Dear Reviewer,

We would like to thank you for your remarks and comments, which have helped us to improve the manuscript. Please find below our answers (marked in bold blue color).

The reviewed manuscript focuses on the development of an inversion framework using spaceborne PRISMA hyperspectral imager to estimate particulate matter emissions from industrial stacks. It addresses the challenges associated with aerosol plume retrieval over continental surfaces, considering, e.g. surface interference and particle characteristics. The study presents a methodology that incorporates surface reflectance estimation, radiative simulations, and an optimal estimation framework to retrieve aerosol properties and emission flow rates. The reliable retrieval of aerosol properties is of vital importance for correct estimate of the ecological footprint of a given factory or country as well as for the public health, so any effort improving our knowledge is laudable, and I think that the manuscript is topical. It is well organized and the presentation is clear. There are few things in the methodological part, though, which I'd like to see improved or clarified, so I've selected "major revision", but the changes I propose below are easy to implement.

General comments:

- The authors did a good job by presenting their model and by providing the optimal estimation formalism. But, the method lacks such an important validation step as a self-consistency study, which would both assure the reliability of the proposed method and estimate its uncertainty. Let me explain: if one has a forward model and the inversion algorithm for some instrument, one should be able to take the initial state of the atmosphere (some synthetic image or a retrieved one, with some prescribed aerosol size distribution), name it a "reference", run the forward simulation and obtain the "would be" multispectral image. Then, this information should be fed to an inversion algorithm, which should be run for several times, each time starting with a different guess state. The retrieved aerosol properties should be compared with the reference ones and the biases and root-mean-square values should be estimated. Presumably, a purely theoretical noise-free retrieval will give the result close to the reference. This will tell us about the reliability of the method. In the next step, the simulated radiances should be modified using some realistic noise and the retrieval should be repeated. The biases and r.m.s. values got at this step would tell us about the real uncertainty of the data presented in the manuscript and obtained from real observations using the suggested method. I saw the discussion of the uncertainties in Section 3.3, but as far as I understand, this does not clarify the self-consistency issues outlined above.

We perfectly understand your point. However we believe that the optimal estimation method is widely used remote sensing application and the link between the measurement noise and the retrieval uncertainties in such a formalism is well controlled (eg. Rodgers and successive applications). We agree that the OEM is sensitive to the prior information and that's why we paid a particular attention to the definition of the prior state vector using literature review. Moreover we control the convergence of the problem using the DOF and iteration number of

the Levenberg-Marquartd algorithm. The main advantage of the OEM is to have a diagnostic of the posterior error on the retrieved values.

- Another validation step I did not find in the manuscript is the comparison with existing observations. All three test cases happened in the period when the CALIOP/CALIPSO lidar was still operable, and the corresponding aerosol information retrieved from CALIOP observations should be available. Another space-borne lidar observing the Earth in the same period was ALADIN/Aeolus, but the comparison of PRISMA retrievals with Aeolus L2A optical data might be difficult due to its long averaging along the track (in this case, the apples-to-apples comparison would require averaging the PRISMA data along the same 87km length). I understand that a good overlapping for a given date is not guaranteed, but I believe that one can pick up another day, which would give the same image as in the figures provided in the manuscript, and at the same time the image would be overlapped with the CALIPSO track. Such a comparison would be a good validation both for the method and for the measurement itself. I do not suggest to redo the whole study, but at least one collocation and one comparison plot for one site are needed.

We thank you for the suggestion. You perfectly mentioned that a good overlapping between CALIPSO and PRISMA for the same location and the same date of our case studies is almost impossible to get. However the emission rate for any of the studied facilities is clearly not steady. It changes on daily and even hourly basis. So we can't validate our results with such a comparison. Moreover CALIPSO/CALIOP has an horizontal resolution of 300 m (distance between acquisition along-track) and vertical resolution of 30 m. The probability to have a valid retrieval for both instruments is very low. Not to mention that the accuracy of the CALIOP retrievals in the lower part of the atmosphere depends on the above atmospheric transmission and on an adequate lidar ratio. Such an analysis is beyond the scope of our paper. However we might consider your suggestion for further investigation using hyperspectral imager and space lidar like EarthCare.

Minor comments and technical corrections

The title mentions only the PRISMA hyperspectral imager whereas the authors use a synergy of PRISMA and S2/MSI multi-spectral observations, and it looks like the results shown in the article would not be achievable using only the PRISMA instrument. Wouldn't it be better to mention this synergy right in the title?

Actually S2/MSI is used to constrain the surface reflectance and not to investigate plume properties. So to avoid any confusion, we prefer to keep the title as it is.

Line 12 and elsewhere: normally, each number should be accompanied with its uncertainty. I hope that the uncertainties will be available after the first general comment above is addressed.

Correct. The uncertainties on the mass flow rate are deduced from the uncertainties on the effective wind speed and the surface mass concentration as presented in section 6.1. The uncertainties are now given in the abstract, Table 3 and the section 6.1. Please note that the number given in the abstract and the Table 3 were not coherent. It's corrected now.

Line 43: “though” -> “through”

Line 47: “impletation” -> “implementation”

Line 49: it would be better to use “smaller than” instead of “less than”

Done. Thank you for your review.

Line 105 and elsewhere: the spectral channels are mentioned, but no information on the kernels is given, whereas it would be interesting to see where does the information on aerosols come from.

We have added the spectral range used on Line 202.

Line 148 and the whole section: the regularized solution implies that there is a certain “safe” solution, around which we try to retrieve the result. How heavily is the solution regularized? What happens if the regularization coefficient is set to zero? Is the task so ill-conditioned that it won’t converge at all? It would be interesting to compare the number of equations with number of unknowns

The only regularization term (dumping factor) used here is in the Levenberg-Marquardt algorithm. The number of iterations is used to assess the convergence of the solution. The DOF are then checked to ensure that the solution is independent of the a priori.

Line 232: it is not surprising that the farther the guess from the real solution is, the longer the iterative process is, but what is the link between the number of iterations and the footprint size, given that the pixels are treated individually?

You are right. Each pixel is treated individually. Above 5 iterations, the size of the retrieved plume footprint becomes independent of the number of iterations. We use an empirical value of max 10 iterations to be sure to get all the pixels for which the algorithm converges.

Line 238: does the solution depend on the initial AOT value?

No. AOT is the component of the initial state vector having the weakest sensitivity (large associated uncertainty).

Line 293: “to estimated” -> “to estimate”

Done