The joint distribution of the error terms can be summarized from these calculated errors across the Monte Carlo simulation. In addition, an estimate of the variance for each of the first three components can be obtained based on a linear approximation and assumed covariance matrices Σ_e and Σ_a . The calculation based on a linear approximation is often called linear “error analysis” in the remote sensing literature [23, 5], and in our experiment we have an opportunity to assess the validity of linear error analysis.

Figure 6 compares the standard deviation of each error component for each treatment, using both the actual errors based on the simulation and the standard deviations computed based on the linear approximation. The variability in the smoothing error and noise error are nearly constant across all treatments, and the simulation-based variability matches that expected from the linear approximation for both smoothing and noise. These two error components reflect variability due to Σ_e and the CO_2 portion of Σ_X , parameters that are not changed across the treatments.

In contrast, the variability of the interference error and the nonlinear error change across treatments. The error budget suggests that different retrieval prior means μ_a will likely lead to different distributions of interference error. The average interference error is related to the difference between the marginal mean μ_X and the retrieval prior mean μ_a for the pressure, aerosol, and albedo components of the state vector. These are the constituents of X_β in the interference term of the error budget (3). Thus the variability in the retrieval prior mean translates to variability in the average interference error. This variability is not present in the calculation based on the linear approximation, where a fixed retrieval prior mean is assumed. The nonlinear error is a component that is difficult to diagnose in operational linear error analysis, but it is available in this Monte Carlo setting. The nonlinear error term can dominate for the treatments with greater uncertainty. As shown in Figure 6, the variability due to nonlinearity is the largest of the error budget terms in the V2 treatments.

Figure 6 also shows the standard deviation of the total error in X_{CO_2} for both the simulation and the linear approximation. The simulation-based standard deviations are computed from the true retrieval errors in the experiment. For the linear approximation, the standard deviation is $\sqrt{E(\widehat{Var}_{X_{CO_2}})}$. The impact of both the nonlinearity and interference error contributions is evident in the simulation-based variability of the total error, especially for the V2 treatments. The bottom panel of Figure 6 shows that the traditional error analysis always yields total variances that are too small, sometimes substantially so.

The total error variance can also be impacted by correlations among the error budget components. Table 4 summarizes these empirical correlations among the terms in the error budget in the MCV0 (control) and MEV2 treatments. This analysis of the correlations among the components of the error budget is possible in the simulation-based setting, but correlations

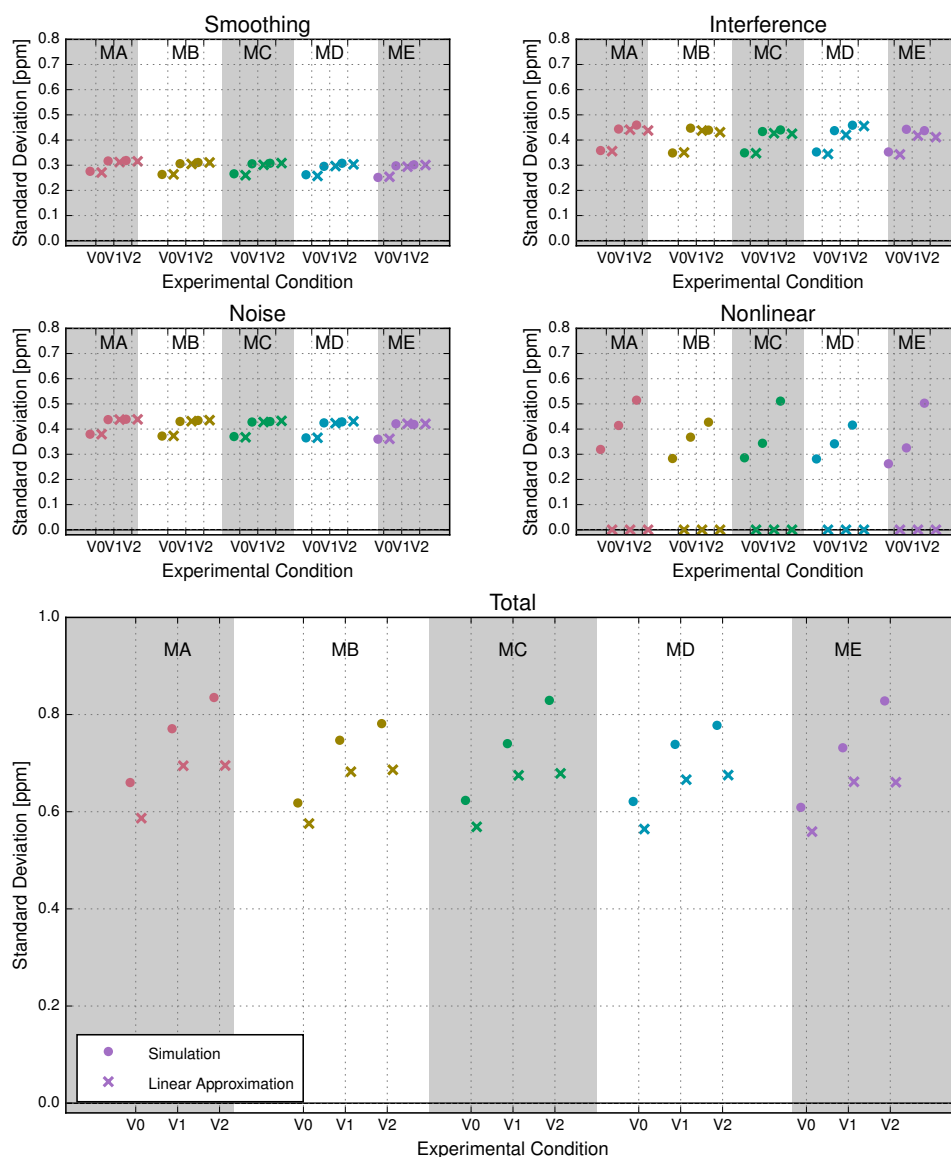


Figure 6. Standard deviation of error-budget components and of the total error for each treatment in the experiment. The four upper panels depict an individual component of the error budget, with the bottom panel depicting the total error. The total error is computed as the standard deviation of the true retrieval errors for the simulation case. For the linear approximation, the total error standard deviation is computed as $\sqrt{E(\widehat{\text{Var}}_{X\text{CO}_2})}$. The treatments are represented in sequence on the horizontal axis. Solid circle (●) symbols represent standard deviations computed from the simulated errors, and × symbols represent standard deviations based on OE's linear approximation.

are not given in traditional linear error analysis. This represents a potential weakness since the variance of the total error is the sum of the variances of individual terms plus twice the sum of covariances between all possible error pairs. Traditional error analysis assumes that

Table 4

Correlations of error-budget components for the MCV0 (control) and MEV2 treatments in the simulation experiment.

	MCV0			
	Smoothing	Interference	Noise	Nonlinear
Smoothing	1.000	−0.039	−0.013	0.075
Interference	−0.039	1.000	0.001	0.081
Noise	−0.013	0.001	1.000	−0.191
Nonlinear	0.075	0.081	−0.191	1.000

	MEV2			
	Smoothing	Interference	Noise	Nonlinear
Smoothing	1.000	−0.011	0.017	0.027
Interference	−0.011	1.000	−0.033	0.043
Noise	0.017	−0.033	1.000	−0.089
Nonlinear	0.027	0.043	−0.089	1.000

the latter component is zero. From Table 4, smoothing, interference and noise errors are essentially uncorrelated with each other. In general, smoothing and interference errors could be correlated with each other if the marginal distribution includes cross-correlations between the CO₂ profile and other components, such as aerosols. The marginal distribution used in this surrogate model experiment does not include correlations between the CO₂ and non-CO₂ components of the state vector. The nonlinear term has modest correlations with the other terms in the control experiment, and the correlation remains, particularly with noise error, in the MEV2 treatment.

5.3. State vector figures of merit (FOMs). An assessment of the error distribution of the full state vector provides additional insight into the behavior of the retrieval algorithm. In particular, a component-by-component look at the retrieval bias and variance can reveal specific state-vector elements that may be more or less problematic in the retrieval. This can be complemented with an investigation of the correlations of retrieval errors across components. Strong correlations, either positive or negative, can suggest combinations of state vector elements that may not be completely identifiable in the retrieval. The Monte Carlo experiment provides the distribution of retrieval errors, $\Delta = \hat{\mathbf{X}} - \mathbf{X}$, and this distribution can be summarized with some key FOMs useful in simultaneous inference [6]. Following the notation of Cressie and Burden [6], we define the retrieval bias and covariance as

$$\begin{aligned} Bias &\equiv E(\hat{\mathbf{X}} - \mathbf{X}) = E(\Delta), \\ Cov &\equiv Cov(\hat{\mathbf{X}} - \mathbf{X}) = Cov(\Delta). \end{aligned}$$

One useful FOM is a unitless normalized bias, or inverse coefficient of variation,

$$Icv = (\text{diag}(Cov))^{-1/2} Bias.$$

Figure 7 illustrates this FOM for the experiment. The behavior of Icv shows some interesting contrasts between \mathbf{X}_α , the CO₂ profile, and \mathbf{X}_β , the other elements of the state vector. In general, larger biases are present for the components \mathbf{X}_β . Some of these errors can compensate

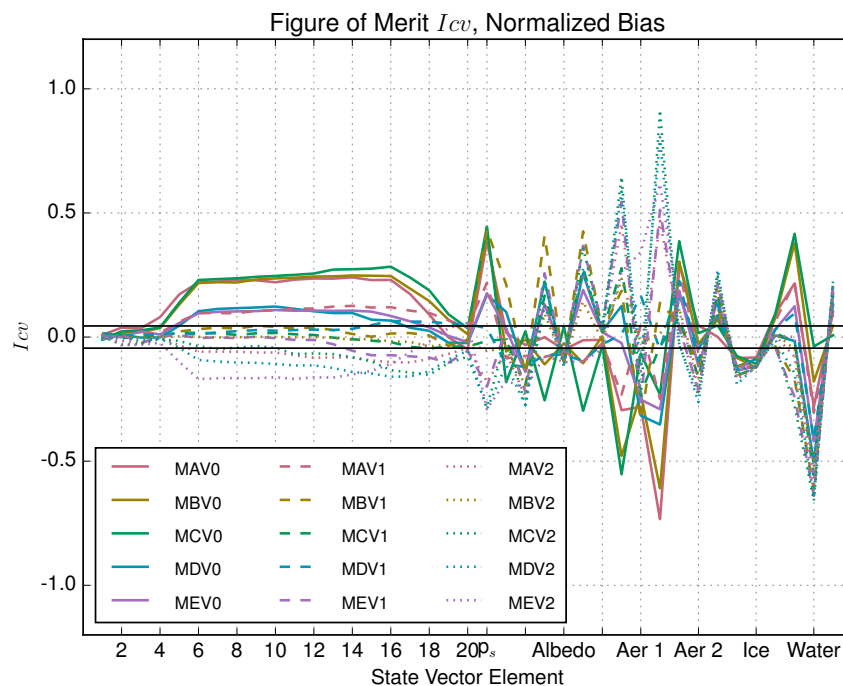


Figure 7. Summary of normalized bias, I_{cv} , for the uncertain inputs experiment.

for each other to an extent; for example, an error in retrieved aerosol can offset an error in retrieved albedo without a substantial impact on CO_2 . Large bias is particularly evident for surface pressure, the band-specific albedo, and the log AOD components for the V1 and V2 treatments. While the V2 treatments have large absolute bias, the variability is most extreme for these treatments as well. The largest errors in \mathbf{X}_α tend to occur in the middle to lower atmosphere, where the CO_2 variability is largest.

Additionally, the correlation matrix of the retrieval errors can provide insight into the relationships among the state vector elements. Figure 8 depicts this matrix for the MCV0 (control) experiment. The upper left 20×20 block represents the correlations among the retrieval errors for the vertical profile of CO_2 . Error correlations for nearby vertical positions are generally positively correlated. The components of the CO_2 profile exhibit modest correlations with other elements of the state vector. The strongest negative correlations exist between the albedo and aerosol components of the state vector. This is an illustration of one of the fundamental challenges for the OCO-2 measurements; surface albedo and aerosol scattering near the surface can give rise to similar spectral signatures. The negative correlation is consistent with the retrieval attempting a trade-off between these contributions.

6. Discussion and conclusion. This study has developed and illustrated a practical framework for quantifying uncertainty in remote sensing retrievals. The combination of a computationally efficient surrogate model and a Monte Carlo framework allows simulation from the retrieval-error distribution under a variety of conditions. These empirical results can be readily compared with OE error analysis based on a linearity assumption. The simulation-based

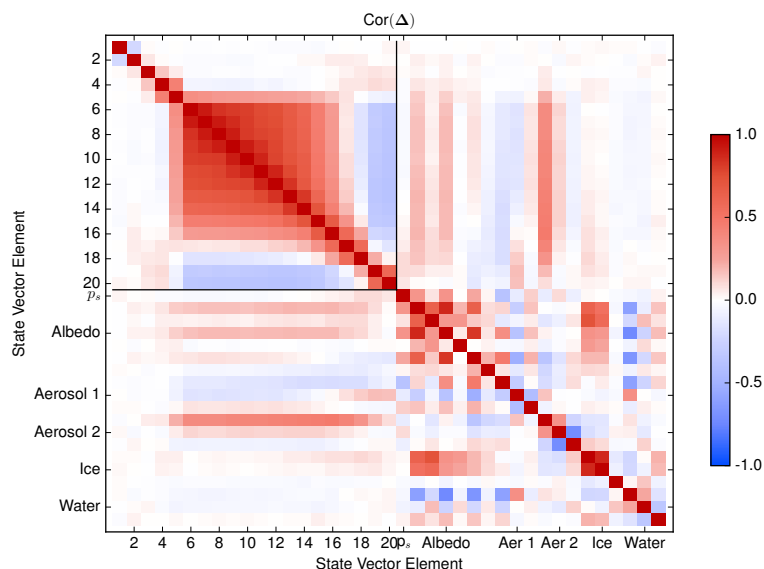


Figure 8. Correlation matrix of retrieval errors, $\Delta = \hat{\mathbf{X}} - \mathbf{X}$, for the MCV0 (control) experiment.

assessment in this study provides a number of insights beyond those obtained from the OE linear error analysis. First, the variability in the error due to nonlinearity can be diagnosed, and it is seen to change across the treatments in the experiment. Second, the simulation reveals that uncertainty in the prior mean μ_a results in a larger interference-error variance than that computed in the linear approximation. Finally, modest correlations among the error budget components are found using the simulation results, which lead to covariances that must be incorporated to achieve an accurate measure of total error.

In addition, the simulation approach provides an overall quantification of the adequacy of the retrieval's uncertainty estimate, and it can also characterize the variability in retrieval errors due to nonlinearity. We find that the combination of systematic misspecification of, and uncertainty in, the prior mean for aerosols and albedo impact the retrieval bias and variance for X_{CO_2} . There is an important interaction between these two factors that leads to large bias and variability when the prior mean of log AOD is high.

The impact of uncertain retrieval-algorithm inputs in general has implications for the community of OCO-2 data-product users. The operational retrieval algorithm reports the approximate posterior variance for X_{CO_2} , called $Var_{X_{CO_2}}$ in this article, which accounts for the variability in the atmospheric state and the radiance residual variability but not uncertainty in the retrieval-algorithm inputs. This can result in a reported uncertainty that underestimates the actual retrieval error variance. Inference for carbon fluxes utilizes remote sensing data along with the reported uncertainties, so a more appropriate characterization of the error variance could lead to improved flux inversion. A geographically and seasonally comprehensive set of UQ experiments could provide guidance to adjusting the reported uncertainty in the operational data products. The results of this study suggest that adjustments would be especially warranted for high AOD conditions.

This study has investigated the impact of uncertainty in the retrieval prior mean μ_a as an algorithm input. We note that the model for uncertainty on μ_a can be written as

$$\mu_a - \mu_{\mathbf{X}} \sim \text{Gaussian}(\delta, \Omega_a)$$

for a given $\mu_{\mathbf{X}}$. Now, if μ_a is fixed, sampling from this distribution would generate uncertainty on the marginal mean, $\mu_{\mathbf{X}}$. Thus, the same MC draws of $\mu_a - \mu_{\mathbf{X}}$ could be used in a simulation experiment that considers uncertainty on the marginal mean, $\mu_{\mathbf{X}}$. The structure of the marginal distribution is also an important consideration; a more general representation than a multivariate Gaussian distribution may be appropriate.

Other key algorithm inputs, especially those linked to aerosols and albedo, likely impact the retrieval uncertainty. The investigation could be extended to incorporate uncertainty in the retrieval prior covariance Σ_a , particularly the portion corresponding to albedo and aerosols. The current OCO-2 operational algorithm uses a constant prior covariance matrix for all retrievals, and the impact of this choice on retrieval error distributions will depend on the spatially and temporally varying nature of the true marginal distribution [19].

The choice of forward-model parameters \mathbf{B} can impact the retrieval uncertainty as well. Several forward-model parameters characterize the wavelength dependence of aerosol absorption and scattering, and uncertainty in these parameters could impact the retrieval-error distribution. In addition, the forward model relies on discrete choices of aerosol types, which cannot perfectly capture the actual aerosol conditions in the atmosphere [9]. There is also potential in using collections of soundings \mathbf{Y} to estimate these forward-model parameters from the data.

This Monte Carlo framework is sufficiently general, and the surrogate model offers an adequate tradeoff between computational efficiency and physical realism to facilitate all of these potential UQ investigations for the OCO-2 OE retrieval. The framework simply requires a statistical model for the atmospheric state, a forward model representing the remote sensing instrument, and a retrieval algorithm for estimating the state given satellite observations. In fact, this framework could be used to provide uncertainty estimates for any retrieval algorithm, whether it is Bayesian or not.

OCO-2's implementation of the OE framework uses a numerical search for the posterior mode and provides a posterior covariance matrix based on a linear approximation. This article has addressed the propagation of uncertainty resulting from uncertain inputs into this specific algorithm and resulting estimator. Section 1 notes that the Bayesian formulation allows for other strategies for inference, including exploration of the full posterior distribution, $[\mathbf{X}|\mathbf{Y}]$. The OCO-2 FP forward model is likely too computationally expensive for posterior inference based on MCMC, for example, but sampling from the posterior distribution is feasible using the more efficient surrogate model developed here. As a reviewer has suggested, the comprehensive results that are efficiently produced with the surrogate model experiments can be compared to a subset of corresponding experiments with the FP forward model. This work is ongoing.

The OE remote sensing retrieval can be framed as an example of prediction in a nonlinear mixed model. This class of statistical models has been applied in a wide range of disciplines from medicine to environmental applications [10], and hence there is the potential to study the properties of predictors for random effects, or of estimators of fixed effects. The error budget

diagnostics developed and illustrated in this paper could be implemented in other applications of nonlinear mixed models.

Appendix A. The effect of linearity on the error budget. If the forward model is linear,

$$\mathbf{Y} = \boldsymbol{\mu} + \mathbf{K}\mathbf{X} + \boldsymbol{\epsilon},$$

then the error budget can be decomposed exactly into contributions from smoothing and noise. For the linear model, the posterior covariance \mathbf{S} , gain \mathbf{G} , and averaging kernel \mathbf{A} are given by

$$\begin{aligned}\mathbf{S} &= [\mathbf{K}^T \boldsymbol{\Sigma}_e^{-1} \mathbf{K} + \boldsymbol{\Sigma}_a^{-1}]^{-1}, \\ \mathbf{G} &= [\mathbf{K}^T \boldsymbol{\Sigma}_e^{-1} \mathbf{K} + \boldsymbol{\Sigma}_a^{-1}]^{-1} \mathbf{K}^T \boldsymbol{\Sigma}_e^{-1} \\ \mathbf{A} &= \mathbf{G}\mathbf{K}.\end{aligned}$$

Assume without loss of generality that $\boldsymbol{\mu} = \mathbf{0}$. For this model, the retrieval is linear,

$$\begin{aligned}\hat{\mathbf{X}} &= [\mathbf{K}^T \boldsymbol{\Sigma}_e^{-1} \mathbf{K} + \boldsymbol{\Sigma}_a^{-1}]^{-1} [\boldsymbol{\Sigma}_a^{-1} \boldsymbol{\mu}_a + \mathbf{K}^T \boldsymbol{\Sigma}_e^{-1} \mathbf{Y}] \\ &= \mathbf{S} \boldsymbol{\Sigma}_a^{-1} \boldsymbol{\mu}_a + \mathbf{G}\mathbf{Y} \\ &= \mathbf{S} \boldsymbol{\Sigma}_a^{-1} \boldsymbol{\mu}_a + \mathbf{G}(\mathbf{K}\mathbf{X} + \boldsymbol{\epsilon}) \\ &= \mathbf{S} \boldsymbol{\Sigma}_a^{-1} \boldsymbol{\mu}_a + \mathbf{A}\mathbf{X} + \mathbf{G}\boldsymbol{\epsilon}.\end{aligned}$$

Now,

$$\begin{aligned}\mathbf{A} + \mathbf{S} \boldsymbol{\Sigma}_a^{-1} &= \mathbf{S} (\mathbf{K}^T \boldsymbol{\Sigma}_e^{-1} \mathbf{K}) + \mathbf{S} \boldsymbol{\Sigma}_a^{-1} \\ &= \mathbf{S} (\mathbf{K}^T \boldsymbol{\Sigma}_e^{-1} \mathbf{K} + \boldsymbol{\Sigma}_a^{-1}) \\ &= \mathbf{S} \mathbf{S}^{-1} \\ &= \mathbf{I},\end{aligned}$$

so

$$\mathbf{S} \boldsymbol{\Sigma}_a^{-1} = \mathbf{I} - \mathbf{A}.$$

Then, the retrieval error can be written as

$$\begin{aligned}\hat{\mathbf{X}} - \mathbf{X} &= \mathbf{S} \boldsymbol{\Sigma}_a^{-1} \boldsymbol{\mu}_a + \mathbf{A}\mathbf{X} - \mathbf{X} + \mathbf{G}\boldsymbol{\epsilon} \\ &= (\mathbf{I} - \mathbf{A}) \boldsymbol{\mu}_a + (\mathbf{A} - \mathbf{I}) \mathbf{X} + \mathbf{G}\boldsymbol{\epsilon}.\end{aligned}$$

This results in the linear error budget

$$\begin{aligned}\boldsymbol{\Delta} &= \hat{\mathbf{X}} - \mathbf{X} \\ &= (\mathbf{A} - \mathbf{I}) (\mathbf{X} - \boldsymbol{\mu}_a) && \text{(smoothing)} \\ &+ \mathbf{G}\boldsymbol{\epsilon} && \text{(noise)}.\end{aligned}$$

Appendix B. Surrogate model description. Some of the key aspects of the surrogate forward model $\mathbf{F}(\mathbf{X}, \mathbf{B})$ include configuration of the atmospheric state vector \mathbf{X} , discretization of the atmospheric profile, trace gas absorption, RT, and viewing geometry.

Formally, the forward model $F_{i,j}(\mathbf{X}, \mathbf{B})$, $i = 1, \dots, n_j, j = 1, 2, 3$, defines the expected radiance as a function of the state \mathbf{X} and parameters \mathbf{B} for wavelength i in spectral band j . Hence, $n = n_1 + n_2 + n_3$. The three spectral bands correspond to the three OCO-2 spectrometers,

- O₂ A-band ($j = 1$), centered near 0.765 μm ,
- weak CO₂ band ($j = 2$), centered near 1.64 μm ,
- strong CO₂ band ($j = 3$), centered near 2.06 μm .

B.1. Vertical profile and state vector. The surrogate model discretizes the *atmospheric vertical profile* into $k = 1, \dots, K$ layers; the surrogate model uses $K = 19$. The atmospheric composition within a layer is assumed homogenous. Layer boundaries are defined by a unitless vertical coordinate $q_k = p_k/p_s$, where p_k is the atmospheric pressure at the top boundary of layer k and p_{k+1} is the pressure at the bottom boundary of layer k . The bottom layer is bounded by the surface, characterized by the surface pressure p_s . The $K + 1$ layer boundaries are fixed at $\{q_1 = 0.0001, q_2 = 1/K, q_3 = 2/K, \dots, q_{K+1} = 1.0\}$.

The *atmospheric state vector* \mathbf{X} includes

- the dry air mole fraction of CO₂, $c_{k'}$, at level $k', k' = 1, \dots, K + 1$;
- the surface pressure p_s ;
- coefficients $b_{0,\ell}, b_{1,\ell}, b_{2,\ell}, \ell = 1, \dots, 4$, representing the vertical profile of each of four atmospheric scattering species, including two composite aerosol types, cloud ice, and cloud water;
- coefficients $a_{0,j}, a_{1,j}$ representing the surface-albedo dependence on wavelength in each of the three spectral bands.

Some additional quantities defined below are functions of these state vector constituents. Any other quantities used are part of the parameter vector \mathbf{B} . These additional parameters include gas absorption coefficients and aerosol extinction and scattering coefficients.

B.2. Intermediate quantities. The surrogate model $F_{i,j}$ can be more conveniently defined in terms of several intermediate quantities, which are functions of \mathbf{X} and \mathbf{B} . The explicit notational expression of this dependence is dropped in subsequent discussion. These intermediate quantities include

- surface albedo $A_{i,j}$;
- vector of layer-specific optical depths $\boldsymbol{\tau}_{i,j} \equiv \{\tau_{i,j,k} : k = 1, \dots, K\}$;
- vector of layer-specific single-scattering albedo $\boldsymbol{\omega}_{i,j}(\boldsymbol{\tau}_{i,j}) \equiv \{\omega_{i,j,k}(\tau_{i,j,k}) : k = 1, \dots, K\}$;
- layer-specific phase function $\mathbf{P}_{i,j}(\boldsymbol{\tau}_{i,j}) \equiv \{P_{i,j,k}(\tau_{i,j,k}) : k = 1, \dots, K\}$.

The layer-specific optical depth $\tau_{i,j,k}$ quantifies the extinction of radiation in layer k . It is the sum of the optical depth for trace gas absorption $\tau_{G,i,j,k}$, from Rayleigh extinction $\tau_{R,i,j,k}$, and from each scattering species $\tau_{M,i,j,k,\ell}$,

$$\tau_{i,j,k} = \tau_{G,i,j,k} + \tau_{R,i,j,k} + \sum_{\ell=1}^4 \tau_{M,i,j,k,\ell}.$$

The optical depth due to *trace gas absorption* is a function of the abundance of the absorbing gas (O₂ or CO₂) in the atmospheric layer and a wavelength-dependent absorption coefficient

$\rho_{i,j,k}$. In the O₂ A-band,

$$\tau_{G,i,j,k} = 0.21 \rho_{i,j,k} \frac{p_s(q_{k+1} - q_k)}{gm_d}, \quad j = 1,$$

where m_d is the molar mass of dry air with units kg mol⁻¹ and g is the gravitational constant. In the weak and strong CO₂ bands,

$$\tau_{G,i,j,k} = \frac{c_k + c_{k+1}}{2} \rho_{i,j,k} \frac{p_s(q_{k+1} - q_k)}{gm_d}, \quad j = 2, 3.$$

The absorption coefficients $\rho_{i,j,k}$, with units m² mol⁻¹, are a set of fixed coefficients that are extracted from the OCO-2 FP absorption coefficient tables. The Rayleigh optical depth is

$$\tau_{R,i,j,k} = \rho_{R,i,j,k} \frac{p_s(q_{k+1} - q_k)}{gm_d},$$

where $\rho_{R,i,j,k}$ is a Rayleigh extinction coefficient, which is assumed known. Note that the quantity

$$\begin{aligned} \Delta p_k &= p_s(q_{k+1} - q_k) \\ &= p_{k+1} - p_k, \end{aligned}$$

is the pressure difference between the bottom and the top of layer k .

The AODs for each of the four scattering species are based on a characteristically shaped aerosol profile, parameterized by the coefficients $b_{0,\ell}$, $b_{1,\ell}$, $b_{2,\ell}$. The characteristic shape mimics a Gaussian probability density function. Then the layer-specific optical depths are defined as

$$\tau_{M,i,j,k,\ell} = e_{i,j,\ell} \exp\{b_{0,\ell}\} \frac{\Phi\left(\frac{q_{k+1}-b_{1,\ell}}{b_{2,\ell}}\right) - \Phi\left(\frac{q_k-b_{1,\ell}}{b_{2,\ell}}\right)}{\Phi\left(\frac{1.0-b_{1,\ell}}{b_{2,\ell}}\right) - \Phi\left(\frac{q_1-b_{1,\ell}}{b_{2,\ell}}\right)},$$

where Φ is the standard Gaussian cumulative distribution function. Each wavelength and scattering species has an extinction efficiency $e_{i,j,\ell}$ that is assumed known, and the shortest wavelength in the O₂ A-band is used as a reference with $e_{1,1,\ell} = 1$. Then $\exp\{b_{0,\ell}\}$ is the total optical depth at this reference wavelength for each scattering species. The coefficient $b_{1,\ell}$ defines the peak height of the aerosol profile, and $b_{2,\ell}$ characterizes the effective depth of the profile.

In addition to extinction from multiple sources, the forward function also incorporates Rayleigh scattering and scattering by the four scattering species. Scattering behavior is quantified by the single-scattering albedo $\omega_{i,j,k}(\tau_{i,j,k})$ and the phase function, $P_{i,j,k}(\tau_{i,j,k})$. The single scattering albedo is defined as

$$\omega_{i,j,k}(\tau_{i,j,k}) = \frac{\tau_{R,i,j,k} + \sum_{\ell=1}^4 \omega_{M,i,j,\ell} \tau_{M,i,j,k,\ell}}{\tau_{i,j,k}}.$$

Each scattering species has its own wavelength-dependent single-scattering albedo, $\omega_{M,i,j,\ell}$, which quantifies the fraction of scattered radiation to extinction, and these parameters are assumed known.

The phase function $P_{i,j,k}(\tau_{i,j,k})$ characterizes angular dependence of scattering,

$$P_{i,j,k}(\tau_{i,j,k}) = \frac{\tau_{R,i,j,k}P_{R,i,j} + \sum_{\ell=1}^4 \omega_{M,i,j,\ell}\tau_{M,i,j,k,\ell}P_{M,i,j,\ell}}{\tau_{R,i,j,k} + \sum_{\ell=1}^4 \omega_{M,i,j,\ell}\tau_{M,i,j,k,\ell}},$$

where $P_{R,i,j}$ and $P_{M,i,j,\ell}$ are known phase functions for Rayleigh scattering and the individual scattering species.

Finally the surface albedo provides a lower boundary condition for the transfer of radiation through the atmosphere. The surrogate model assumes a Lambertian surface and the wavelength dependence of albedo is represented by

$$A_{i,j} = a_{0,j} + a_{1,j}(\nu_{i,j} - \nu_j^{(0)}),$$

where $\nu_{i,j}$ is the wavenumber of interest and $\nu_j^{(0)}$ is a predefined reference wavenumber for each band.

B.3. Radiative transfer. The surface albedo, optical depth, single-scattering albedo, and phase function are inputs to computational routines for RT. A variety of routines of varying complexity and numerical accuracy are available for solving the RT equation, which is an integro-differential equation for the intensity of radiation as a function of the path through the atmosphere. Additional inputs for RT include the solar geometry and satellite viewing geometry (zenith and azimuth angles). Vector RT routines solve for the full Stokes vector, which incorporates scalar intensity along with polarization. The surrogate model $F_{i,j}$ includes a fully polarized first order of scattering (FO) routine and a scalar two-stream (2S) approximation for the contribution from multiple scattering. The FO routine outputs the top of atmosphere (TOA) Stokes vector $(I_{FO,i,j}, Q_{FO,i,j}, U_{FO,i,j})$, and the 2S routine outputs a (TOA) multiple scattering intensity $I_{2S,i,j}$. This RT implementation is one key distinction between the surrogate model and the OCO-2 FP forward model, where the latter utilizes more numerically accurate second-order of scattering and a larger number of streams for multiple scattering [16, 18].

The instrument geometry defines the Stokes coefficients (M_I, M_Q, M_U) , and the expected radiance can be computed as

$$\begin{aligned} F_{i,j}(\mathbf{X}, \mathbf{B}) = & M_I I_{FO,i,j}(A_{i,j}, \tau_{i,j}, \omega_{i,j}(\tau_{i,j}), \mathbf{P}_{i,j}(\tau_{i,j})) \\ & + M_I I_{2S,i,j}(A_{i,j}, \tau_{i,j}, \omega_{i,j}(\tau_{i,j}), \mathbf{P}_{i,j}(\tau_{i,j})) \\ & + M_Q Q_{FO,i,j}(A_{i,j}, \tau_{i,j}, \omega_{i,j}(\tau_{i,j}), \mathbf{P}_{i,j}(\tau_{i,j})) \\ & + M_U U_{FO,i,j}(A_{i,j}, \tau_{i,j}, \omega_{i,j}(\tau_{i,j}), \mathbf{P}_{i,j}(\tau_{i,j})). \end{aligned}$$

Acknowledgments. The authors thank Hai Nguyen, Dejian Fu, James McDuffie, Vijay Natraj, Brian Connor, Lesley Ott, Chris O'Dell, Tommy Taylor, and Annmarie Eldering for valuable suggestions and technical assistance.

REFERENCES

- [1] A. BACCINI, S. J. GOETZ, W. S. WALKER, N. T. LAPORTE, M. SUN, D. SULLA-MENASHE, J. HACKLER, P. S. A. BECK, R. DUBAYAH, M. A. FRIEDL, S. SAMANTA, AND R. A. HOUGHTON, *Estimated carbon dioxide emissions from tropical deforestation improved by carbon-density maps*, *Nature Clim. Change*, 2 (2012), pp. 182–185.
- [2] V. BUCHARD, A. M. DA SILVA, P. R. COLARCO, A. DARMENOV, C. A. RANDLES, R. GOVINDARAJU, O. TORRES, J. CAMPBELL, AND R. SPURR, *Using the OMI aerosol index and absorption aerosol optical depth to evaluate the NASA MERRA aerosol reanalysis*, *Atmos. Chem. Phys.*, 15 (2015), pp. 5743–5760.
- [3] P. CLAIS, C. SABINE, G. BALA, L. BOPP, V. BROVKIN, J. CANADELL, A. CHHABRA, R. DEFRIES, J. GALLOWAY, M. HEIMANN, C. JONES, C. L. QUÉRÉ, R. B. MYNENI, S. PIAO, AND P. THORNTON, *Carbon and other biogeochemical cycles*, in *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, T. F. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S. K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. M. Midgley, eds., Cambridge University Press, Cambridge, 2013, pp. 465–570.
- [4] B. CONNOR, H. BÖSCH, J. MCDUFFIE, T. TAYLOR, D. FU, C. FRANKENBERG, C. O'DELL, V. H. PAYNE, M. GUNSON, R. POLLOCK, J. HOBBS, F. OYAFUSO, AND Y. JIANG, *Quantification of uncertainties in OCO-2 measurements of XCO₂: Simulations and linear error analysis*, *Atmos. Measure. Tech.*, 9 (2016), pp. 5227–5238.
- [5] B. J. CONNOR, H. BOESCH, G. TOON, B. SEN, C. MILLER, AND D. CRISP, *Orbiting carbon observatory: Inverse method and prospective error analysis*, *J. Geophys. Res. Atmos.*, 113 (2008), D05305, <https://doi.org/10.1029/2006JD008336>.
- [6] N. CRESSIE AND S. BURDEN, *Figures of merit for simultaneous inference and comparisons in simulation experiments*, *Stat.*, 4 (2015), pp. 196–211.
- [7] N. CRESSIE AND R. WANG, *Statistical properties of the state obtained by solving a nonlinear multivariate inverse problem*, *Appl. Stoch. Models Bus. Ind.*, 29 (2013), pp. 424–438.
- [8] N. CRESSIE, R. WANG, M. SMYTH, AND C. E. MILLER, *Statistical bias and variance for the regularized inverse problem: Application to space-based atmospheric CO₂ retrievals*, *J. Geophys. Res. Atmos.*, 121 (2016), pp. 5526–5537.
- [9] D. CRISP, H. BOESCH, L. BROWN, R. CASTANO, M. CHRISTI, B. CONNOR, A. ELDERING, C. FRANKENBERG, M. GUNSON, J. MCDUFFIE, C. MILLER, V. NATRAJ, C. O'DELL, D. O'BRIEN, I. POLONSKY, F. OYAFUSO, M. SMYTH, D. THOMPSON, G. TOON, AND R. SPURR, *Orbiting Carbon Observatory (OCO)-2 Level 2 Full Physics Algorithm Theoretical Basis Document*, Technical report, OCO D-65488, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 2014.
- [10] M. DAVIDIAN AND D. M. GILTINAN, *Nonlinear models for repeated measurement data: An overview and update*, *J. Agric. Biol. Environ. Stat.*, 8 (2003), pp. 387–419.
- [11] A. ELDERING, R. POLLOCK, R. LEE, R. ROSENBERG, F. OYAFUSO, D. CRISP, L. CHAPSKY, AND R. GRANAT, *Orbiting Carbon Observatory (OCO)-2 Level 1b Theoretical Basis Document*, Technical report, OCO D-55206, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 2015.
- [12] H. HAARIO, M. LAINE, M. LEHTINEN, E. SAKSMAN, AND J. TAMMINEN, *Markov chain Monte Carlo methods for high dimensional inversion in remote sensing*, *J. Roy. Statist. Soc. Ser. B*, 66 (2004), pp. 1–17.
- [13] J. R. HARRING AND J. LIU, *A comparison of estimation methods for nonlinear mixed effects models under model misspecification and data sparseness: A simulation study*, *J. Mod. Appl. Stat. Meth.*, 15 (2016), pp. 539–569.
- [14] A. HARTFORD AND M. DAVIDIAN, *Consequences of misspecifying assumptions in nonlinear mixed effects models*, *Comput. Statist. Data Anal.*, 34 (2000), pp. 139–164.
- [15] P. K. KITANIDIS, *Quasi-linear geostatistical theory for inversing*, *Water Res. Res.*, 31 (1995), pp. 2411–2419.
- [16] V. NATRAJ AND R. J. D. SPURR, *A fast linearized pseudo-spherical two orders of scattering model to account for polarization in vertically inhomogeneous scattering-absorbing media*, *J. Quant. Spectros.*, 107 (2007), pp. 263–293.

- [17] NRC, *Assessing the Reliability of Complex Models: Mathematical and Statistical Foundations of Verification, Validation, and Uncertainty Quantification*, The National Academies Press, Washington, DC, 2012.
- [18] C. W. O'DELL, *Acceleration of multiple-scattering, hyperspectral radiative transfer calculations via low-streams interpolation*, J. Geophys. Res. Atmos., 115 (2010), D10206, <https://doi.org/10.1029/2009JD012803>.
- [19] C. W. O'DELL, B. CONNOR, H. BOESCH, D. O'BRIEN, C. FRANKENBERG, R. CASTANO, A. ELDERING, B. FISHER, M. GUNSON, J. MCDUFFIE, C. E. MILLER, V. NATRAJ, F. OYAFUSO, I. POLONSKY, M. SMYTH, T. TAYLOR, G. C. TOON, P. O. WENNBERG, AND D. WUNCH, *The ACOS CO₂ retrieval algorithm – Part 1: Description and validation against synthetic observations*, Atmos. Meas. Tech., 5 (2012), pp. 99–121.
- [20] A. O'HAGAN, *Probabilistic uncertainty specification: Overview, elaboration techniques and their application to a mechanistic model of carbon flux*, Environ. Model. Assess., 36 (2012), pp. 35–48.
- [21] L. E. OTT, S. PAWSON, G. J. COLLATZ, W. W. GREGG, D. MENEMENLIS, H. BRIX, C. S. ROUSSEAU, K. W. BOWMAN, J. LIU, A. ELDERING, M. R. GUNSON, AND S. R. KAWA, *Assessing the magnitude of CO₂ flux uncertainty in atmospheric CO₂ records using products from NASA's Carbon Monitoring Flux Pilot Project*, J. Geophys. Res. Atmos., 120 (2015), pp. 734–765.
- [22] C. D. RODGERS, *Inverse Methods for Atmospheric Sounding*, World Scientific, Hackensack, NJ, 2000.
- [23] C. D. RODGERS AND B. J. CONNOR, *Intercomparison of remote sounding instruments*, J. Geophys. Res. Atmos., 108 (2003), 4116, <https://doi.org/10.1029/2002JD002299>.
- [24] J. TAMMINEN, *Validation of nonlinear inverse algorithms with Markov chain Monte Carlo method*, J. Geophys. Res. Atmos., 109 (2004), 019303, <https://doi.org/10.1029/2004JD004927>.
- [25] J. TAMMINEN AND E. KYRÖLÄ, *Bayesian solution for nonlinear and non-Gaussian inverse problems by Markov chain Monte Carlo method*, J. Geophys. Res. Atmos., 106 (2001), pp. 14,377–14,390.
- [26] W. WANNER, A. H. STRAHLER, B. HU, P. LEWIS, J.-P. MULLER, X. LI, C. L. B. SCHAAF, AND M. J. BARNESLEY, *Global retrieval of bidirectional reflectance and albedo over land from EOS MODIS and MISR data: Theory and algorithm*, J. Geophys. Res. Atmos., 102 (1997), pp. 17,143–17,161.
- [27] J. X. WARNER, Z. WEI, L. L. STROW, C. D. BARNET, L. C. SPARLING, G. DISKIN, AND G. SACHSE, *Improved agreement of AIRS tropospheric carbon monoxide products with other EOS sensors using optimal estimation retrievals*, Atmos. Chem. Phys., 10 (2010), pp. 9521–9533.