Atmospheric Chemistry and Physics Discussions

Interactive comment on "An evolution strategy to estimate emission source distributions on a regional scale from atmospheric observations" *by* P. O'Brien et al.

M. Kanakidou (Editor)

mariak@chemistry.uoc.gr

Received and published: 16 October 2003

This paper presents a new computational technique for the inference of surface sources of a trace gas given concentration data. This is a widely used technique for inferring such sources, generally known as atmospheric inversion. The technique has found uses from regional to global scales and for species ranging from short-lived pollutants to long-lived and well-mixed gasses.

The new technique is an evolutionary strategy, one of the stochastic optimization procedures broadly akin to a genetic algorithm. The method is compared with a version of the singular value decomposition approach widely used in such problems. The test used is the recovery of a set of known sources, the so-called identical twin experiment.

Full Screen / Esc

Print Version

Interactive Discussion

Discussion Paper

© EGU 2003

the strength of the paper is, just as it should be, the presentation and demonstration of the new technique. The results are surprising indeed. Firstly the optimization method appears to work well (relative to SVD) even in the problem domain in which the linear SVD technique should be expected to perform best. The ES, and all similar global optimization strategies, are targeted to applications with nonlinear behaviour and hence complicated structures for objective functions. This is not the case here at all. Unless a nonnegativity constraint or other threshold is applied, the source determination problem here is completely linear. Still the ES performs well, both computationally and, more surprisingly, producing better fits to the data, i.e lower cost functions. I am curious about the first of these findings and sceptical about the second. If I understand the procedure correctly, the ES is likely to involve calculation, on average of the cost function for about 20% of the population each iteration. In 300 iterations with a population size of 40 this is about 2,400 calculations of the forward model. In this case the forward model consists, I think, of a matrix-vector product followed by a dot-product. This is, indeed, a pretty simple calculation. I can imagine this being computationally competitive with the SVD. But what happens when the forward model gets more complex? What, for example, if there were some nonlinear chemistry in the model? At one level we might expect this to favour the ES since we must now use the SVD (or some other nonlinear version of a gradient descent method) iteratively while the ES does not change. However the relative cost might well. This does not invalidate the paper at all but I believe the authors need to be claear on the problem domain for which the computational advantages of the ES will hold. The second behaviour, with the ES producing lower cost functions than the SVD, is surprising and, indeed, concerning. The SVD is, at least when used without a cut-off, an implementation of the solution of the linear least squares problem via the pseudo-inverse. In a linear problem it should find the global minimum, upto adition of any member of the null space. Unless I am misunderstanding something, failure to do this would have me inspecting my implmentation. Note, for example, that not all SVD algorithms are comparably stable. If I *am* misunderstanding something and the bahaviour is in fact correct, then the authors need to

Full Screen / Esc

Print Version

Interactive Discussion

Discussion Paper

© EGU 2003

explain it further in the paper.

There is one more serious concern about the treatment of both algorithms in the paper, the discussion, or lack of it, of uncertainty and ill-conditioning. In several places the authors describe the sVD as "failing" when it does not recover the target flux field. In at least some, and perhaps all of these cases this failure is a reflection of the physical problem the authors seek to solve. This is most clear for the under-determined problem. This under-determinacy is not just an algorithmic nuisance. There really *is* more than one solution possible here which produces the same, minimum, value of the cost function. This is clear from investigation of the various diagnostics of the inversion, especially the structure of the null space. I would rather be concerned, here, that the ES is hitting a single solution when a family should be possible. The same under-determinacy can occur even when the problem is not obviously rank-deficient. Imagine two gridcells which are only ever traversed together, that is if a trajectory passes over either one of them it must pass over the other. In a formulation in which the signature of the sources seen at the detector is the sum of contributions into the parcel along a trajectory, these two gridcells, *must*, I think, be indistinguishable in the inversion. We can simply trade-off emissions in one gridcell against those in the other without impacting the sum. As soon as there is a trajectory which only passes over one of the cells we have independent information but otherwise this problem seems fundamental. Much more complicated versions of the same logic can apply in real cases. This is a likely explanation for some of the spurious sources seen in both algorithms. The SVD allows various methods of diagnosing such problems. In particular, the uncertainties returned by uncertainty propagation from the data will likely show that various of these apparent failures occur in regions where the algorithm can hardly discern any information about the source distribution. I do not regard this as an algorithhm failure. The algorithm fails when either one cannot ascribe any confidence interval to the source distribution or when the method predicts "wrong" answers with high confidence.

There is one other point in the paper which either leaves me confused or concerned

Full Screen / Esc

Print Version

Interactive Discussion

Discussion Paper

about the validity of the comparison. The authors' discussion of the threshold on the initial population for the ES seems either unclear or incomplete. It is rather surprising to me that the choice of the initial population has much impact on the solution after 300 generations. It would suggest that the estimate returned by the ES was sensitive to the random initial condition. Is this true and, if so, does it provide a means to investigate the families of possible solutions for ill-conditioned problems?

Overall, I believe this paper should be published -after revisions- because, although I suspect the useful domain of application of the ES is limited, it is important that a potential alternative technique be presented to the atmospheric inverse community. I do believe the paper would be greatly improved by a more careful analysis of the SVD results themselves. I recommend, in particular, that the authors guarantee then explain why the ES produces lower cost functions than the SVD and that the authors pay some note to the potential for multiple solutions to the ill-conditioned inverse problem, either by considering generated uncertainties or by an analysis of quantities like the model resolution matrix.

Interactive comment on Atmos. Chem. Phys. Discuss., 3, 1333, 2003.