Parametric 3D Atmospheric Reconstruction in Highly Variable Terrain with Recycled Monte Carlo Paths and an Adapted Bayesian Inference Engine

Ian Langmore[a], Anthony B. Davis[b], Guillaume Bal[a], and Youssef M. Marzouk[c]

[a] *Columbia University, Department of Applied Physics & Applied Mathematics, 500 W. 120th Street, New York, NY, 10027, USA*
[b] *Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Mail Stop 233-200, Pasadena, Ca, 91109, USA*
[c] *Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, 77 Massachusetts Avenue, Cambridge, MA, 02139, USA*

Abstract. We describe a method for accelerating a 3D Monte Carlo forward radiative transfer model to the point where it can be used in a new kind of Bayesian retrieval framework. The remote sensing challenge is to detect and quantify a chemical effluent of a known absorbing gas produced by an industrial facility in a deep valley. The available data is a single low-resolution noisy image of the scene in the near IR at an absorbing wavelength for the gas of interest. The detected sunlight has been multiply reflected by the variable terrain and/or scattered by an aerosol that is assumed partially known and partially unknown. We thus introduce a new class of remote sensing algorithms best described as "multi-pixel" techniques that call necessarily for a 3D radaitive transfer model (but demonstrated here in 2D); they can be added to conventional ones that exploit typically multi- or hyper-spectral data, sometimes with multi-angle capability, with or without information about polarization. The novel Bayesian inference methodology uses adaptively, with efficiency in mind, the fact that a Monte Carlo forward model has a known and controllable uncertainty depending on the number of sun-to-detector paths used.

Keywords: 3D radiative transfer, variable terrain, aerosol optics, gaseous absorption, gas effluents, remote sensing, Monte Carlo methods, path recycling techniques, Bayesian techniques, Markov Chain Monte Carlo algorithms.
PACS: 42.68.Ay, 42.68.Ca, 42.68.Sq, 42.68.Wt, 92.60.H-, 92.60.Kc, 92.60.Ta

INTRODUCTION, CONTEXT & MOTIVATION

Physics-based retrievals of atmosphere and/or surface properties are generally multi- or hyper-spectral in nature; some use multi-angle information as well. Recently, polarization has been added to the available input from sensors and accordingly modeled with vector radiative transfer (RT). At any rate, a single pixel is processed at a time using a forward RT model predicated on 1D transport theory. Neighboring pixels are sometimes considered but, generally, just to formulate statistical constraints on the inversion based on spatial context. Here, we demonstrate the power that could be harnessed by adding bona fide multi-pixel techniques to the toolbox. This report is a digest of a full-length paper by Langmore, Davis and Bal [1]. The geometry of the hypothetical 2D remote sensing problem is described in Fig. 1 (left panel) and in Table 1. It is inspired by observations that could be used for certain nuclear treaty verification applications looking for signatures of covert processing. They key challenge comes from the highly variable terrain.

We thus need a forward RT model in 2D—sufficient for this demo, and easily extended to 3D—for the response of a single-wavelength imaging sensor. The remote sensing data is a single image, as plotted in Fig. 1 (right panel). It is used to infer position, size and opacity of an absorbing atmospheric plume somewhere in the valley in the presence of a partially-known/partially-unknown aerosol. The first necessary innovation we describe is the speed-up of a forward 2D RT model that predicts images with cross-pixel radiative fluxes fully accounted for. In spite of its reputation for inefficiency, we use a Monte Carlo technique. However, the adopted scheme is highly accelerated without loss of accuracy by using "recycled" Monte Carlo paths, a parametric technique recently used in biomedical imaging that is akin to dependent sampling methods [2]. This forward model is then put to work in a new kind of Bayesian inversion adapted to this kind of RT model where it is straightforward to trade precision and efficiency. Retrievals target the plume properties and the specific amount of aerosol. In spite of the limited number of pixels (namely, 15) and low signal-to-noise-ratio (SNR $\approx$5), we show that multi-pixel processing adds value to existing multi- or hyper-spectral remote sensing programs, and we suggest further developments that could be focused on specific sponsor interests.

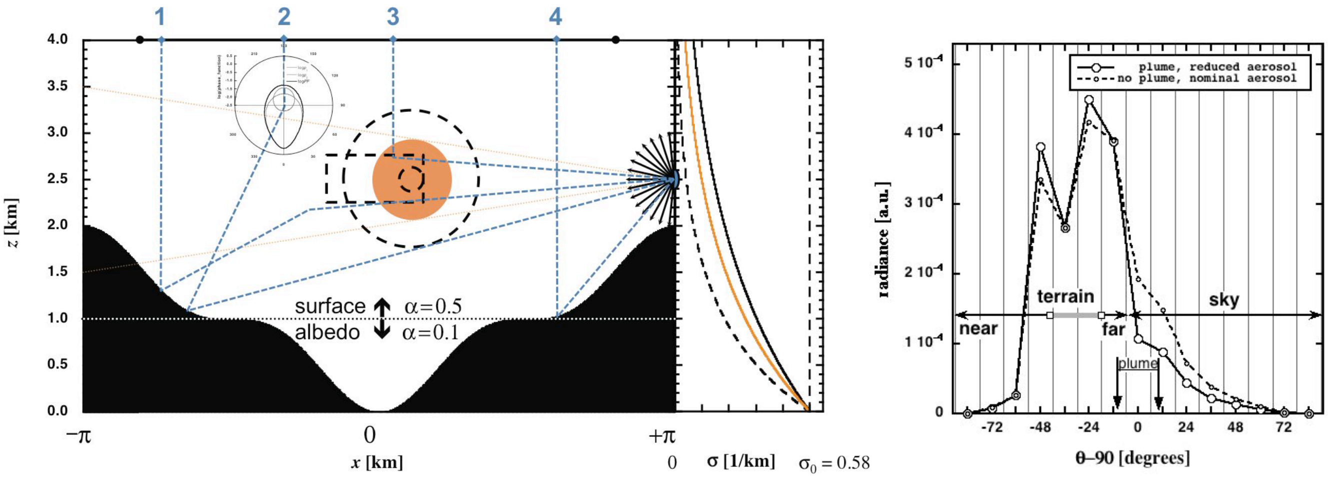

FIGURE 1. Left: Schematic of the parameterized 2D scene and observation described in the main text. Bayesian extremal values for the plume geometry and aerosol profile are plotted with black dashes; the baseline aerosol is in solid black while the "true" plume and aerosol are rendered in an orange hue. Four typical Monte Carlo "particle" trajectories from the (solar) source to the imaging sensor are in blue: some intersect the actual (orange) plume, as well as many potential plumes, one grazes a number of possible large ones, and one intersects no plumes whatsoever. **Right:** Two flatland "images," which are 1D functions: without (---) and with (—) the actual plume and $-^{\mathrm{ve}}$ aerosol perturbation present; representative line-of-sight features are indicated (near/far terrain, plume, sky).

FAST MONTE CARLO COMPUTATION WITH RECYCLED PATHS ILLUSTRATED WITH THE PARAMETERIZED 2D SCENE

2D RT Equation and Associated Boundary Conditions (BCs)

Let $\boldsymbol{x} = (x,z)^{\mathrm{T}}$ be position in the 2D domain M = $\{\boldsymbol{x} \in (0,H_{\mathrm{sky}})\times(-L/2,+L/2);\ z/H_{\mathrm{ter}} > 1-\cos^3(2\pi x/L)\}$, i.e., limited below by a "1–cos^3" terrain, and let $\boldsymbol{\Omega}(\theta) = (\cos\theta,\sin\theta)^{\mathrm{T}}$ be a direction reckoned clockwise from the upward z-axis. We take $H_{\mathrm{sky}} = 2H_{\mathrm{ter}} = 4$ km and $L = 2\pi$ km; see Fig. 1 (left panel). We seek numerical estimates of certain values of the radiance field $I(\boldsymbol{x},\theta)$, which obeys the 2D RT equation in integro-differential form:

$$\left[\boldsymbol{\Omega}(\theta)\bullet\nabla + \sigma(\boldsymbol{x})\right] I = \sigma(\boldsymbol{x})\varpi(\boldsymbol{x})\int_{-\pi}^{+\pi} p_{\mathrm{v}}(\boldsymbol{x},\theta'-\theta)\, I(\boldsymbol{x},\theta')\mathrm{d}\theta' + Q_{\mathrm{v}}(\boldsymbol{x},\theta), \tag{1}$$

where $\sigma(\boldsymbol{x}) = \sigma_{\mathrm{aer}}(z)+\sigma_{\mathrm{gas}}(x,z)$ is the extinction coefficient, $\varpi(\boldsymbol{x}) = 0.9\,\sigma_{\mathrm{aer}}(z)/\sigma(\boldsymbol{x})$ is the single scattering albedo, $p_{\mathrm{v}}(\boldsymbol{x},\theta_{\mathrm{s}}) \equiv p_{\mathrm{v}}(\theta_{\mathrm{s}})$ is the scattering phase function, dependent only on the scattering angle θ_{s}, and $Q_{\mathrm{v}}(\boldsymbol{x},\theta)$ is the volume source term, taken here to vanish identically; $I(\boldsymbol{x},\theta)$ thus models total (direct+diffuse) radiance. For the background aerosol, we take $\sigma_{\mathrm{aer}}(z) = \sigma_0\exp[-(c_0+\delta c)z]$ with $c_0 = 0.5$ km^{-1} (2 km scale height), δc being an unknown aerosol perturbation, $\sigma_0 = 0.58$ km^{-1}; this yields to an aerosol optical depth of unity when $\delta c = 0$. For $p_{\mathrm{v}}(\theta_{\mathrm{s}})$, we assume a double Henyey–Greenstein model for 2D RT, $f p_{\mathrm{HG}}(g_{\mathrm{f}},\theta_{\mathrm{s}})+(1-f)p_{\mathrm{HG}}(g_{\mathrm{b}},\theta_{\mathrm{s}})$ where $p_{\mathrm{HG}}(g,\theta_{\mathrm{s}}) = (1-g^2)/(1+g^2-2g\cos\theta_{\mathrm{s}})/2\pi$ [3]; by setting $(f,g_{\mathrm{f}}, g_{\mathrm{b}}) = (0.9,+0.8,-0.4)$ the phase function looks like that of a typical fine-mode particle population.

The RT equation in (1) is complemented by the appropriate BCs. For $\boldsymbol{x} \in \partial$M, the boundary of M, we have

$$\boldsymbol{\Omega}(\theta)\bullet\boldsymbol{n}(\boldsymbol{x})\, I = \alpha(\boldsymbol{x}) \int_{\theta_{\boldsymbol{n}(\boldsymbol{x})}-\pi/2}^{\theta_{\boldsymbol{n}(\boldsymbol{x})}+\pi/2} p_{\mathrm{s}}(\boldsymbol{x},\theta'\to\theta)\, I(\boldsymbol{x},\theta')\mathrm{d}\theta' + \boldsymbol{\Omega}(\theta)\bullet\boldsymbol{n}(\boldsymbol{x})\, Q_{\mathrm{s}}(\boldsymbol{x},\theta), \tag{2}$$

where $\alpha(\boldsymbol{x})$ is the local value of the surface albedo, $p_{\mathrm{s}}(\boldsymbol{x}, \theta_{\mathrm{inc}}\to\theta_{\mathrm{out}}) \equiv p_{\mathrm{s}}(\theta_{\mathrm{inc}}\to\theta_{\mathrm{out}})$ is the phase function for reflection (a.k.a. surface scattering), and $Q_{\mathrm{s}}(\boldsymbol{x},\theta)$ is the surface source term. We take: $\alpha(\boldsymbol{x}) \equiv \alpha(z) = 0.1$ for $0 \le z < H_{\mathrm{ter}}/2$, 0.5 for $H_{\mathrm{ter}}/2 \le z < H_{\mathrm{ter}}$, and 0 for $H_{\mathrm{ter}} \le z < H_{\mathrm{sky}}$; and diffusely reflecting surfaces are assumed Lambertian, which in 2D reads as $p_{\mathrm{s}}(\boldsymbol{x},\theta_{\mathrm{inc}}\to\theta_{\mathrm{out}}) \equiv \boldsymbol{\Omega}(\theta_{\mathrm{out}})\bullet\boldsymbol{n}(\boldsymbol{x})/2$. Lastly, we have $Q_{\mathrm{s}}(\boldsymbol{x},\theta) = \delta(\theta-\pi)$ for $z = H_{\mathrm{sky}}$, and 0 otherwise, for an overhead sun.

Apart from the aerosol property δc, the remote sensing unknowns define the purely absorbing chemical plume, which is parameterized as $\sigma_{\mathrm{gas}}(x,z) = k_{\mathrm{p}}$ if $(x-x_{\mathrm{p}})^2+(z-z_{\mathrm{p}})^2 < \rho_{\mathrm{p}}^2$, and 0 otherwise. The choice of the single wavelength captured by the small imaging detector illustrated in Fig. 1 is presumably informed by the known absorption spectrum of the molecule of interest. Table 1 summarizes the unknowns of the remote sensing problem, along with some attributes of their Bayesian prior distribution. Key combinations are its optical diameter $2k_{\mathrm{p}}\rho_{\mathrm{p}}$ and "mass" $\propto k_{\mathrm{p}}\rho_{\mathrm{p}}^2$.

TABLE 1. Summary of the unknown parameters in the simulated remote sensing problem.

Parameter to retrieve	Symbol	Reference	Truth	Min	Max	Unit	Prior	Mean	St. Dev.
Plume's *x*-position	x_p	arbitrary	+0.35	–0.5	+0.5	km	uniform	0	0.204
Plume's *z*-position	x_p	arbitrary	2.5	2.2	2.7	km	uniform	2.45	0.102
Plume's radius	ρ_p	arbitrary	0.5	0.15	0.85	km	uniform	0.5	0.143
Plume's absorption coefficient	k_p	0	0.5	0	∞	1/km	Gamma	0.25	$1/\sqrt{8}$
Aerosol perturbation	δc	0	+0.15	$-c_0$	$+c_0$	1/km	uniform	0	$c_0/\sqrt{6}$

The Path Recycling Technique for Circular Gaseous Plumes and a Varying Aerosol Load

In path-recycling Monte Carlo, as well as in dependent-trajectory/sampling methods, one revisits a given and fixed path traced in space, here 2D space, and computes the probabilities of what actually happened as well as the probability of some alternative based on a change of certain parameters of interest. The ratio of these probabilities is used to reweight the path, hence to recompute the predicted signal, in this case, the value of one of the 15 pixels in the image. The key difference between path recycling and dependent-trajectory methods is that, in the former case, all the paths of interest (i.e., those that end in detection) are stored on file, and this file is reread and reprocessed for each instantiation of the varying scene. In contrast, conventional dependent trajectory schemes perform the reweighting "on the fly," i.e., while the data in still in RAM. Consequently, one can only compute a finite number of predefined parameter values within the usual dependent sampling framework while in path recycling the whole continuum of parameter values is accessible, and this is key to the Bayesian retrieval methodologies we explored.

Specifically, to change plume size, position, and opacity, path recycling is just a matter of computational geometry: one only needs to know how to compute efficiently the intersection of a given segment and a given circle. For the aerosol perturbation it is more involved, and details are described in Ref. [1]. If there is a specific sponsor need, there is no reason why path recycling cannot be implemented inside far more capable 3D Monte Carlo codes such as the I3RC community Monte Carlo code http://i3rc.gsfc.nasa.gov/I3RC_community_model_new.htm or DIRSIG http://dirsig.org.

FAST MARKOV CHAIN MONTE CARLO FOR BAYESIAN SAMPLING OF THE 5-DIMENSIONAL PARAMETER SPACE

One could seek the set of single values of the 5 parameters in Table 1, with uncertainties, that jointly minimize the L_2 difference between the observed signal—synthesized using the "truth"—and the one predicted iteratively by the fast path-recycling Monte Carlo. Instead, we used Bayes' rule [4] to determine the 5-dimensional histogram of plausible parameter values based on a weak prior assumption on it (cf. Table 1) and the likelihood of the parameter-space sample, given the data. This calls for an effective sampling of the 5-dimensional parameter space using a Markov Chain Monte Carlo (MCMC) technique such as the classic Metropolis–Hastings algorithm, which however proved to be too slow for the present problem. This prompted the development of a novel MCMC algorithm specifically adapted to Monte Carlo-based forward models (the "MC^3" algorithm) where the cost of choosing to move to a new point in parameter space is managed by using predictions of lesser precision by known factors [5]. The resulting speedup is demonstrated in Fig. 2.

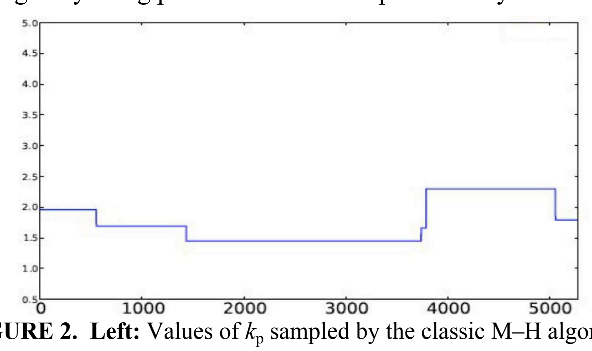

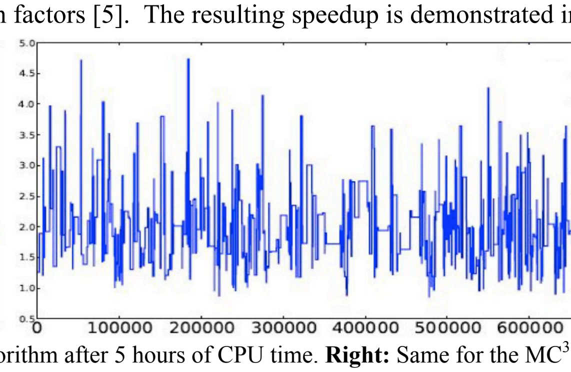

FIGURE 2. Left: Values of k_p sampled by the classic M–H algorithm after 5 hours of CPU time. **Right:** Same for the MC^3 method.

The marginal distributions we obtained are displayed in Fig. 3, along with the samples used for the joint marginal distribution of k_p and ρ_p where we note an anti-correlation originating from the direct dependence of the signal on k_p times the chord of the path segment inside the plume. Interestingly, the aerosol is better constrained than the plume.

The MC^3 algorithm is straightforward to extend to other parameters of the present 2D "demo" scene such as inherent aerosol properties (σ_0 and $\{f, g_f, g_b\}$) or surface albedo (α), as well as to 1D aerosol-only scenes or to 3D cloud scenes.

ACKNOWLEDGMENTS

This research was supported by grants from DOE (DE-FG52-08NA28779, to Columbia & JPL) and NSF (DMS-0804696, to Columbia). Part of it was carried out at JPL/Caltech, under a contract with the National Aeronautics and Space Administration. ©2012 California Institute of Technology. Government sponsorship acknowledged.

REFERENCES

1. I. Langmore, A. B. Davis, and G. Bal, "Multi-pixel retrieval of structural and optical parameters in a 2D scene with a path-recycling Monte Carlo forward model and a new Bayesian inference engine," *IEEE Trans. Geosci. Remote Sens.* (in press, 2012).
2. G. I. Marchuk et al., *The Monte Carlo Methods in Atmospheric Optics*. New York (NY): Springer-Verlag, 1980, pp. 45–48.
3. J. Heino, S. R. Arridge, J. Sikora, and E. Somersalo, *Phys. Rev. E* **68**, 031908-1-031908-8, (2003).
4. W. M. Bolstad, *Introduction to Bayesian Statistics*, 2nd ed. New York (NY): John Wiley, 2007.
5. G. Bal, I. Langmore, and Y. Marzouk, "Bayesian inverse problems with Monte Carlo forward models," *Inverse Problems and Imaging* (submitted, 2012).

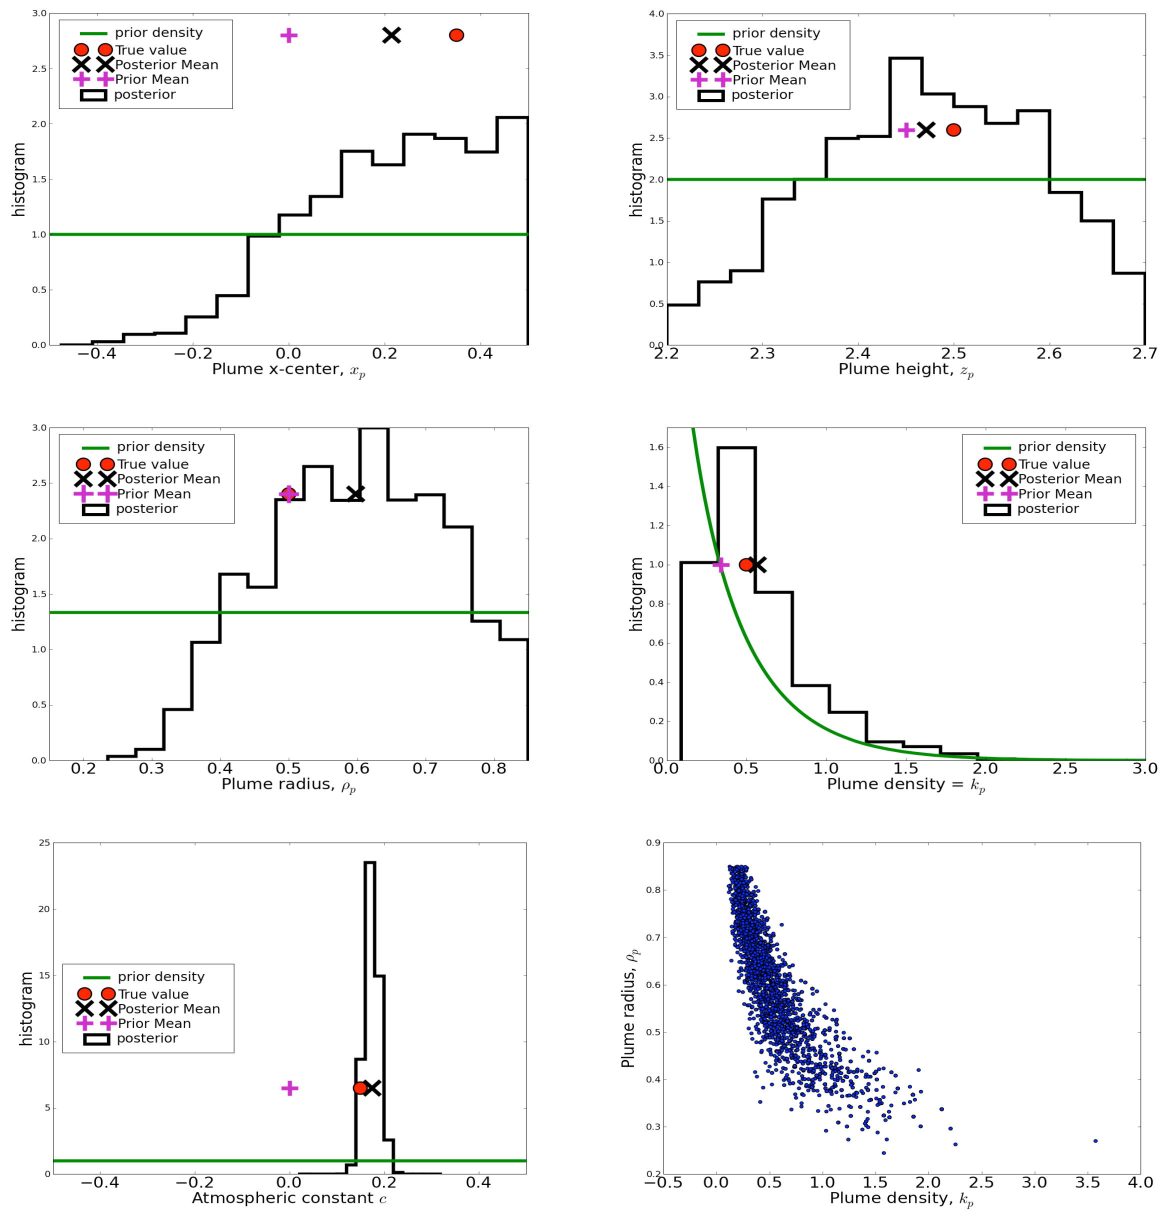

FIGURE 3. Top: Marginal posteriors (x_p,z_p). **Middle:** Same for (ρ_p,k_p). **Bottom:** Same for δc and samples for (ρ_p,k_p).