



Inference over radiative transfer models using variational and expectation maximization methods

Daniel Heestermans Svendsen¹  · Daniel Hernández-Lobato² · Luca Martino³ · Valero Laparra¹ · Álvaro Moreno-Martínez¹ · Gustau Camps-Valls¹

Received: 20 May 2020 / Revised: 25 January 2021 / Accepted: 14 May 2021 /

Published online: 8 June 2021

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2021

Abstract

Earth observation from satellites offers the possibility to monitor our planet with unprecedented accuracy. Radiative transfer models (RTMs) encode the energy transfer through the atmosphere, and are used to model and understand the Earth system, as well as to estimate the parameters that describe the status of the Earth from satellite observations by inverse modeling. However, performing inference over such simulators is a challenging problem. RTMs are nonlinear, non-differentiable and computationally costly codes, which adds a high level of difficulty in inference. In this paper, we introduce two computational techniques to infer not only point estimates of biophysical parameters but also their joint distribution. One of them is based on a variational autoencoder approach and the second one is based on a Monte Carlo Expectation Maximization (MCEM) scheme. We compare and discuss benefits and drawbacks of each approach. We also provide numerical comparisons in synthetic simulations and the real PROSAIL model, a popular RTM that combines land vegetation leaf and canopy modeling. We analyze the performance of the two approaches for modeling and inferring the distribution of three key biophysical parameters for quantifying the terrestrial biosphere.

Keywords Variational autoencoder · Expectation maximization · Radiative transfer model · Inverse modeling · Density estimation

Editors: Dino Ienco, Thomas Corpetti, Minh-Tan Pham, Sebastien Lefevre, Roberto Interdonato.

✉ Daniel Heestermans Svendsen
daniel.svendsen@uv.es

¹ Image Processing Laboratory (IPL), Universitat de València, Valencia, Spain

² Dep. Computer Science, Universidad Autónoma de Madrid, Madrid, Spain

³ Dep. Signal Theory and Telecommunications, Universidad Rey Juan Carlos de Madrid, Madrid, Spain

1 Introduction

In many areas of science and engineering, systems are analyzed by running computer code simulations, which act as convenient approximations to reality. Depending on the body of literature, they are known as physics-based, processed-oriented and mechanistic models, or simply just *simulators* (Santner et al., 2003; Wescott, 2013). Simulators are ubiquitous in physics, brain, social, Earth and climate sciences (Raissi et al., 2017; Sandberg, 2013; Svendsen et al., 2020). Model simulations are needed to understand system behaviour, but also to perform counterfactual studies.

In Earth sciences the use of simulators is of paramount importance. Earth observation (EO) from airborne and satellite remote sensing platforms along with in-situ observations play a fundamental role in monitoring our planet (Lillesand et al., 2008; Liang, 2008; Rodgers, 2000). Remote sensing simulators of the involved processes are known as *radiative transfer models* (RTMs). These models describe the complex interactions of scattering and absorption of radiation with the constituents of the atmosphere, water, vegetation and soils. RTMs are useful because they allow us to translate (map) a set of parameter¹ values describing the state of soil, leaf, canopy and atmosphere to at-sensor reflectance or radiance. Such simulations allow for modeling, understanding, and predicting parameters related to the state of the land cover, water bodies and atmosphere.

While modeling and characterizing the involved processes is key, in practice one is typically interested in solving the so-called *inverse problem*; that is, for example, inferring the set of atmospheric or canopy biophysical properties, so that the computed reflectances best fit the remotely sensed ones (Tarantola, 2005; Svendsen et al., 2018; Zurita-Milla et al., 2015). The problem of inverting the forward model is in general highly *ill-posed* (Combal et al., 2003): Different sets of parameter values can map into the same reflectance, thus making it difficult to recover the true set of parameters given a remotely sensed reflectance. This issue has been largely reported in the literature (Verstraete & Pinty, 1996; Knyazikhin, 1999; Liang, 2008) and is, together with the complexity and computational cost of the RTMs as well as the scarcity of labeled data, the main reasons why the inverse problem is a difficult and unresolved one.

Many methods have been proposed for model inversion. Early approaches considered minimizing the (e.g. least squares) error between observations and model simulations stored in big Look-Up-Tables (LUTs). Comparing each observed spectrum with all spectra stored in the LUT proved impractical, leading to gradient-descent techniques in combination with emulators to be proposed in the literature (Vicent et al., 2019). More advanced approaches have recently exploited machine learning regression algorithms, such as random forests (Liang et al., 2015; Campos-Taberner et al., 2018), neural networks (Baret, 2007; Djamai et al., 2019) and Gaussian processes (Svendsen et al., 2020, 2018; Camps-Valls et al., 2018, 2019) to achieve improved multidimensional interpolation capabilities. Treating the inverse problem purely as a regression problem, however, only leads to point-wise estimates, and not a joint probability distribution of the parameters. We argue in this work that when two or more physical variable configurations result in the same spectrum, a conventional inversion method will

¹ Note that both the field of remote sensing parameter retrieval and the field of Bayesian inference make heavy use of the word *parameter*. In the former case it refers to the physical parameters over which we are inferring distributions in this work, and in the latter case it refers to distribution parameters (mean, covariance, etc.). The meaning is always clear from context, however.

perform poorly. Such methods that try to predict the cause from the effect will simply predict something in between the possible causes that gave rise to the effect. This problem corresponds to having a multimodal posterior and can be addressed with one of the proposed probabilistic frameworks in this paper.

A Bayesian formalism is useful in order to (1) generate probability density functions (PDFs) of the parameters, and hence account for all moments and uncertainty on the retrieved parameters (Pinty et al., 2011; Lewis et al., 2012), (2) to incorporate constraints in the form of a-priori parameter distributions, which are often subject of intense debate in the literature (Combal et al., 2003; Baret, 2007; Atzberger & Richter, 2012), and (3) to overcome the limited potential of iterative steepest-descent optimization procedures to locate globally optimal solutions (Bacour et al., 2002; Zhang et al., 2005). Furthermore, deriving a generative model provides a straightforward way to perform outlier detection, by measuring the probability of the observed data under the fitted model.

Given some vector of physical causes \mathbf{c} (atmospheric or canopy properties), the forward RTM model induces a likelihood function $p(\mathbf{e}|\mathbf{c})$, which links the causes with the physical effects \mathbf{e} (reflectance spectra). In this work, we address a general problem: Learning the distribution of the physical parameters or *causes*, instead of only providing a pointwise estimation of these parameters (by statistical or numerical inversion). Provided a dataset of observed effects \mathbf{e}' , our goal is twofold: learning the marginal density $p(\mathbf{c})$ and obtaining an approximation of the conditional distribution $p(\mathbf{c}|\mathbf{e}')$, which in a Bayesian setting represents the posterior density of the causes given the effects. Note that $p(\mathbf{c}|\mathbf{e}')$ also represents a probabilistic inverse model, i.e., given \mathbf{e}' we can obtain a prediction of the causes \mathbf{c} and related uncertainty measures. Probabilistic inverse modelling, although not so widely used in Remote Sensing applications, has proven to be a powerful tool, providing more general (and hence potentially more valuable) solutions than point-wise approaches, and can help in better understanding the problem itself (Zhang et al., 2005; Coccia et al., 2015; Ma et al., 2017).

Since RTMs are generally complex, non-differentiable (i.e. having non-analytical Jacobian) and computationally costly models, mathematical tractability is typically compromised, especially when the aim is to combine RTMs and Bayesian methods. Here, we propose and compare two different approaches which allows us to infer parameters for a non-differentiable simulator. One approach is based on Monte Carlo Expectation Maximization (MCEM) (Wei & Tanner, 1990) and the other is based on Variational Autoencoders (VAEs) (Kingma & Welling, 2013). We will show that each approach has different pros and cons. While the MCEM approach is mathematically elegant, flexible and has good convergence properties, its application in practice is computationally demanding. On the other hand the proposal based on a simple version of VAE obtains good results and is fast, yet it is not able to describe multimodal distributions. While possible, its extension for multimodal distributions makes the approach more complicated, reducing the good computational properties (see, e.g., Mescheder et al., 2017). We illustrate these properties in several toy examples of varying sample sizes and complexity, as well as with the PROSAIL RTM (Baret et al., 1992). PROSAIL is the combination of the PROSPECT (Jacquemoud & Baret, 1990) leaf optical properties model and the SAIL (Verhoef, 1984) (Scattering by Arbitrary Inclined Leaves) canopy reflectance model. In particular, we compare the approaches for inferring the distribution of three key parameters for quantifying the terrestrial biosphere.

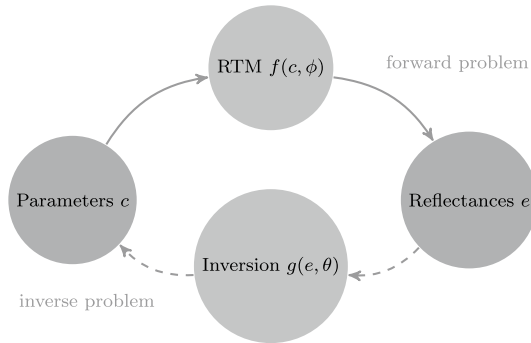


Fig. 1 The forward problem in Earth observation involves taking the system’s structural state as an input, defining representative bio-geo-physical parameters (e.g. vegetation canopy or leaf characteristics), then propagating the solar energy through the atmosphere medium and producing a simulated at-sensor reflectance. The inverse problem involves performing inference over the forward model. We use \mathbf{c} to denote the model parameters or causes, \mathbf{f} is an RTM or simulator, and \mathbf{e} is the simulated output reflectance or effect. Both the forward RTM \mathbf{f} and the inverse model \mathbf{g} are nonlinear functions parameterized by θ and ϕ , respectively

2 Proposed methodology

2.1 Forward and inverse modeling

Notationally, an RTM \mathbf{f} operating in *forward mode* generates a multidimensional *reflectance/radiance observation (or effect)* $\mathbf{e} \in \mathbb{R}^{D_e}$ as observed by the sensor given a multidimensional *parameter state vector (or cause)* $\mathbf{c} \in \mathbb{R}^{D_c}$, see Fig. 1. Running forward simulations yields a look-up-table (LUT) of input-output pairs, $\mathcal{D} = \{(\mathbf{c}_i, \mathbf{e}_i)\}_{i=1}^n$. Solving the *inverse problem* using machine learning implies learning the function \mathbf{g} using \mathcal{D} , to return an estimate \mathbf{c}_* each time a new satellite observation \mathbf{e}_* is acquired.

2.2 Problem setting

Notationally, let us consider then the vector of effects $\mathbf{e} \in \mathbb{R}^{D_e}$ and vector of causes $\mathbf{c} \in \mathcal{C} \subseteq \mathbb{R}^{D_c}$, an RTM model represents the underlying mapping from \mathbf{c} to \mathbf{e} , that we denote as $\mathbf{f}(\mathbf{c}) : \mathbb{R}^{D_c} \rightarrow \mathbb{R}^{D_e}$. The complete observation model is given by

$$\mathbf{e} = \mathbf{f}(\mathbf{c}) + \epsilon, \quad \epsilon \sim \mathcal{N}(\epsilon|\mathbf{0}, \sigma^2\mathbf{I}), \tag{1}$$

where \mathbf{I} is a unit $D_e \times D_e$ matrix. The observation model defines the likelihood function as

$$p(\mathbf{e}|\mathbf{c}) = \mathcal{N}(\mathbf{e}|\mathbf{f}(\mathbf{c}), \sigma^2\mathbf{I}). \tag{2}$$

Note that by fixing \mathbf{c} , the conditional probability $p(\mathbf{e}|\mathbf{c})$ is Gaussian, but as a function of \mathbf{c} the likelihood is a highly non-linear function due to the dependence on the RTM with the causes, i.e. $\mathbf{f}(\mathbf{c})$. We assume an Gaussian prior over \mathbf{c} ’s,

$$p(\mathbf{c}) = \mathcal{N}(\mathbf{c}|\mathbf{m}, \mathbf{S}), \tag{3}$$

where $\mathbf{m} \in \mathbb{R}^{D_c}$ and the $D_c \times D_c$ covariance matrix \mathbf{S} are considered unknown. The posterior density given the observed data \mathbf{e} over the causes can be expressed as

$$p(\mathbf{c}|\mathbf{e}) \propto p(\mathbf{e}|\mathbf{c})p(\mathbf{c}) = \mathcal{N}(\mathbf{e}|\mathbf{f}(\mathbf{c}), \sigma^2\mathbf{I})\mathcal{N}(\mathbf{c}|\mathbf{m}, \mathbf{S}) \quad (4)$$

Our goals are: (a) learn the prior parameters, vector \mathbf{m} and matrix \mathbf{S} , and (b) obtain an approximation of the posterior $p(\mathbf{c}|\mathbf{e})$, which serves as an inverse probabilistic mapping from \mathbf{e} to \mathbf{c} . We assume that some set of data \mathbf{e} is given. The two main ways of approaching this problem are a Variational inference (VI²) scheme on the one hand, and an expected maximization method on the other. For the VI method we follow the approach of Kingma and Welling (2013) and substitute the decoder network with the generative model of Eq. (1). For the MC-based approach we use MC Expectation Maximization (Wei & Tanner, 1990).

2.3 Variational inference method

The idea of variational inference is to optimize the parameters of a *variational posterior* in order to come as close as possible to the *true posterior*. Following Kingma and Welling (2013) we choose a Gaussian variational posterior,

$$q(\mathbf{c}|\mathbf{e}) = \mathcal{N}(\mathbf{c}|\boldsymbol{\mu}_{\text{NN}}(\mathbf{e}), \boldsymbol{\Sigma}_{\text{NN}}(\mathbf{e})), \quad (5)$$

where $\boldsymbol{\mu}_{\text{NN}}(\mathbf{e})$ and $\boldsymbol{\Sigma}_{\text{NN}}(\mathbf{e})$ are obtained by tuning a Neural Network (NN) with parameters $\boldsymbol{\phi}$. These parameters $\boldsymbol{\phi}$ are also referred to as the variational parameters. In order to tune the NN parameters $\boldsymbol{\phi}$, we minimize the Kullback-Leibler (KL) divergence between $q(\mathbf{c}|\mathbf{e})$ and the true posterior $p(\mathbf{c}|\mathbf{e})$, i.e.,

$$\begin{aligned} \text{KL}[q(\mathbf{c}|\mathbf{e})||p(\mathbf{c}|\mathbf{e})] &= -\mathbb{E}_{q(\mathbf{c}|\mathbf{e})} \left[\log \frac{p(\mathbf{c}|\mathbf{e})}{q(\mathbf{c}|\mathbf{e})} \right] \\ &= -\mathbb{E}_{q(\mathbf{c}|\mathbf{e})} \left[\log \frac{p(\mathbf{c}, \mathbf{e})}{q(\mathbf{c}|\mathbf{e})} - \log p(\mathbf{e}) \right] \\ &= -\mathcal{L} + \log p(\mathbf{e}), \end{aligned} \quad (6)$$

where we have used $\mathbb{E}_{q(\mathbf{c}|\mathbf{e})}[\log p(\mathbf{e})] = \log p(\mathbf{e})$. Since $\log p(\mathbf{e})$ is constant w.r.t. the variational parameters, in order to minimize the KL divergence, we have to maximize

$$\mathcal{L} = \mathbb{E}_{q(\mathbf{c}|\mathbf{e})} \left[\log \frac{p(\mathbf{c}, \mathbf{e})}{q(\mathbf{c}|\mathbf{e})} \right] = \mathbb{E}_{q(\mathbf{c}|\mathbf{e})} \left[\log \frac{p(\mathbf{e}|\mathbf{c})p(\mathbf{c})}{q(\mathbf{c}|\mathbf{e})} \right],$$

called the *Evidence Lower Bound* (ELBO). We can split the ELBO into two terms: the first one represents the expected log-likelihood with respect to the variational posterior and the second one is the KL divergence between the variational posterior and the prior, i.e.,

² Please note that we use *VI* to abbreviate *variational inference* as in most machine learning literature. This is not to be confused with the abbreviation of *vegetation indices* in the remote sensing literature.

$$\begin{aligned}
\mathcal{L} &= \mathbb{E}_{q(\mathbf{c}|\mathbf{e})} \left[\log \frac{p(\mathbf{e}|\mathbf{c})p(\mathbf{c})}{q(\mathbf{c}|\mathbf{e})} \right], \\
&= \mathbb{E}_{q(\mathbf{c}|\mathbf{e})} \left[\log p(\mathbf{e}|\mathbf{c}) \right] - \left[\frac{q(\mathbf{e}|\mathbf{c})}{p(\mathbf{c}|\mathbf{e})} \right], \\
&= \mathbb{E}_{q(\mathbf{c}|\mathbf{e})} \left[\log p(\mathbf{e}|\mathbf{c}) \right] - \text{KL}[q(\mathbf{c}|\mathbf{e})||p(\mathbf{c})].
\end{aligned} \tag{7}$$

As opposed to the approach in Kingma and Welling (2013) we place a deterministic forward model in stead of a decoder network and fix a low value of noise variance in the likelihood (Eq. 2) in order to reflect the trust in the forward model. This approach is akin to that of McCarthy et al. (2017). In order to optimize this expression, we perform a Monte Carlo estimation of the expected value (i.e., the first term) (Robert & Casella, 2013). The second term has a simple analytical form as it is the KL divergence between two Gaussians.

Importantly, maximizing \mathcal{L} with respect to ϕ should make $\text{KL}[q(\mathbf{c}|\mathbf{e})||p(\mathbf{c}|\mathbf{e})]$ fairly small and hence $\mathcal{L} \approx \log p(\mathbf{e})$. The maximization of \mathcal{L} with respect to the prior parameters, $\theta = \{\mathbf{m}, \mathbf{S}\}$, is hence expected to maximize $\log p(\mathbf{e})$, which is the maximum likelihood principle for parameter estimation. In practice, we maximize \mathcal{L} simultaneously with respect to θ and ϕ .

The previous approach can be easily extended to the case of having several observed data instances $\{\mathbf{e}_i\}_{i=1}^N$. In that case the objective is simply the sum of \mathcal{L}_i , for $i = 1, \dots, N$, where \mathcal{L}_i is the lower bound corresponding to \mathbf{e}_i , i.e., the i th data instance. This sum can be approximated using mini-batches and optimized using stochastic optimization techniques such as the ADAM algorithm (Kingma & Ba, 2015). In this study, we use a mini-batch size of 1 for all experiments. For a proof of convergence of stochastic optimization see Robbins & Monro (1951). The variational approach is expected to find reasonable values for the prior parameters θ , using approximate maximum likelihood estimation, and to provide a recognition model $q(\mathbf{c}|\mathbf{e})$ that can be used to infer the potential values of \mathbf{c} given \mathbf{e} .

2.4 Monte Carlo expectation maximization

Another method which can be used to address the learning goals described in Sect. 2, i.e. to infer the prior parameters from the observed data, and to generate samples from the posterior distribution $p(\mathbf{c}|\mathbf{e})$, is the Monte Carlo expectation maximization (MCEM) method (Wei & Tanner, 1990).

We begin by briefly describing the expectation maximization (EM) algorithm, which can be used to maximize the likelihood function in models that involve latent variables (Dempster et al., 1977). This is precisely the scenario considered in Sect. 2. Namely, given some observed data $\{\mathbf{e}_i\}_{i=1}^N$, we would like to maximize

$$\prod_{i=1}^N p(\mathbf{e}_i|\theta) = \prod_{i=1}^N \int_{\mathcal{C}} p(\mathbf{e}_i|\mathbf{c}_i)p(\mathbf{c}_i|\theta)d\mathbf{c}_i, \tag{8}$$

as a function of the prior parameters $\theta = \{\mathbf{m}, \mathbf{S}\}$. Direct optimization of (8) is intractable, since we cannot marginalize the latent variables \mathbf{c}_i . The EM algorithm uses the fact that the complete likelihood function $p(\mathbf{e}_i, \mathbf{c}_i|\theta) = p(\mathbf{e}_i|\mathbf{c}_i)p(\mathbf{c}_i|\theta)$ is tractable. Consider the following decomposition of the logarithm of (8)

$$\log \sum_{i=1}^N p(\mathbf{e}_i | \boldsymbol{\theta}) = \sum_{i=1}^N \mathcal{L}(q_i, \boldsymbol{\theta}) + \text{KL}(q_i || p_i), \quad (9)$$

where we have introduced an approximate distribution $q_i(\mathbf{c}_i)$ and

$$\mathcal{L}(q_i, \boldsymbol{\theta}) = \int_{\mathcal{C}} q_i(\mathbf{c}_i) \log \frac{p(\mathbf{e}_i, \mathbf{c}_i | \boldsymbol{\theta})}{q_i(\mathbf{c}_i)} d\mathbf{c}_i, \quad (10)$$

$$\text{KL}(q_i || p_i) = - \int_{\mathcal{C}} q_i(\mathbf{c}_i) \log \frac{p(\mathbf{c}_i | \mathbf{e}_i, \boldsymbol{\theta})}{q_i(\mathbf{c}_i)} d\mathbf{c}_i. \quad (11)$$

Note that (10) is the Kullback Leibler divergence between q_i and the exact posterior $p(\mathbf{c}_i | \mathbf{e}_i, \boldsymbol{\theta})$ for the instance \mathbf{e}_i .

The EM algorithm maximizes (9) in a two stage iterative process. Assume the current parameter vector is $\boldsymbol{\theta}^{\text{old}}$. In the E step, the lower bound $\sum_{i=1}^N \mathcal{L}(q_i, \boldsymbol{\theta}^{\text{old}})$ is maximized with respect to each q_i assuming $\boldsymbol{\theta}^{\text{old}}$ to be fixed. Because $\sum_{i=1}^N \log p(\mathbf{e}_i | \boldsymbol{\theta})$ does not depend on each q_i , the solution to this problem consists in setting each $q_i(\mathbf{c}_i)$ equal to $p(\mathbf{c}_i | \mathbf{e}_i, \boldsymbol{\theta}^{\text{old}})$, minimizing $\text{KL}(q_i || p_i)$ in consequence. In the subsequent M step, each $q_i(\mathbf{c}_i)$ is held fixed, and $\sum_{i=1}^N \mathcal{L}(q_i, \boldsymbol{\theta}^{\text{old}})$ is maximized with respect to $\boldsymbol{\theta}$, to give new prior parameters $\boldsymbol{\theta}^{\text{new}}$. This will cause the lower bound $\sum_{i=1}^N \mathcal{L}(q_i, \boldsymbol{\theta}^{\text{old}})$ to increase, which will in turn increase the log-likelihood $\sum_{i=1}^N \log p(\mathbf{e}_i | \boldsymbol{\theta})$. Critically, $q_i(\mathbf{c}_i)$ will be computed in this step using $\boldsymbol{\theta}^{\text{old}}$, which is fixed. Therefore, the only required integral to evaluate in the M step is

$$\mathcal{L}(q_i, \boldsymbol{\theta}) = \int_{\mathcal{C}} q_i(\mathbf{c}_i) \log p(\mathbf{c}_i | \boldsymbol{\theta}) d\mathbf{c}_i + \text{const}. \quad (12)$$

A difficulty, however, is that the posterior $p(\mathbf{c}_i | \mathbf{e}_i, \boldsymbol{\theta}^{\text{old}})$ is intractable, which makes computing q_i and hence the integral in (12) challenging. Monte Carlo EM (MCEM), provides a solution to this problem (Wei & Tanner, 1990). The intractable integral in (12) is simply approximated by a Monte Carlo average over several samples drawn from q_i . Namely,

$$\mathcal{L}(q_i, \boldsymbol{\theta}) \approx \frac{1}{S} \sum_{s=1}^S \log p(\mathbf{c}_i^s | \boldsymbol{\theta}) + \text{const}. , \quad (13)$$

where \mathbf{c}_i^s has been generated from q_i and S is the number of generated samples. The convergence properties of MCEM are analyzed in Neath (2013).

Recall that the approximate distribution q_i is targeting the exact posterior $p(\mathbf{c}_i | \mathbf{e}_i, \boldsymbol{\theta}^{\text{old}})$. So ideally, we should generate the samples \mathbf{c}_i^s from the exact posterior. For this, we use Hamilton Monte Carlo (HMC) (Neal, 2011) as in Kingma and Welling (2013). HMC is a Markov chain Monte Carlo (MCMC) method that can be used to generate (correlated) samples from some target distribution (Martino & Elvira, 2017). More specifically, a Markov chain is generated whose stationary distribution coincides with the target distribution. By running the Markov chain for a sufficiently large number of steps one can obtain an approximate independent sample from $p(\mathbf{c}_i | \mathbf{e}_i, \boldsymbol{\theta}^{\text{old}})$. HMC has the advantage, when well-tuned, reduces substantially the correlation among samples (Martino & Elvira, 2017). For this, it simulates a dynamical system that uses information about the gradient of the posterior, i.e., $\nabla_{\mathbf{c}_i} \log p(\mathbf{c}_i | \mathbf{e}_i, \boldsymbol{\theta}^{\text{old}})$, to sample from regions of high posterior probability. In our implementation of MCEM, the HMC procedure consists of 20 leapfrog steps with small step-size (i.e., 5×10^{-4}) which guarantees that the acceptance rate is high enough. In practice, we only use just one sample to

approximate (13). Each time, the Markov chain is initialized at the mode of the posterior distribution, which is found using quasi-newton optimization methods (i.e., L-BFGS). Of course, after optimizing the prior parameters θ using MCEM, HMC can be used to generate samples from the approximate posterior distribution $\hat{p}(\mathbf{c}|\mathbf{e}, \theta) \propto p(\mathbf{c}|\mathbf{e})\hat{p}_\theta(\mathbf{c})$.

2.5 Important considerations

Note that both, the variational and MCEM methods, provide an estimation of the parameters θ of the prior. Thus, we obtain a Gaussian approximation of the prior, which is denoted here as $\hat{p}_\theta(\mathbf{c})$. Therefore, both techniques provide the following posterior approximation

$$\hat{p}(\mathbf{c}|\mathbf{e}, \theta) \propto p(\mathbf{c}|\mathbf{e})\hat{p}_\theta(\mathbf{c}). \quad (14)$$

However, the variational algorithm provides another posterior approximation given in Eq. (5), i.e.,

$$q(\mathbf{c}|\mathbf{e}) = \mathcal{N}(\mathbf{c}|\boldsymbol{\mu}_{\text{NN}}(\mathbf{e}), \boldsymbol{\Sigma}_{\text{NN}}(\mathbf{e})), \quad (15)$$

which yields an important advantage with respect the previous one: given one \mathbf{e} , using $q(\mathbf{c}|\mathbf{e})$ we can easily, and at low computational cost, produce a predictive mean $\boldsymbol{\mu}_{\text{NN}}(\mathbf{e})$ and covariance $\boldsymbol{\Sigma}_{\text{NN}}(\mathbf{e})$. The approximation $\hat{p}(\mathbf{c}|\mathbf{e}, \theta)$ on the other hand would require the use of additional Monte Carlo schemes for obtaining a predictive mean and variance, for each new observation vector \mathbf{e} . Another advantage of the variational approach is the computational speed compared to the MCEM method. However, one advantage of the MCEM scheme is that it can directly handle more practical scenarios (e.g., problems involving multiple posterior modes, heavy tailed distributions, etc.) leading to better performance in terms of smaller error in the parameter estimation of the prior. The variational approach described here would require a different and more general derivation for addressing these scenarios, see e.g. Mescheder et al. (2017). These features of each method are confirmed by the results obtained in our experiments. Implementations of the two approaches can be found at github.com/dhsvendsen/rtm_vi_mcem_inference.

3 Experiments

We illustrate the strengths and weaknesses of the two approaches, first by means of informative toy experiments: One that studies the computational efficiency of the respective methods, and another which analyzes their ability to handle forward models leading to multimodal posteriors. Following this, we show how these approaches can be used to perform inference over biophysical parameters using an RTM as the forward model.

3.1 On the computational efficiency

In order to analyze the computational efficiency of the two approaches, we consider a simple forward model

$$\mathbf{f}(\mathbf{c}) = f(c_1, c_2) = [2c_1, 2c_2],$$

for which both approaches converge to the true values of the parameters of the prior. We draw the training data $\{\mathbf{c}_n\}_{n=1}^N$ from the prior




Fig. 2 Marginal log-likelihood of test dataset as a function of training time for different inference methods. Four different sizes of training datasets are used, showing that the VI method is computationally more efficient than the MCEM method for larger datasets

$$p(\mathbf{c}) = \mathcal{N}\left(\begin{bmatrix} 4 \\ 6 \end{bmatrix}, \begin{bmatrix} 1 & 0.6 \\ 0.6 & 1 \end{bmatrix}\right),$$

and pass it through the nonlinear mapping \mathbf{f} in order to generate the training data $\{\mathbf{e}_n\}_{n=1}^N$. Datasets of several sizes $N = \{50, 500, 1000, 2000\}$ are used for training the models. The model likelihood noise is in all experiments fixed at a negligible value, with $\sigma^2 = 10^{-7}$, in order to reflect the trust in the knowledge encoded in the RTMs.

In Fig. 2 we plot an estimate of the average log marginal likelihood of each method on a test dataset as a function of training time (averaged over 40 repetitions). The marginal likelihood is computed using the estimator described in Appendix A. Observing the test log-likelihood, which is computed after each epoch, we see that the MCEM method converges after 1 epoch (one iteration of the E and M steps). With a training dataset of 50 points, each epoch of training is sufficiently fast that a parallelized version of the MCEM method (in which each E step is done in parallel) converges faster than the VI method. For the non-parallelized algorithm, this is not the case. For larger datasets, VI converges before the completion of 1 epoch of the MCEM algorithm. Since this is a simple toy problem, larger learning rates can be used in the VI method, leading to earlier convergence (just after 1 epoch) for datasets of 1000 and 2000 points. We can conclude from these experiments that the VI approach (as a consequence of stochastic optimization) has a better scaling properties with respect to the dataset size than MCEM.

3.2 Dealing with multimodal posteriors

We have seen that when faced with sufficiently large amounts of training data, variational inference performs faster than Monte Carlo sampling methods. However, since




Fig. 3 Contour plots of samples from the true posterior conditioned on the observation $\mathbf{e} = [4, 4]^T$. HMC samples using the prior parameters learned by the MCEM method shown in blue, and samples from the learned variational posterior in orange. The density (left) is so sharply peaked around the two modes that it is more informative to study the log-density (right)

the form of the variational posterior assumed in Eq. (5) is unimodal, we cannot expect it to be able to capture any multimodality in the true posterior. Consider for instance the forward mapping (with $\mathbf{c} = [c_1, c_2]$),

$$\mathbf{f}(\mathbf{c}) = [c_1^2, c_1 c_2].$$

For a given observed $\mathbf{e} = [e_1, e_2] \geq [0, 0]$ there will always be two possible solutions, namely $\mathbf{c}^{(1)} = [\sqrt{e_1}, e_2/\sqrt{e_1}]$ and $\mathbf{c}^{(2)} = [-\sqrt{e_1}, -e_2/\sqrt{e_1}]$ making the posterior inherently multimodal. As stated in the previous sections, we consider $\epsilon \sim \mathcal{N}(\epsilon|\mathbf{0}, 10^{-7}\mathbf{I})$. In this example, the prior density is Gaussian with parameters

$$p(\mathbf{c}) = \mathcal{N}\left(\begin{bmatrix} 1 \\ 2 \end{bmatrix}, \begin{bmatrix} 1 & 0.6 \\ 0.6 & 1 \end{bmatrix}\right),$$

from which 500 samples are drawn and passed through \mathbf{f} to generate the training dataset.

In the process of maximizing the ELBO, the expected log-likelihood with respect to the variational posterior is computed. We can see from Fig. 3, however, that the variational posterior, upon convergence, only captures the positive mode at $\mathbf{c}' = [2, 2]^T$ of the true posterior given the observation $\mathbf{e}' = [4, 4]^T$. On the other hand, the MCEM algorithm computes the expected complete log-likelihood with respect to the true posterior as approximated with HMC. As opposed to the variational posterior, HMC does manage to capture both the modes of the true posterior as shown in Fig. 3. The learning algorithm of the MCEM method is therefore more likely to converge to the true parameters of the prior if the posterior is multimodal.

We can see the inability of the variational method to capture the multimodality of the problem from the results of the converged methods given in Table 1. The fitted parameters of the prior are far from the true ones when compared to the results of the MCEM method which as also reflected in the KL divergence between the fitted and true prior distributions. Multimodality such as this is likely to be observed in the remote sensing experiment latter, as it has been remarked before that different configurations of inputs can lead to the same output making it an *ill-posed inversion problem* (Gómez-Dans et al., 2016).

Table 1 Comparison of methods for inference on a forward model which leads to a bimodal posterior

Method	VI	MCEM
Mean	$\begin{bmatrix} 1.20 \\ 1.76 \end{bmatrix}$	$\begin{bmatrix} 1.027 \\ 2.061 \end{bmatrix}$
Covariance	$\begin{bmatrix} 0.619 & 0.682 \\ 0.682 & 1.600 \end{bmatrix}$	$\begin{bmatrix} 1.001 & 0.617 \\ 0.617 & 0.945 \end{bmatrix}$
D_{KL}	0.315	0.00581

The first and second columns show the estimated mean vector and covariance matrix of the prior. The third column shows the KL divergence between the fitted and the true prior

3.3 PROSAIL experiment

We now turn to inference in a remote sensing setting using one of the most widely used RTM over the last almost three decades in the field as our physical forward model (Jacquemoud et al., 2009). PROSAIL is a canopy reflectance model which allows us to relate fundamental vegetation canopy properties, such as, the Leaf Area Index (LAI), and leaf chemical and structural properties, to the scene reflectance for a given set of illumination and sensor (observation) geometry conditions (Liang, 2005). To perform its simulations, PROSAIL combines two sub-models: PROSPECT (Feret et al., 2008), which models the optical properties of the leaves; and SAIL (Verhoef, 1984), which models bidirectional reflectances considering the scattering by arbitrarily inclined canopy leaves in a turbid medium (Fang et al., 2019). This combination of models requires the following set of input parameters:

- 1) A set of leaf optical properties (PROSPECT), given by the mesophyll structural parameter (N), leaf chlorophyll (Chl), dry matter (Cm), water (Cw), carotenoid (Car) and brown pigment (Cbr) contents.
- 2) A set of canopy level and geometry characteristics (SAIL), determined by leaf area index (LAI), the average leaf angle inclination (ALA), the hot-spot parameter (Hotspot), the solar zenith angle (θ_s), view zenith angle (θ_v), and the relative azimuth angle between both angles ($\Delta\Theta$).

We consider PROSAIL for simulating Landsat-8 spectra. This satellite has been widely used in many applications such as cryosphere monitoring, aquatic science and surface water mapping, and vegetation monitoring (Wulder et al., 2019). Landsat 8's Operational Land Imager (OLI) includes nine spectral bands with wavelengths ranging from 0.433 to 1.390 μm , leaving us with an output-dimension of $D_e = 9$ for our problem. In our experimental setup, we have chosen to work with the most relevant leaf-level parameters to monitor vegetation status and functioning included in PROSAIL, namely Cw, Cm and Chl, resulting in an input dimension of $D_c = 3$. The remaining parameters were set constant during our experiments and their values were obtained from previous studies (Svendsen et al., 2020) to be representative of realistic cases. Their values can be found in Table 2.

Constraining the radiative transfer models with realistic and representative distributions of their inputs is a key part of the RTM inversion process. To facilitate this, in this work we relied on the largest global plant traits database available, the TRY

Table 2 Characteristics of the simulations using the PROSAIL model

<i>Leaf</i>	N	Car	Cbr			
	1.5	8 g/cm ²	0			
<i>Canopy</i>	ALA	Hotspot	θ_s	θ_v	$\Delta\Theta$	LAI
	Spherical	0.01	30°	10°	0	4

database (Kattge et al., 2011, 2020), which contains thousands of leaf data records measured at unprecedented spatial and climatological coverage. Using these data we computed the following empirical mean vector and covariance matrix which was used to sample 2000 values of \mathbf{c} and pass them through PROSAIL to generate the training data. The empirical mean and covariance (to be compared with the results in Table 3) of the samples are

$$\hat{\mathbf{m}} = \begin{bmatrix} 9.76e^3 \\ 1.77e^2 \\ 46.2 \end{bmatrix}, \hat{\mathbf{S}} = \begin{bmatrix} 6.42e^5 & 5.06e^5 & 3.68e^2 \\ 5.06e^5 & 1.34e^4 & 2.86e^3 \\ 3.68e^2 & 2.86e^3 & 288 \end{bmatrix}.$$

The units of the parameters are g/cm² for Cm and Cw, and μg/cm² for Chl respectively. Note that the ground truth prior estimated from the TRY database has some probability density in the negative region of parameter space. This is not physically meaningful, but serves the point of illustrating the capabilities of the inference methods. We alter PROSAIL so that it sets every negative parameter to 0 before mapping into spectral space to get a modified likelihood that will lead to more multimodality (since all negative values in \mathbf{c} will be mapped into the same value, i.e. 0, and then through PROSAIL into a spectrum).

The results of the variational approach to inference over PROSAIL are summarized in Fig. 4. We see that the parameters of the prior are fitted well, which can also be confirmed in Table 3 quantitatively, even though the variational posterior is not able to produce predictive means in the negative domain. It is interesting to note that the modification of PROSAIL to truncate negative data, which leads to multimodality, does not prevent the variational approach from estimating the parameters of the prior well.

Nevertheless, the MCEM method is somewhat more accurate than the VI method, obtaining a KL divergence with to the true prior of 1.23×10^{-2} compared to 2.08×10^{-2} obtained using the VI approach. This is to be expected since, as we have seen, the MCEM approach handles multimodality better. We especially foresee a clear difference in results in future work the LAI variable which is difficult to estimate due to its multimodal posterior distribution as pointed out elsewhere (Gómez-Dans et al., 2016).

Once the VI method has converged, the neural network which parameterizes the variational posterior can be used as a fast inverse model that maps from observed satellite spectra to biophysical variables. Using the mean outputs that model the mean value of the variational posterior we can obtain good predictive accuracy on a test set as shown in the scatter plots of Fig. 5. Despite the promising results, it is very important to note that we run our experiments using a simplified PROSAIL configuration, keeping some of the input parameters static (see Table 2) and that results can vary greatly in more realistic modeling scenarios.

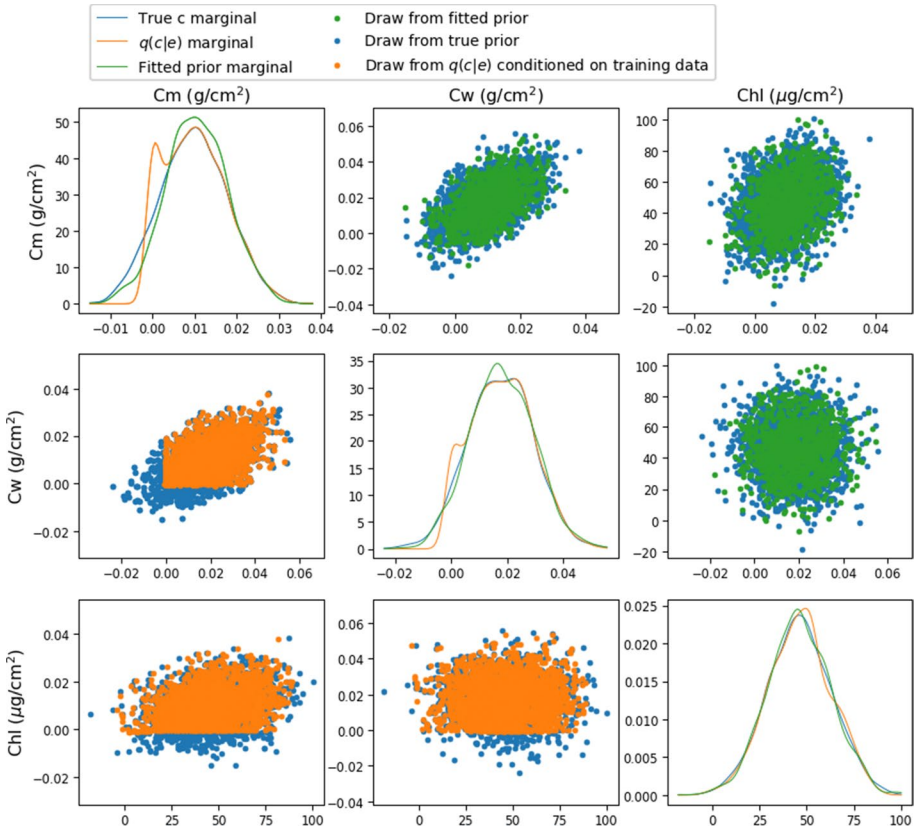


Fig. 4 Results of the variational approach to inference over PROSAIL. The blue points are c 's from the training set, while the green points are draws from the fitted prior. The orange points are draws from the variational posterior conditioned on the training e 's. The diagonal shows KDE plots of c using samples from the ground truth prior (blue), the variational posterior conditioned on training data (orange) and the fitted prior (green)

Table 3 Comparison of methods for inference on biophysical parameters using a radiative transfer forward model

Method	VI	MCEM
Mean	$\begin{bmatrix} 1.02e^2 \\ 1.80e^2 \\ 45.9 \end{bmatrix}$	$\begin{bmatrix} 1.01e^2 \\ 1.81e^2 \\ 46.6 \end{bmatrix}$
Covariance	$\begin{bmatrix} 5.32e^5 & 4.31e^5 & 3.52e^2 \\ 4.31e^5 & 1.19e^4 & 2.18e^3 \\ 3.52e^2 & 2.18e^3 & 280 \end{bmatrix}$	$\begin{bmatrix} 5.49e^5 & 4.45e^5 & 3.31e^2 \\ 4.45e^5 & 1.25e^4 & 1.74e^3 \\ 3.31e^2 & 1.74e^3 & 292 \end{bmatrix}$
D_{KL}	0.0208	0.0123

The first and second columns show the mean vector and covariance matrix respectively of the true and the estimated prior over the causes, using $e = \times 10$ notation for space. The third column shows the Kullback–Leibler divergence between the fitted and the true prior




Fig. 5 True values of RTM parameters in test dataset versus mean of variational posterior conditioned on spectra in test dataset. The trained encoder network can thus be used as an effective predictive model

4 Discussion and conclusions

In this work, we approached the long-standing inverse problem in remote sensing of estimating biophysical parameters from observational reflectances. Unlike previous works, we focus on estimating not only the particular parameter point estimates but its full multivariate distribution. We evaluated two different approximations that include an RTM forward model to enforce the inverse estimations to be physically consistent.

Both proposed techniques have different advantages and shortcomings that we illustrated with toy examples and with simulations from the PROSAIL RTM. The MCEM-based approach admits more flexible models while the VAE is computationally more efficient. For instance, while MCEM deals easily with multimodal distributions, this is a challenge for VAE. On the other hand, the convergence time of VAE is orders of magnitude faster depending on the problem. Moreover, the VAE scheme provides a posterior approximation, with a predictive mean and a covariance matrix, implicitly defined by the trained neural network that can be readily evaluated. The experiment involving PROSAIL shows that, while the accuracy of the VAE and MCEM are deemed similar, the computational simplicity of the VAE approach is critical in this problem. Note that including the RTM PROSAIL in the forward-inverse modeling loop increases the time of computation and combining it with MCEM makes it unfeasible especially for large data sets.

We anticipate a wide interest in these techniques for inferring the parameter densities from simulations and then, as further work, from observational satellite data. This will require more accurate and realistic priors; for this we plan to explore mixtures of Gaussians for modeling the prior of the causes as a generalization of the simplified Gaussian model assumed in this work. Likewise, more sophisticated computational methods and variational approaches (e.g., Bugallo et al., 2015; Martino & Elvira, 2017; Mescheder et al., 2017) could be explored in the future.

Finally, we are well aware of the fact that this problem is ubiquitous in other domains of Earth observation and geosciences, and may have implications in climate science too. Inferring parameters is a transversal important topic, not only attached to terrestrial biosphere processes but to the atmosphere, cryosphere and the ocean modeling too. For instance parametrization of small-scale processes such as clouds or biological processes (that are important at the land surface for the exchange of energy and carbon) cannot be explicitly resolved. In this context, learning appropriate parametrizations directly from

data may reduce the sources of uncertainties in current models, eventually leading to further advances in climate modeling.

A Marginal likelihood Estimation by Reverse Importance Sampling

In order to evaluate the performance of the methods described in the paper, we use an estimator of marginal likelihood on a test-dataset. More precisely, we use the Reverse Importance Sampling (RIS) estimator (Llorente et al., 2019) described below:

1. Sample L values $\{\mathbf{c}_l\}_{l=1}^L$ from the posterior with an MCMC-method. We use Hamiltonian Monte Carlo.
2. Fit a density estimator $q(\mathbf{c})$ to the samples $\{\mathbf{c}_l\}_{l=1}^L$. In this work we fit a Gaussian mixture model, doing cross validation in order to find the best number of components.
3. Sample M new values $\{\mathbf{c}_m\}_{m=1}^M$ from the posterior to be inserted in the following estimator:

$$p(\mathbf{e}) \simeq \left(\frac{1}{M} \sum_{m=1}^M \frac{q(\mathbf{c}_m)}{p(\mathbf{c}_m)p(\mathbf{e}|\mathbf{c}_m)} \right)^{-1} \quad \text{where } \mathbf{c}_m \sim p(\mathbf{c}|\mathbf{e}).$$

For the proof and more details see Llorente et al. (2019). It is important to remark that the function $q(\mathbf{c})$ must be a valid probability density for which there are several possible choices. If one chooses $q(\mathbf{c}) = p(\mathbf{c})$, RIS becomes the so-called harmonic mean estimator (this name is due to the fact that the corresponding estimator is the harmonic mean of the likelihood values). However, it has been shown that this does not lead to a good estimator. It is possible to show that, in order to ensure finite variance of the resulting estimator, the density $q(\mathbf{c})$ should have equal or lighter tails than the posterior $p(\mathbf{c}|\mathbf{e})$ [e.g., see first numerical example in Llorente et al. (2019)]. Gaussian mixture approximations and kernel density estimators of $p(\mathbf{c}|\mathbf{e})$ are suitable choices for $q(\mathbf{c})$. Different alternative estimators of the marginal likelihood are possible mixing MCMC and importance sampling schemes [see Llorente et al. (2019) and Martino et al. (2017)].

Acknowledgements Research partly funded by the ERC under the ERC-CoG-2014 project (Grant Agreement 647423) and the ERC-SyG-2019 USMILE Project (Grant Agreement 855187), the Ministerio de Ciencia e Innovación Project PID2019-106827GB-I00/AEI/10.13039/501100011033, and the Projects TEC2016-77741-R, DPI2017-89867-C2-2-R, RTI2018-096765-A-I00, PID2019-109026RB-I00 and PID2019-105032GB-I00

References

- Atzberger, C., & Richter, K. (2012). Spatially constrained inversion of radiative transfer models for improved LAI mapping from future sentinel-2 imagery. *Remote Sensing of Environment*, 120, 208–218.
- Bacour, C., Jacquemoud, S., Leroy, M., Hauteceur, O., Weiss, M., Prévot, L., et al. (2002). Reliability of the estimation of vegetation characteristics by inversion of three canopy reflectance models on airborne POLDER data. *Agronomie*, 22(6), 555–565.
- Baret, F., et al. (2007). LAI, fAPAR and fCover cyclopes global products derived from vegetation: Part 1: Principles of the algorithm. *Remote Sensing of Environment*, 110(3), 275–286.
- Bugallo, M., Martino, L., & Corander, J. (2015). Adaptive importance sampling in signal processing. *Digital Signal Processing*, 47, 36–49.

- Campos-Taberner, M., Moreno-Martínez, A., García-Haro, F. J., Camps-Valls, G., Robinson, N. P., Kattge, J., & Running, S. W. (2018). Global estimation of biophysical variables from google earth engine platform. *Remote Sensing*, *10*(8), 1167.
- Camps-Valls, G., Martino, L., Svendsen, D. H., Campos-Taberner, M., Munoz-Mari, J., Laparra, V., et al. (2018). Physics-aware Gaussian processes in remote sensing. *Applied Soft Computing*, *28*, 69–82.
- Camps-Valls, G., Runge, J., Sejdinovic, D., & Reichstein, M. (2019). A perspective on Gaussian processes for earth observation. *National Science Review*, *6*, 616–618.
- Coccia, G., Siemann, A. L., Pan, M., & Wood, E. F. (2015). Creating consistent datasets by combining remotely-sensed data and land surface model estimates through Bayesian uncertainty post-processing. The case of land surface temperature from HIRS. *Remote Sensing of Environment*, *170*, 290–305.
- Combal, B., Baret, F., Weiss, M., Trubuil, A., Mace, D., Pragnere, A., et al. (2003). Retrieval of canopy biophysical variables from bidirectional reflectance: Using prior information to solve the ill-posed inverse problem. *Remote Sensing of Environment*, *84*(1), 1–15.
- Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, *39*, 1–22.
- Djamai, N., Fernandes, R., Weiss, M., McNairn, H., & Goita, K. (2019). Validation of the sentinel simplified level 2 product prototype processor (SL2P) for mapping cropland biophysical variables using sentinel-2/MSI and landsat-8/OLI data. *Remote Sensing of Environment*, *225*, 416–430.
- Fang, H., Baret, F., Plummer, S., & Schaepman-Strub, G. (2019). An overview of global leaf area index (LAI): Methods, products, validation, and applications. *Reviews of Geophysics*, *57*(3), 739–799.
- Feret, J.-B., François, C., Asner, G. P., Gitelson, A. A., Martin, R. E., Bidel, L. P. R., et al. (2008). Prospect-4 and 5: Advances in the leaf optical properties model separating photosynthetic pigments. *Remote Sensing of Environment*, *112*(6), 3030–3043.
- Frédéric, B., Jacquemoud, S., Guyot, G., & Leprieux, C. (1992). Modeled analysis of the biophysical nature of spectral shifts and comparison with information content of broad bands. *Remote Sensing of Environment*, *41*(2–3), 133–142.
- Gómez-Dans, J. L., Lewis, P. E., & Disney, M. (2016). Efficient emulation of radiative transfer codes using Gaussian processes and application to land surface parameter inferences. *Remote Sensing*, *8*(2), 119.
- Jacquemoud, S., & Baret, F. (1990). Prospect: A model of leaf optical properties spectra. *Remote Sensing of Environment*, *34*(2), 75–91.
- Jacquemoud, S., Verhoef, W., Baret, F., Bacour, C., Zarco-Tejada, P. J., Asner, G. P., et al. (2009). PROSPECT+ SAIL models: A review of use for vegetation characterization. *Remote Sensing of Environment*, *113*, S56–S66.
- Kattge, J., Bönsch, G., Díaz, S., Lavorel, S., Prentice, I. C., Leadley, P., et al. (2020). Try plant trait database—enhanced coverage and open access. *Global Change Biology*, *26*, 119–188.
- Kattge, J., Díaz, S., Sandra Lavorel, I., Prentice, C., Leadley, P., Bönsch, G., et al. (2011). Try-a global database of plant traits. *Global Change Biology*, *17*(9), 2905–2935.
- Kingma, D. P., & Ba, J. (2015). ADAM: A method for stochastic optimization. *International conference on learning representations* (pp. 1–15).
- Kingma, D.P., & Welling, M. (2013) Auto-encoding variational Bayes. arXiv preprint [arXiv:1312.6114](https://arxiv.org/abs/1312.6114).
- Knyazikhin, Y. (1999). *MODIS leaf area index (LAI), and fraction of photosynthetically active radiation absorbed by vegetation FPAR*. GSFC/NASA Technical report <http://eosps0.gsfc.nasa.gov/atbd/modis/tables.html>.
- Lewis, P., Gómez-Dans, J., Kaminski, T., Settle, J., Quaife, T., Gobron, N., et al. (2012). An earth observation land data assimilation system (EO-LDAS). *Remote Sensing of Environment*, *120*, 219–235.
- Liang, S. (2005). *Quantitative remote sensing of land surfaces* (Vol. 30). Hoboken: Wiley.
- Liang, S. (2008). *Advances in land remote sensing: System, modeling, inversion and applications*. Berlin: Springer.
- Liang, L., Di, L., Zhang, L., Deng, M., Qin, Z., Zhao, S., & Lin, H. (2015). Estimation of crop LAI using hyperspectral vegetation indices and a hybrid inversion method. *Remote Sensing of Environment*, *165*, 123–134.
- Lillesand, T. M., Kiefer, R. W., & Chipman, J. (2008). *Remote sensing and image interpretation*. New York: Wiley.
- Llorente, F., Martino, L., Delgado, D., & Lopez-Santiago, J. (2019) Marginal likelihood computation for model selection and hypothesis testing: An extensive review (pp. 1–58). *viXra:2001.0052*.
- Ma, C., Li, X., Notarnicola, C., Wang, S., & Wang, W. (2017). Uncertainty quantification of soil moisture estimations based on a Bayesian probabilistic inversion. *IEEE Transactions on Geoscience and Remote Sensing*, *55*(6), 3194–3207.
- Martino, L., & Elvira, V. (2017). *Metropolis sampling* (pp. 1–15). Wiley StatsRef: Statistics Reference Online.

- Martino, L., Elvira, V., Luengo, D., & Corander, J. (2017). Layered adaptive importance sampling. *Statistics and Computing*, 27(3), 599–623.
- McCarthy, A., Rodriguez, B., & Mincholé, A. (2017). Variational inference over non-differentiable cardiac simulators using Bayesian optimization. arXiv preprint [arXiv:1712.03353](https://arxiv.org/abs/1712.03353).
- Mescheder, L., Nowozin, S., & Geiger, A. (2017). Adversarial variational Bayes: Unifying variational autoencoders and generative adversarial networks. In *International conference on machine learning* (pp. 2391–2400).
- Neal, R. M., et al. (2011). MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2(11), 2.
- Neath, R. C. (2013). On convergence properties of the Monte Carlo EM algorithm. In *Advances in modern statistical theory and applications: A festschrift in honor of Morris L. Eaton* (pp. 43–62).
- Pinty, B., Andreadakis, I., Clerici, M., Kaminski, T., Taberner, M., Verstraete, M. M., Gobron, N., Plummer, S., & Widowski, J.-L. (2011). Exploiting the MODIS albedos with the two-stream inversion package (JRC-TIP): 1. Effective leaf area index, vegetation, and soil properties. *Journal of Geophysical Research: Atmospheres*, 116(D9).
- Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2017). Machine learning of linear differential equations using Gaussian processes. *Journal of Computational Physics*, 348, 683–693.
- Robbins, H., & Monro, S. (1951). A stochastic approximation method. *The Annals of Mathematical Statistics*, 22, 400–407.
- Robert, C., & Casella, G. (2013). *Monte Carlo statistical methods*. Berlin: Springer.
- Rodgers, C. D. (2000). *Inverse methods for atmospheric sounding: Theory and practice*. Singapore: World Scientific Publishing Co. Ltd.
- Sandberg, A. (2013). Feasibility of whole brain emulation. In *Philosophy and theory of artificial intelligence* (pp. 251–264). Springer.
- Santner, T. J., Williams, B. J., & Notz, W. I. (2003). *The design and analysis of computer experiments*. Berlin: Springer.
- Svendsen, D. H., Martino, L., Campos-Taberner, M., García-Haro, F. J., & Camps-Valls, G. (2018). Joint Gaussian processes for biophysical parameter retrieval. *IEEE Transactions on Geoscience and Remote Sensing*, 56(3), 1718–1727.
- Svendsen, D. H., Martino, L., & Camps-Valls, G. (2020). Active emulation of computer codes with Gaussian processes: Application to remote sensing. *Pattern Recognition*, 100, 107103.
- Tarantola, A. (2005). *Inverse problem theory and methods for model parameter estimation* (Vol. 89). Philadelphia: SIAM.
- Verhoef, W. (1984). Light scattering by leaf layers with application to canopy reflectance modeling: The sail model. *Remote Sensing of Environment*, 16(2), 125–141.
- Verstraete, M. M., & Pinty, B. (1996). Potential and limitations of information extraction the terrestrial biosphere from satellite remote sensing. *Remote Sensing of Environment*, 58, 201–214.
- Vicent, J., Alonso, L., Martino, L., Sabater, N., Verrelst, J., Camps-Valls, G., & Moreno, J. (2019). Gradient-based automatic look-up table generator for radiative transfer models. *IEEE Transactions on Geoscience and Remote Sensing*, 57(2), 1040–1048.
- Wei, G. C. G., & Tanner, M. A. (1990). A Monte Carlo implementation of the EM algorithm and the poor man's data augmentation algorithms. *Journal of the American Statistical Association*, 85, 699–704.
- Wescott, B. (2013). *Every computer performance book* (1st ed.). Create Space Independent Publishing Platform, USA.
- Wulder, M., et al. (2019). Remote sensing of environment: Current status of landsat program, science, and applications. *Remote Sensing of Environment*, 225, 127–147.
- Zhang, Q., Xiao, X., Braswell, B., Linder, E., Baret, F., & Berrien Moore, I. I. I. (2005). Estimating light absorption by chlorophyll, leaf and canopy in a deciduous broadleaf forest using MODIS data and a radiative transfer model. *Remote Sensing of Environment*, 99(3), 357–371.
- Zurita-Milla, R., Laurent, V., & van Gijsel, J. (2015). Visualizing the ill-posedness of the inversion of a canopy radiative transfer model: A case study for sentinel-2. *International Journal of Applied Earth Observation and Geoinformation*, 43, 7–18.