General Comments

This paper describes and demonstrates a system for estimating PM2.5 emissions given concentration measurements. This classical inverse approach has been widely used for greenhouse gases, less commonly used for chemically reactive gases and rarely for aerosols. The reasons for this are two-fold: Firstly the modelling of aerosol transport and modification/removal is difficult so introduces potentially confounding uncertainty into the calculations. Furthermore this uncertainty is hard to characterise, making the specification of the required PDFs difficult. Secondly the specification of background concentrations is difficult.

The authors have noted the background concentration problem though I think their method of dealing with it is debatable. They haven't dedicated as much attention to the aerosol modelling problem or its impact on their calculations.

The paper is clearly written with the methods fairly thoroughly described. It is certainly within scope for the journal.

Since this is a comment rather than final review I will concentrate on more general points.

Major Points

Modelling Aerosol Transport and Production/Loss The study treats PM2.5 as an inert tracer with no explicit surface loss processes (more on this below). Is this true? It's clearly a question of scale. Probably the right question is whether the transit time is short compared to the atmospheric lifetime. At the least it should introduce an extra uncertainty into the problem. The authors remove observations strongly influenced by scavenging which removes one important process. Surface deposition, though remains. If it is significant it might help explain the findings of the "real data" experiment. The inversion solves for net surface fluxes, a combination of sources and deposition. The prior, using only sources, misses part of the story. Could part of the reduced posterior flux be the influence of deposition? Finally there is no treatment of secondary aerosol. This is regarded as a significant contribution elsewhere so its importance should be at least discussed.

Background Concentration As the authors note, this is a bugbear for most regional inverse studies. The authors have chosen an elaboration of the "upwind background" approach often used. This requires that the aerosol field in the background air is homogeneous enough that the contribution of the background can be described by a single site (or a single site per wind direction here). It's pretty easy to construct cases where this won't hold, e.g. a source close enough to the city perimeter that its plume is missed by the background site but seen by the sites used in the inversion. The real world probably contains less artificial versions of this problem and we can't really tell how serious they are. There are approaches which treat the boundaries explicitly (e.g. Ziehn et al., 2014) but these depend on the model's ability to calculate sensitivity with respect to boundary cells. In any case the background should not be ignored in the OSSEs where errors in the background concentration will introduce correlated errors into the enhancements which should be treated in the observational covariance.

Minor Points

Negative Fluxes the solution method used by the authors is the classical linear approach of Tarantola (2005). This is quite capable of producing negative fluxes. Does it do so here? If so how are they treated? One should not truncate the emissions as positive post hoc since the resultant emissions map no longer minimises the Bayesian cost function. One can solve the problem subject to a positive constraint but this usually applies iterative solutions rather than the one-shot matrix method described here.

Calculation of Uncertainties The authors describe a Monte Carlo approach for doing this. If they are using a pure matrix method for the Maximum Likelihood Estimate (MLE) solution they don't need such an approach, the direct solution of the posterior covariance will work better. Furthermore it will not introduce potentially spurious correlations arising from the small ensemble size (usually ten in this study I think). the fact that the authors do use this Monte Carlo approach suggests there is some kind of positivity constraint for the MLE but more explanation is warranted.

Regularisation On P13 the authors describe the use of exponential correlation functions for their prior emissions. The justification is that these are needed over and above the specification of a prior covariance i.e. the Bayesian approach itself) in order to regularise the problem. I think this is unnecessary and not well supported by the reasons given. The Bayesian problem in the linear Gaussian case should always return a solution. Its uncertainty might be large and allow for an MLE that doesn't look very nice. That's a pity but still a fair statement of what the data allows us to say. the prior covariance shouldn't be used as a numerical device to avoid this. Rather, as required by the Bayesian formulation, the prior covariance should encapsulate the prior PDF for emissions. Positive covariance between neighbouring grid cells implies that an error of the *prior* estimates in one grid cell suggests an error of the same sign in neighbouring grid cells. Just as likely is an error in grid cells governed by the same emissions process, no matter how far they are away. I would like to see at least one test case where the prior covariances were removed.

References

Tarantola, A.: Inverse problem theory and methods for model parameter estimation, no. 89 in Other Titles in Applied Mathematics, Society for Industrial and Applied Mathematics, 2005.

Ziehn, T., Nickless, A., Rayner, P. J., Law, R. M., Roff, G., and Fraser, P.: Greenhouse gas network design using backward Lagrangian particle dispersion modelling: Part 1: Methodology and Australian test case, Atmospheric Chemistry and Physics, 14, 9363–9378, doi:10.5194/acp-14-9363-2014, URL `http://www.atmos-chem-phys.net/14/9363/2014/`, 2014.