IRS2012-200
International Radiation Symposium 2012
Dahlem Cube, Berlin, Germany, 06 – 10 August 2012

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

I. Langmore (1), A.B. Davis (2), G. Bal (3), and Y.M. Marzouk (4)

(1) Opera Solutions, New York, NY, USA, (2) Jet Propulsion Laboratory / California Institute of Technology, Pasadena, Ca, USA (anthony.b.davis@jpl.nasa.gov, +1-818-393-4619), (3) Columbia University, New York, NY, USA, (4) Massachusetts Institute of Technology, Cambridge, Mass, USA

The vast majority of physics-based retrieval algorithms used in remote sensing of atmosphere and/or surface properties are multi- or hyper-spectral, some use multi-angle information; recently, polarization diversity has been added to the available input from sensors and accordingly modeled with vector radiative transfer codes. At any rate, a single pixel is processed at a time using a forward radiative transfer model predicated on 1D transport theory. Neighboring pixels are sometimes considered, but almost always to only formulate a statistical constraint in the inversion based on spatial context. We demonstrate here the potential power that could be harnessed by adding bona fide multi-pixel techniques to the mix.

To this effect, we use forward and inverse radiative transfer modeling in two spatial dimensions (sufficient for this demo) of a single-wavelength imaging sensor's response used to infer the position, size and opacity of an absorbing atmospheric plume somewhere in a deep valley in the presence of a partly-known/partly-unknown aerosol assumed to have an exponential profile with altitude.

We first describe the necessary innovation in forward multi-dimensional radiative transfer. In spite of its notorious reputation for inefficiency, we use a Monte Carlo technique. However, the adopted scheme is highly accelerated without loss of accuracy by using efficiently "recycled" Monte Carlo paths, a methodology adapted from ongoing research in biomedical imaging.

We then put the 2D forward model to work in a multi-pixel Bayesian inversion scheme that infers more than just the most probable values of the parameterized 2D scene's five unknown quantities: the plume's position, radius and density, and the specific aerosol amount. It also delivers an estimate of their 5-dimensional probability density function, hence means and associated uncertainty bounds from its marginal distributions.

This Bayesian inference is made possible by refining Monte Carlo Markov Chain (MCMC) methodology that, in essence, is based on generating random walks through the 5-dimensional parameter space. Here again the standard MCMC algorithms (e.g., classic Metropolis–Hastings) are too slow to converge in the present application. We therefore developed a much accelerated sampling algorithm that capitalizes on the fact that we can recycle only some of the Monte Carlo paths with a known penalty in precision. In many instances, a forward model prediction with degraded precision is sufficient to decide whether or not to accept the new sample. In the new "MC^3" algorithm, the savings in forward model prediction cost are thus immediately re-invested in an improved parameter space sampling.

In the end, the numerical evidence shows that the basic parameters of plumes in such scenes can be inferred, within known uncertainty bounds that add value in some important applications of airborne and space-based remote sensing.