

Review reply for “Estimating volcanic ash emissions using retrieved satellite ash columns and inverse ash transport modelling using VolcanicAshInversion v1.2.1, within the operational eMEP volcanic plume forecasting system (version rv4_17)”

Dear editor and reviewers,

Thank you for the comments and remarks on our paper. We have considered them all and made many changes to the manuscript to accommodate most of the requested changes. In our opinion, your comments have improved the paper and forced us to clarify certain phrasings and parts. We submit our updated manuscript with the changes highlighted. In addition, we have responded to the comments and remarks in this document.

On behalf of the authors, I wish you a happy holiday.

Best wishes,

André Brodtkorb

Comments from referees (CEC1): Dear authors,

I would like to bring to your attention an issue with the "Code and data availability" section in your manuscript. Currently, it reads that you have stored the code and data in GitHub. The problem here is double:

- First, GitHub is not a suitable repository for scientific publication. GitHub itself instructs authors to use other alternatives for long-term archival and publishing, such as Zenodo. In this way, your statement in this section is misleading and poorly educates readers that can assume that GitHub is a valid repository. Therefore, please, modify it in any potentially reviewed version of your manuscript, and change it to "Zenodo".
- Secondly, the assets of your manuscript are actually stored in Zenodo; however, you include this information in the list of references, with their links and DOIs. This is wrong. You must include in the Code and data availability section the DOIs and links, instead of including them as cites and references.

Therefore, please, in any potentially reviewed version of your manuscript, address these issues.

Regards,

Juan A. Añel

Geosci. Model Dev. Executive Editor

Author's response: Dear Juan A. Añel,

thank you for pointing out the inaccuracy in our code and data availability section - your comment is most welcome. We have changed the text in the manuscript to the following:

The source code and data used in this work is released under an open licenses, and archived versions are available on Zenodo:

Software package: André R. Brodtkorb. (2022). VolcanicAshInversion: v1.2.1 (v1.2.1). Zenodo. <https://doi.org/10.5281/zenodo.8073110>.

Satellite observations: André Rigland Brodtkorb. (2020). Eyjafjallajökull satellite observations (1.0) [Data set]. Zenodo. <https://doi.org/10.5281/zenodo.3855526>.

Forward eMEP simulations: André Rigland Brodtkorb, Alvaro Valdebenito, & eMEP contributors. (2020). Three-hourly gridded volcanic ash emissions for the Eyjafjallajökull 2010 eruption [Data set]. Zenodo. <https://doi.org/10.5281/zenodo.3818196>.

Comments from referees (CC1): Review of the article “Estimating volcanic ash emissions using retrieved satellite ash columns and inverse ash transport modelling using VolcanicAshInversion v1.2.1, within the operational eMEP volcanic plume forecasting system (version rv4_17)” by Brodtkorb et. al.

The paper presents a procedure for estimating ash emission rates by inverting retrieved satellite ash columns along with a priori information. Although the article introduces some innovative concepts, I believe that a more consistent and formal presentation would enhance its quality.

One notable concern is the absence of a description regarding the utilized observations. It remains unclear how many observations were employed and what the dimensions of the domain are. Additionally, it is not specified whether the domain is fixed in size and location, and it is unclear whether all emissions were inverted simultaneously after all observations became available.

It is crucial for the authors to explicitly state that the ash being considered does not possess a size distribution and is instead simulated using a single tracer. Moreover, it is essential to address whether the gravitational settling of ash has been taken into account in the study.

A significant issue that requires attention is the absence of colorbar units in most of the plots. Including these units would greatly enhance the interpretability of the figures.

In my opinion, the paper necessitates revisions to its text before it can be considered for publication.

Line 51-52: What do authors mean by “non-zero observations” and “zero observations”?

Line 62-64: Please rephrase this sentence and avoid using “our”, and what is meant by “our satellite images”?

Line 75: what is “system matrix”?

Sec 2.2 Is not clear what was used for the inversion. Column loadings or “...three-dimensional simulation results”?

Line 98: Not clear “unit of ash”. I saw it in other places too.

Fig 4b caption: Source receptor matrix M contains ‘footprints’ of each simulation. But in the text, it says “each row in the matrix corresponds to one observation of ash”. Not clear, which observation?

To be in agreement with the caption, I think that plot needs to be rotated by 90 degrees CW (Fig. 6a too). And the text on the arrows needs to be swapped as well in this case. I believe the procedure uses ash column loadings, not concentrations. SEVIRI observations first time mentioned, please allocate a separate section in the text, where you describe used observations.

Line 127: The role of Tikhonov regularisation is not correctly presented.

Line 148: "... discontinuous in the vertical dimension.". Maybe the authors mean "discontinuity in the solution vector"? I wonder how the authors apply "smoothing"? To the whole vector or to the 19-element sections of the vector?

Line 155: Formally, the transition to this expression was skipped in the text.

Line 165: $\gamma \rightarrow \gamma_0$?. I think this can be rephrased: "... are column loadings expressed in kg/m²."

Line 169: What is "LSQR" matrix?

Line 180: Not clear, how to compute G without storing M. Storing where? I imagine, that M is assembled from vectors when forward modeling is finished. It means that M is stored somewhere. Please clarify it, as you refer to this as one of the novelties of the paper.

Fig 6: missing colorbar units

Fig 7: missing colorbar

Line 233: Please replace 'satellite images' with 'satellite observations' everywhere

Line 246: Tikhonov relaxation -> Tikhonov regularisation

Line 346: Mt ->MT

Dr. Alexander Ukhov

Author's response: Dear Alexander Ukhov,

thank you for your review and feedback on the manuscript. We have tried to respond to all your comments as follows:

Comment: "One notable concern is the absence of a description regarding the utilized observations. It remains unclear how many observations were employed and what the dimensions of the domain are. Additionally, it is not specified whether the domain is fixed in size and location, and it is unclear whether all emissions were inverted simultaneously after all observations became available."

Response: Thank you for noticing this oversight. This information is "hidden" in the supplementary material/data. We have added a paragraph in both the validation and verification subsections to describe the data used. The mathematical formulation inverts all observations simultaneously, though in the "Summary and discussions", we write that "Our framework is also well tailored to an operational setting with an eruption going on for several weeks. It is technically possible to use a least squares matrix G1 representing e.g., the first three days of the eruption, and simply add another matrix G2 that may represent a subsequent four days (and equivalently for the B vector), instead of assembling the full seven day period from scratch. With satellite imagery coming in continuously, this may save large amounts of computational time to create up-to-date estimates."

Comment: "It is crucial for the authors to explicitly state that the ash being considered does not possess a size distribution and is instead simulated using a single tracer. Moreover, it is essential to address whether the gravitational settling of ash has been taken into account in the study."

Response: The advection model is not the focus of this paper, as it is merely a tool used to create input data for the inversion. However, we have added text to clarify.

Comment: "A significant issue that requires attention is the absence of colorbar units in most of the plots. Including these units would greatly enhance the interpretability of the figures."

Response: "Figures 6 and 7 were missing colorbars, as we did not feel the added value of a colorbar in these illustrative figures (the point of these figures is to show the matrix/vector structure and illustrate what the observations look like compared to the visual truth). We have, however, added a colorbar to figures 6 and 7."

Comment: "Line 51-52: What do authors mean by “non-zero observations” and “zero observations”?"

Response: "Text has been updated to “In this paper we use a similar approach by using observations of non-zero ash mass from the ground up to the detected plume top, and observations of zero ash mass above the observed ash cloud-top.”"

Comment: "Line 62-64: Please rephrase this sentence and avoid using “our”, and what is meant by “our satellite images”?"

Response: "We have rephrased, removing the use of “our”, and using “satellite observations” instead of “satellite images”."

Comment: "Line 75: what is “system matrix”?"

Response: "Added a reference to section 2.2 here, which explains in detail. "

Comment: "Sec 2.2 Is not clear what was used for the inversion. Column loadings or “...three-dimensional simulation results”?"

Response: "The mathematical notation uses two-dimensional coordinates, see eqn (3)."

Comment: "Line 98: Not clear “unit of ash”. I saw it in other places too."

Response: "Added a reference to figure 3, that perhaps explains it best."

Comment: "Fig 4b caption: Source receptor matrix M contains ‘footprints’ of each simulation. But in the text, it says “each row in the matrix corresponds to one observation of ash”. Not clear, which observation? "

Response: "Rephrased the caption to make it a bit more clear, “Linear system of equations. (a) shows the vector of observations originating from the SEVIRI instrument, in which a single location is highlighted with its corresponding location in the observation vector. Note that the image shows only part of the domain, and many of the detected ash pixels are outside the highlighted area. (b) shows the source-receptor matrix M. Each row in the matrix contains all simulated ash mass loadings at the observation coordinate of the observation in the same row in the observation vector (see also Equation (3)). Equivalently, each column in the matrix corresponds one individual emission simulation, S_j . The matrix shown here has 60 observations (columns), three emission time points, and 19 emission altitudes (3×19 rows). White means no ash in the simulation, and the colored elements correspond to the concentration of ash in the individual simulations. See also Figure 5 which shows the full matrix”"

Comment: "To be in agreement with the caption, I think that plot needs to be rotated by 90 degrees CW (Fig. 6a too). And the text on the arrows needs to be swapped as well in this case. I believe the

procedure uses ash column loadings, not concentrations. SEVIRI observations first time mentioned, please allocate a separate section in the text, where you describe used observations."

Response: "We believe the figures to be correct. We solve the system $Mx = y_0$, in which y_0 is the column vector shown in subfigure (a), and M is shown in subfigure (b). The SEVIRI dataset is described in detail in the data availability section with a reference to the actual data and its provenance."

Comment: "Line 127: The role of Tikhonov regularisation is not correctly presented."

Response: "We have rephrased to "The aim is to find a smooth solution that accounts for both observational and modeling errors.""

Comment: "Line 148: "... discontinuous in the vertical dimension.". Maybe the authors mean "discontinuity in the solution vector"? I wonder how the authors apply "smoothing"? To the whole vector or to the 19-element sections of the vector?"

Response: "Smoothing the whole vector would lead to significant artefacts close to ground/top of simulation domain, and is therefore unwanted. The formulation using the second derivative penalizes discontinuities in the vertical for 19-element sections of the vector."

Comment: "Line 155: Formally, the transition to this expression was skipped in the text."

Response: "The transformation from just before line 155 to just after is relatively straight forward linear algebra, though we will be happy to include it if requested."

Comment: "Line 165: $y \rightarrow y_0$?. I think this can be rephrased: "... are column loadings expressed in kg/m^2 .""

Response: "Thank you for spotting the missing subscript! Rephrased the sentence also."

Comment: "Line 169: What is "LSQR" matrix?"

Response: "Expanded to least squares matrix."

Comment: "Line 180: Not clear, how to compute G without storing M . Storing where? I imagine, that M is assembled from vectors when forward modeling is finished. It means that M is stored somewhere. Please clarify it, as you refer to this as one of the novelties of the paper."

Response: "All of the information required to build the "huge" matrix M lies in collocated netcdf-files. The traditional approach is to assemble the "huge" M matrix in order to compute the "small" matrix G using equation (11). This, however, quickly uses an enormous amount of memory. The "magic sauce" is to replace the traditional expression (see Equation (12)) with a sum of outer products (Equation (13)), thus reducing the memory requirement to a fraction."

Comment: "Figures 6 and 7 missing colorbar"

Response: Added

Comment: "Line 233: Please replace 'satellite images' with 'satellite observations' everywhere""

Response: "Replaced "images" with "observations where it makes sense."

Comment: "Line 246: Tikhonov relaxation \rightarrow Tikhonov regularization"

Response: Rephrased

Comment: "Line 346: Mt - >MT"

Response: Thank you, fixed typo.

Comments from referees (RC1): This manuscript presents an inversion algorithm employed by the Norwegian Meteorological Institute to estimate ash emission rates using satellite data on ash column loadings and a-priori information. Ash emission rate is a crucial parameter for accurate modeling of volcanic ash dispersion. The procedure can be summarized in five key stages: (1) generating preliminary emission estimates, (2) conducting forward simulations with various emission profiles, (3) aligning observations with simulation results, (4) constructing a system of linear equations, and (5) solving this over-determined system. The paper delves into the underlying mathematical principles of this inversion process and assesses its performance in both simulated and real-world scenarios.

The manuscript is well structured and written. The research is generally suited for publication in GMD. But there few concerns listed below addressing which requires major revision.

1- Ash emission rate is a crucial parameter but not the only one as the authors mention in the text. But there is lack of discussion how the uncertainties associated with ash properties (particle size distribution, density, shape) and processes (aggregation, sedimentation) might affect the results. In general, there is no concrete and quantitative analysis on the uncertainties of the method which a major drawback of the manuscript.

2- In many aspects, the authors ignore the recent advancements made in the field of volcanic ash dispersion modeling. In particular, there are several novel approaches on plume height estimation (Horváth et al. 2021), ESP calculations (Bruckert et al. 2022) and dispersion modeling (de Leeuw et al. 2021; Muser et al. 2020) all documented in the special issue of ACP/AMT/GMD on Reikoke eruption 2019 (https://acp.copernicus.org/articles/special_issue1104.html). The authors should include and discuss these references, where suited.

3- The authors claim "A novelty in this paper is the use of observed altitude from e.g., the SLSTR instrument." But there is no info and discussions about the pros and cons of this method compared to other methods (e.g. stereo retrievals). A critical aspect is differentiation between ash cloud and meteorological clouds. The view of satellites is often 'polluted' by meteorological clouds as one can see in Figure 7 for example. Besides, both mass and height estimations are usually problematic for dense ash plumes i.e. in the first few hours of the eruption. These aspects must be discussed in the paper and if possible, the corresponding uncertainties should be analyzed quantitatively.

4- Figure 11 is crucial for validation but not so much informative in the current form. It is impossible to make a quantitative conclusion based on this figure. Authors should use quantitative measures like structure, amplitude and location score (SAL) which is recently adopted for validation of ash dispersion forecasts (de Leeuw et al. 2021; Muser et al. 2020).

5- There is no information on how eMEP model considers the ash removal processes as a function of particle density, PSD, shape and aggregation. Ten years ago, it was perhaps fine if the models neglected some of these aspects but nowadays we know more than enough to accept such simplifications. Especially because the ash mass loading is significantly affected by the removal processes. This should be discussed in the paper.

6- Label the y-axis in Figures 9, 10, 12.

References:

Horváth, Á., Carr, J. L., Girina, O. A., Wu, D. L., Bril, A. A., Mazurov, A. A., Melnikov, D. V., Hoshyaripour, G. A., and Buehler, S. A.: Geometric estimation of volcanic eruption column height from GOES-R near-limb imagery – Part 1: Methodology, *Atmos. Chem. Phys.*, 21, 12189–12206, <https://doi.org/10.5194/acp-21-12189-2021>, 2021.

Bruckert, J., Hoshyaripour, G. A., Horváth, Á., Muser, L. O., Prata, F. J., Hoose, C., and Vogel, B.: Online treatment of eruption dynamics improves the volcanic ash and SO₂ dispersion forecast: case of the 2019 Raikoke eruption, *Atmos. Chem. Phys.*, 22, 3535–3552, <https://doi.org/10.5194/acp-22-3535-2022>, 2022.

de Leeuw, J., Schmidt, A., Witham, C. S., Theys, N., Taylor, I. A., Grainger, R. G., Pope, R. J., Haywood, J., Osborne, M., and Kristiansen, N. I.: The 2019 Raikoke volcanic eruption – Part 1: Dispersion model simulations and satellite retrievals of volcanic sulfur dioxide, *Atmos. Chem. Phys.*, 21, 10851–10879, <https://doi.org/10.5194/acp-21-10851-2021>, 2021.

Muser, L. O., Hoshyaripour, G. A., Bruckert, J., Horváth, Á., Malinina, E., Wallis, S., Prata, F. J., Rozanov, A., von Savigny, C., Vogel, H., and Vogel, B.: Particle aging and aerosol–radiation interaction affect volcanic plume dispersion: evidence from the Raikoke 2019 eruption, *Atmos. Chem. Phys.*, 20, 15015–15036, <https://doi.org/10.5194/acp-20-15015-2020>, 2020.

Author's response: Dear reviewer,

thank you for your comments and remarks on our paper. The following tries to respond to your review.

1- Ash emission rate is a crucial parameter but not the only one as the authors mention in the text. But there is lack of discussion how the uncertainties associated with ash properties (particle size distribution, density, shape) and processes (aggregation, sedimentation) might affect the results. In general, there is no concrete and quantitative analysis on the uncertainties of the method which a major drawback of the manuscript.

Response: Parameter sensitivity analysis is outside the scope of this work (and some of it handled in Steensen et al. 2017), though, such a work would be very interesting to pursue and a natural continuation that warrants significant effort in itself. The most important source of uncertainty in our work is most probably the difference in modeled and true meteorology. Errors in wind directions and speeds, vertical shear in both wind speed and direction at high altitudes, as well as numerical errors in the advection scheme itself will place the plume in a different spatial location than the observation. See for example (Raskob, 2020) and (Harvey et al., 2022) for some more discussion on uncertainties affecting the inversion. We added the following text to section 2.5, Iterative inversion procedure:

“Because there are large uncertainties in both the meteorology and the satellite observations (see e.g., (Harvey et al., 2022))”

Raskob, W., Beresford, N. A., Duranova, T., Korsakissok, I., Mathieu, A., Montero, M., ... & Woda, C. (2020). CONFIDENCE: project description

Steensen, B., Kylling, A., Kristiansen, N., and Schulz, M.: Uncertainty assessment and applicability of an inversion method for volcanic ash forecasting, *Atmospheric Chemistry and Physics*, 17, 9205–9222, 2017a

Harvey, N. J., Dacre, H. F., Saint, C., Prata, A. T., Webster, H. N., & Grainger, R. G. (2022). Quantifying the impact of meteorological uncertainty on emission estimates and the risk to aviation using source inversion for the Raikoke 2019 eruption. *Atmospheric Chemistry and Physics*, 22(13), 8529-8545.

2- In many aspects, the authors ignore the recent advancements made in the field of volcanic ash dispersion modeling. In particular, there are several novel approaches on plume height estimation (Horváth et al. 2021), ESP calculations (Bruckert et al. 2022) and dispersion modeling (de Leeuw et al. 2021; Muser et al. 2020) all documented in the special issue of ACP/AMT/GMD on Raikoke eruption 2019 (https://acp.copernicus.org/articles/special_issue1104.html). The authors should include and discuss these references, where suited.

Response: Some of the papers mentioned are more towards modeling the physics in the dispersion process, and we have not included these, as we feel they are not central to our contribution on the inversion. However, we have added Horváth et al. 2021 to the text – thank you for the suggestions. The added text is “The ash cloud-top height can be estimated using the SLSTR instrument, using photogrammetric parallax between different satellites (Zakšek et al., 2013), or other geometric approaches (Horváth et al., 2021))

3- The authors claim “A novelty in this paper is the use of observed altitude from e.g., the SLSTR instrument.” But there is no info and discussions about the pros and cons of this method compared to other methods (e.g. stereo retrievals). A critical aspect is differentiation between ash cloud and meteorological clouds. The view of satellites is often ‘polluted’ by meteorological clouds as one can see in Figure 7 for example. Besides, both mass and height estimations are usually problematic for dense ash plumes i.e. in the first few hours of the eruption. These aspects must be discussed in the paper and if possible, the corresponding uncertainties should be analyzed quantitatively.

Response: The discussion of SLSTR vs stereo retrievals and quality of input data is not the focus of our work, as our focus is on the inversion procedure. Our method should work with altitude observations reconstructed from stereo retrievals as well.

4- Figure 11 is crucial for validation but not so much informative in the current form. It is impossible to make a quantitative conclusion based on this figure. Authors should use quantitative measures like structure, amplitude and location score (SAL) which is recently adopted for validation of ash dispersion forecasts (de Leeuw et al. 2021; Muser et al. 2020).

Response: Thank you for the information about the SAL metric. A thorough discussion on the SAL/FSS together with a parameter sensitivity analysis would be a natural continuation of our work.

5- There is no information on how eMEP model considers the ash removal processes as a function of particle density, PSD, shape and aggregation. Ten years ago, it was perhaps fine if the models neglected some of these aspects but nowadays we know more than enough to accept such simplifications. Especially because the ash mass loading is significantly affected by the removal processes. This should be discussed in the paper.

Response: The eMEP model is outside the scope of this work, and our inversion framework is agnostic to the forward simulation model. The eMEP model is discussed in more detail in the listed references. A continuation of our work could be to look at how it behaves with updated eMEP models, different forward models (including e.g., Flexpart) and different meteorologies (which is more influential than the specific advection and ash removal model ref. our comment above on uncertainties).

6- Label the y-axis in Figures 9, 10, 12.

Response: The y-axis label has been omitted due to space constraints (it is mentioned in the caption in all these figures), especially for figure 9. However, we would be happy to add it if that is a strong wish.

Comments from referees: Review of "Estimating volcanic ash emissions using retrieved satellite ash columns and inverse ash transport modelling using VolcanicAshInversion v1.2.1, within the operational eMEP volcanic plume forecasting system" by Brodtkorb et al.

This paper presents a procedure for estimating volcanic ash emission as a function of altitude and time, from satellite observations. The overall approach is similar to that used before for such estimations, namely a regularised least-squares method. The focus of the manuscript is a description of this method and its application to synthetic data and my review therefore focuses on the method itself, rather than the challenges of applying it in real-world scenarios.

The paper is well written and is relevant for the journal, but I would like to see quite a bit more detail given on the performance of the algorithm, and consequence of various choices made in algorithm and its parameters.

- Around line 65: Negative emissions are avoided by iteratively reducing the uncertainty of the a priori for emissions that are initially negative. How does this method compare to the direct non-negative least squares solutions used by Pelley et al (2021)? In particular, does the chosen algorithm always converge, and how quickly? Is it robust to small changes in the observations? Does the algorithm used have a greater requirement for an accurate a priori than the direct non-negative least squares methods?

- Around line 148: What does 'discontinuous in the vertical direction' mean in this context? (emission sources at adjacent altitudes having very different fluxes? Why is it nonphysical for there to be rapid altitude variation in the source flux?)

- Around line 150:

-- is the matrix D really diagonal? In AshInversion.py, D appears to be defined as tridiagonal, not diagonal.

-- With respect to what variable is D taking the second derivative? The smoothing applied by the matrix D as implemented in the code seems to act sometimes across altitude and sometimes across time (if I am right in thinking that the n_{emis} rows/columns of D correspond to emissions at both different altitudes and different times). Can this be right? Some more clarification is needed.

-- What are the dimensions of D and of epsilon?

- Figure 5 caption: The matrix M here (~736k rows) is somewhat smaller than the one discussed in section 2.4 (56.5M rows). Can you clarify what the difference is?

- Around line 210: could you expand a little on the details of the memory:

-- in the light of the comment about efficient dense linear algebra routines, do you use dense or sparse matrix representation for M, D, σ_x and σ_o , when they occur in equations (12), (13), (15)?

-- What is meant by 'memory expansion' on line 210?

- Around line 235:

Are the synthetic satellite observations, at random points in space and time, representative of real observations, which presumably have more structure (images at periodic time intervals?). Does such a difference influence the performance of the inversion system?

- Around line 245:

-- is 'Tikhonov relaxation' a synonym for 'Tikhonov regularisation'?

-- Am I right in thinking that the difference between a posteriori and a priori in case A is only due to the smoothing term ($\epsilon D^T D$) in (11), and without this the a posteriori would be identical to the a priori? If so, then it is intriguing that the a posteriori here (with $\epsilon=10^{-3}$) looks *less* smooth (peak flux is higher) than the a priori. In this case A, it would be useful to explain in a bit more detail why and how the difference between a priori and a posteriori depends on the Tikhonov regularisation, and what the influence of epsilon is.

- Around line 275:

-- In both cases D and E, the inferred a posteriori has a distinct increase in inferred flux in emissions at the second highest altitude level, persistent across all times. What is the reason for this?

-- In cases D and E, if a priori error σ_x is increased (reflecting the lack of confidence in the a priori in this case), might the inversion code work better? That is, is it possible that there is sufficient information to invert accurately in cases D and E, but the a posteriori is simply being degraded by an inaccurate prior?

- Figure 12 caption: "lelve" -> "level"

- Line 390: This link is now to zenodo, rather than Github

Author's response: Thank you for your review and comments. The following tries to address your comments.

- Around line 65: Negative emissions are avoided by iteratively reducing the uncertainty of the a priori for emissions that are initially negative. How does this method compare to the direct non-negative least squares solutions used by Pelley et al (2021)? In particular, does the chosen algorithm always converge, and how quickly? Is it robust to small changes in the observations? Does the algorithm used have a greater requirement for an accurate a priori than the direct non-negative least squares methods?

Response: This is a very interesting point, as an NNLS could possibly avoid the "brute force" nature of forcing a negative a posteriori closer to the a priori by reducing the uncertainty. We have experimented with NNLS in this regard, but unfortunately only at a preliminary stage. In our experiments, the approach of iterative reduction of uncertainties is quite robust with respect to small changes in the a priori as well as the observations. We have not added the use of NNLS (and similar approaches) this to the paper, but it is a natural extension of the work.

- Around line 148: What does 'discontinuous in the vertical direction' mean in this context? (emission sources at adjacent altitudes having very different fluxes? Why is it nonphysical for there to be rapid altitude variation in the source flux?)

Response: Thank you for this question, we have paraphrased the sentence in question. Whilst it may be physically OK to have a sharp gradient in the vertical dimension, the solutions produced without such a penalization term is typically deemed as nonphysical solutions. We rephrased the text to "Unfortunately, solving this minimization problem often results in implausible solutions with non-

physical discontinuities in the vertical dimension. To avoid such nonphysical solutions we can add a smoothness minimization term,”

- Around line 150: is the matrix D really diagonal? In AshInversion.py, D appears to be defined as tridiagonal, not diagonal.

Response: Thank you for spotting this. It is as you say tridiagonal to represent the second order derivative. We have fixed the mistake in the text.

-- With respect to what variable is D taking the second derivative? The smoothing applied by the matrix D as implemented in the code seems to act sometimes across altitude and sometimes across time (if I am right in thinking that the n_emis rows/columns of D correspond to emissions at both different altitudes and different times). Can this be right? Some more clarification is needed.

Response: The minimization term is $J_3 = \epsilon \|D\tilde{x}\|$, though in the Tikhonov regularization it enters as $\epsilon D^t D$. Thus, in the code it will look as though it operates on both rows and columns.

-- What are the dimensions of D and of epsilon?

Response: epsilon is a scalar ($1.0e-3$ in our experiments - see line 151), and D is the $n \times n$ where n is the number of a priori estimates (typically a few hundred to a few thousand).

- Figure 5 caption: The matrix M here (~736k rows) is somewhat smaller than the one discussed in section 2.4 (56.5M rows). Can you clarify what the difference is?

Response: The difference here is that figure 5 only considers observations of non-zero ash, i.e., detected ash cloud. This disregards all clear sky observations in the dataset, which are important observations in the inversion. We have added a footnote to clarify this: “The full Eyjafjallajökull eruption case has 876906 observations of ash and 55.7 million observations of no ash. In addition, there are uncertain observations e.g., due to cloud coverage, that are not used.”

- Around line 210: could you expand a little on the details of the memory:

Response: If we examine equation (12), we see that the inner product formulation creates the matrix $(M^t \sigma_o^2 M)$, which has dimension $n \times n$, in which n is the number of observations. This gives the 56M x 56M-shaped matrix which typically exhausts memory. However, we can partition the matrix M (and σ_o) into a subset of observations, e.g., 1000 observations at a time. This yields a 1000x1000 matrix, but we have to add/assemble 56 thousand of these matrices into the G matrix in order to fully assemble the full 56 million observations. We can also elaborate a bit more on this in the paper if that is a request.

-- in the light of the comment about efficient dense linear algebra routines, do you use dense or sparse matrix representation for M, D, σ_x and σ_o , when they occur in equations (12), (13), (15)?

Response: We do not represent M at all in our computations, as it is such a large matrix. We represent D as a dense matrix (though, it is only $n \times n$ in which n is the number of a priori values to estimate). σ_x is represented as a vector, though converted to a dense matrix in the calculation (again, size $n \times n$), and σ_o is represented as a vector.

-- What is meant by 'memory expansion' on line 210?

Response: Memory expansion is poorly phrased – thanks for noticing. It means that the outer product formulation requires computing the outer product of an n -long vector (expanding the vector to an $n \times n$ matrix) for each observation, in which n is the number of a priori values. We have rephrased to “as well as the expensive outer product formulation required for Equation (13)”.

- Around line 235: Are the synthetic satellite observations, at random points in space and time, representative of real observations, which presumably have more structure (images at periodic time intervals?). Does such a difference influence the performance of the inversion system?

Response: Having uniformly distributed observations is probably better than real-world observations. In real-world observations, you can have (water) clouds partly covering the ash cloud. Thus in our synthetic observations we probably have a better set of observations to cover the extent of the ash cloud than what we see in real-world data.

- Around line 245: is 'Tikhonov relaxation' a synonym for 'Tikhonov regularisation'?

Response: Thank you for noticing. We have changed the text to regularization.

-- Am I right in thinking that the difference between a posteriori and a priori in case A is only due to the smoothing term ($\epsilon D^T D$) in (11), and without this the a posteriori would be identical to the a priori? If so, then it is intriguing that the a posteriori here (with $\epsilon=10^{-3}$) looks *less* smooth (peak flux is higher) than the a priori. In this case A, it would be useful to explain in a bit more detail why and how the difference between a priori and a posteriori depends on the Tikhonov regularisation, and what the influence of ϵ is.

Response: We would not (in general) expect to be able to perfectly recreate the a priori in our a posteriori, even without smoothing and with using all possible observations in the inversion. Consider for example a uniform wind field (also in the vertical dimension). Even with perfect observations (and observing all locations), it is not possible for us to determine at what altitude the ash is emitted as we only observe vertically integrated values. Thus, our inverted estimates have an inherent uncertainty that depends not only on observations, but also on the weather conditions.

- Around line 275: In both cases D and E, the inferred a posteriori has a distinct increase in inferred flux in emissions at the second highest altitude level, persistent across all times. What is the reason for this?

Response: Good question, which is difficult to answer definitively. Our hypothesis is that this is because the upmost level has a reduced a priori (due to the vertical discretization in which an 8.4 km plume ends in the middle of a layer).

-- In cases D and E, if a priori error σ_x is increased (reflecting the lack of confidence in the a priori in this case), might the inversion code work better? That is, is it possible that there is sufficient information to invert accurately in cases D and E, but the a posteriori is simply being degraded by an inaccurate prior?

Response: In our experiments, increasing the uncertainty has not lead to increased skill in the inversion. The inherent difference in true meteorological conditions and modeled conditions is in our experience one of the major contributors to this.

- Figure 12 caption: "lelve" -> "level"

Response: Embarrassing. Thank you.

- Line 390: This link is now to zenodo, rather than Github

Response: This has been updated also – thank you!