Response

Reviewer #2

We thank Reviewer #2 for the very constructive and detailed feedback. We included almost all the suggestions which substantially improved our manuscript! In the following we reply to questions and comments in blue while new or changed text is coloured in green.

Specific comments

Posterior emissions are analysed across the whole of Europe, despite the authors noting that emission sensitivity from the observations is poor for most of eastern and southern Europe. To address this, a more detailed comment should be included about the uncertainty of results beyond the region of good sensitivity. Or, results should only be analysed and presented over a smaller region that does have good sensitivity to emissions. Showing spatial emission results over a smaller region would also make all of the spatial emission figures easier to read.

Yes, we agree! We initially chose the full European domain because it includes both regions that are well constrained by the observations and regions with poor observational coverage. In eastern and southern Europe, where the emission sensitivity is low, the improvements through the inversion is limited and the posterior emission are more uncertain. This is, for instance, clearly to see in our synthetic experiments, where perturbations in the prior can almost not be corrected in these regions. For the purpose of this study, we wanted to investigate the behaviour of the algorithm under both strong and weak observational constraint as limited observational coverage is a common challenge in atmospheric inversions. In particular, we wanted to study the transition from the lognormal prior towards a more normally distributed posterior, and its dependence on emission sensitivity. As illustrated in Fig. 8, this transition occurs predominantly in well-constrained regions, which is consistent with our expectation that normally distributed observational uncertainties exert a relatively stronger influence on the shape of the posterior distribution where the observational constraint is larger.

To make this clear we added:

Emission sensitivities are substantially lower in eastern and southern Europe, and thus, the improvements achieved by the inversion in these regions are limited and posterior emission estimates are more uncertain.

In contrast, in regions with weaker observational constraints, the transition is weak or absent, so that the posterior largely retains the lognormal shape of the prior.

Some topics that would be unfamiliar to non-experts (such as using atmospheric models to link surface emissions to atmospheric concentrations) are treated as assumed knowledge, which is a reasonable choice considering the scope and target audience of this paper. If the authors want to make this paper relevant to a wider audience, they would need to introduce some terms (such as atmospheric transport) in more detail, but I will leave this to the authors discretion.

Thank you for this comment. We agree that some topics are treated as assumed knowledge. As the topic is rather specific, the primary target audience of this study is the atmospheric inversion community, particularly researchers interested in implementing and using lognormal priors in atmospheric inversions. We therefore feel that a more detailed introduction to these established concepts would go beyond the scope of the paper, and we would prefer to retain the current level of detail. We nevertheless appreciate the suggestion and agree that a more extensive introduction could be useful for a broader audience.

Technical corrections and questions

Line 13: missing word after "however"? Should this read "...however, they avoid..."

Yes, thank you - >done

Line 19: this line may be clearer if "based on its measurement in the atmosphere" is replaced with "based on measurements of its concentration in the atmosphere".

Yes, agreed - >done

Line 23: The repeat of "(or mole fractions)" could be removed, as this is already included on line 22.

Yes, agreed, thank you - >done

Line 24: replace "determining" with "determine".

Yes - >done

Line 27: what does "variation, especially due to climate" mean here? Spatial or temporal variation or trends in emissions?

We included "spatiotemporal"

Line 37: "and prior estimate of **x**" should be added after "observations, **y**".

done

Line 49: is there a reference for the line "The prior emission estimates for many atmospheric species have a lognormal distribution, strongly suggesting that the errors follow a lognormal distribution as well"?

Yes, we cited Super et al. 2020 which seems appropriate, who studied uncertainties of emission inventories, suggesting lognormal distributed errors.

Line 51: "cost function" has not been introduced yet. A line could be added earlier in this section, describing how a cost function is used to find a most likely estimate of **x**.

Yes, that is a great idea: We added:

"The maximum of the posterior distribution P (x | y), corresponding to the most probable solution, can be obtained by minimizing a cost function J given by the negative logarithm of the posterior probability"

after introducing Bayes' Theorem.

Line 79: add a line here, explaining why the PDF needs to be expressed in terms of the error, otherwise the reader has to wait for the presentation of the cost function for this explanation.

Yes, that is a good idea, thank you. We added:

"The probability distribution is expressed in terms of the errors ε, as the uncertainty is associated with deviations from the prior estimate. For the lognormal distribution these errors are geometric, meaning……"

Line 81: add brackets to these equations (e.g. $\ln(\varepsilon) = \ln(x) - \ln(xb)$) to maintain consistency with where this equation is used again later (e.g. line 101).

Yes, done!

Line 108: I think the footnote here could be moved to after equation 13, as this is useful information and it would not break the flow of the text here. However, this is just a personal preference, so I'll leave this to the authors' discretion.

Thank you, we agree and changed the text accordingly.

Line 115: some discussion of the choice of mean/median/mode in relation to atmospheric inversions could be included here. How does this choice differ compare to the data assimilation context in the quoted references?

Yes, that's an excellent idea! We added:

While these results were derived in the context of data assimilation, the choice of estimator is also relevant to atmospheric inversions, where x_b represents a prior emission estimate rather than a model forecast and the observational constraint can vary substantially across the domain. The performance of the mean, median, and mode may therefore depend on both the accuracy of the prior emission estimate and the strength of the observational constraint.

which also bridges to our synthetic tests, using different uncertainties!

Line 125: ‘transport function H(x)” is referenced here, when H has only been referred to as an “operator” before this point. To resolve this, more detail on H could be added at line 90, explaining how this could contain output from a Lagrangian or Eulerian transport model.

Agreed, we added/changed:

..... where H is an operator that maps the state to the observation space. For atmospheric inversions, H represents the forward model that relates the emissions to the observed concentrations, using an atmospheric transport model, such as a Lagrangian particle dispersion model (LPDM) or a Eulerian transport model.

Line 127: reference or derivation for this equation? Other equations have been derived in detail, (e.g. with clear references to what substitutions and transformations are being used) but there is less detail on the derivation of Equation 15 from Equation 14.

Yes, we add more details to the derivation:

To obtain the gradient with respect to $\mathbf{z}$, we need to apply the chain rule to the observation term: $\frac{\partial \mathcal{H}}{\partial \mathbf{z}} = \frac{\partial \mathcal{H}}{\partial \mathbf{x}} \frac{\partial \mathbf{x}}{\partial \mathbf{z}}$ with $\frac{\partial \mathbf{x}}{\partial \mathbf{z}} = \mathrm{diag}(\mathbf{x}_b \exp(\mathbf{z}))$. When the transport function $\mathcal{H}(\mathbf{x})$ is linear, as is commonly the case for Lagrangian models, it can be expressed as a matrix operator $\mathbf{H}$. In that case $\mathcal{H}^* := \left(\frac{\partial \mathcal{H}}{\partial \mathbf{x}}\right)^T$ reduces to $\mathbf{H}^T$, with T denoting the matrix transpose and the gradient of the cost function reads:

$$J'_{\mathrm{mod}}(\mathbf{z}) = \mathbf{1} + \boldsymbol{\Sigma}^{-1}\mathbf{z} + \mathrm{diag}(\mathbf{x}_b \exp(\mathbf{z}))\mathbf{H}^T\mathbf{R}^{-1}\left(\mathbf{H}(\mathbf{x}_b \exp(\mathbf{z})) - \mathbf{y}\right). \tag{15}$$

Line 132: why does using this substitution improve the convergence rate?

The substitution acts as a preconditioning of the optimization problem by transforming the prior covariance into the identity matrix. This makes the quadratic part of the prior curvature equal in all transformed directions (I think of it as geometrically turning an elongated prior ellipse into a circle), thereby improving the conditioning of the cost function, generally accelerating the convergence.

Line 158: some more explanation of how this equation was derived would be helpful. &

Line 161: you could add a real-world example of what this correlation means, presumably spatial correlation between emissions from different grid cells?

Yes, thank you, we agree with both points and revised the section to be more detailed.

In inverse modelling, prior emission uncertainties σ_b are often specified as relative uncertainties σ with respect to the corresponding prior emissions $\mathbf{x}_b$.

$$\sigma_b = \sigma \cdot \mathbf{x}_b. \tag{24}$$

For the lognormal formulation, these uncertainties need to be expressed in the transformed log space, where the errors are assumed to follow a Gaussian distribution. We define the log-space uncertainty s such that a deviation of s in log space corresponds to the specified relative uncertainty in linear space. Thus, considering $\mathbf{x} = \mathbf{x}_b(1+\sigma)$ and $\mathbf{x} = \mathbf{x}_b \exp(\mathbf{z})$, s can be expressed as

$$s = \ln(1+\sigma). \tag{25}$$

This formulation allows the relative uncertainty specified for the prior to be consistently represented as a symmetric uncertainty in log space. The error covariance matrix Σ, describing the uncertainty s of the log-transformed variable $\mathbf{z}$, can then be expressed as

$$\Sigma = \begin{pmatrix} s_1^2 & c_{12}s_1s_2 & \cdots & c_{1N}s_1s_N \\ c_{21}s_2s_1 & s_2^2 & \cdots & c_{2N}s_2s_N \\ \vdots & \vdots & \ddots & \vdots \\ c_{N1}s_Ns_1 & c_{N2}s_Ns_2 & \cdots & s_N^2 \end{pmatrix}, \tag{26}$$

where c_{ij} represents the correlation between the i-th and j-th elements, accounting for spatiotemporal correlations between emission errors in different grid cells.

Line 165: add a reference for the statement " land sinks are negligibly small".

Yes, we cited Harnisch and Eisenhauer, 1998!

Lines 166 and 167: add the full names of the EDGAR and GAINS databases.

Done ->

Figure 1 represents the a priori SF_6 flux distribution in Europe provided by the well-established bottom-up inventories from the Emissions Database for Global Atmospheric Research (EDGAR version 8; EDGAR, 2023; Crippa et al., 2023) and the Greenhouse Gas and Air Pollution Interactions and Synergies (GAINS) model (Purohit and Höglund-Isaksson, 2017), regridded to a spatial resolution of 0.25°.

Line 169: for reproducibility, it might be useful to explain how you produced the normal and lognormal fit lines, discussed at this point in the text and shown in Figure 1.

Yes, thank you , we add:

The orange dashed lines show the maximum-likelihood fits of the respective distributions, with the resulting probability density functions scaled to match the histogram counts.

to the capture of Figure 1.

Line 183: the footnote here could be included in the caption for Figure 2 instead.

Agreed, done!

Line 185: when reading this section, I wondered what observational uncertainty you were using, but this is covered later in sections 3.3 and 3.4. So it might be helpful to add a line here stating where observational error is defined.

Alright, we added:

The observational uncertainties are defined in Sects. 3.3 and 3.4.

Line 193: add a reference for the statement "SF_6 is almost inert up to the middle stratosphere".

Done -> we cited Kovács et al., 2017

Line 233: did you use any methods to test for convergence, in either the synthetic or real data tests?

We assessed convergence by monitoring the evolution of the cost function. We initially set the maximum number of iterations to 20, which is often sufficient, but observed considerable changes in the cost function and therefore increased the limit to 40 iterations. At 40 iterations, the cost function changed only marginally. For the real-data experiments, we nevertheless increased the maximum number of iterations to 100 as a precaution, to ensure that the differences investigated in our analysis were not caused by insufficient convergence.

Line 243: you could add some information on whether these prior and observation uncertainties are realistic, compared to the uncertainties 'traditionally' placed on these terms in inversions. Where does model error (e.g. uncertainties from the transport model) fit into these definitions of uncertainty?

&

Line 257: why was this observation error of 0.09 ppt chosen? And is it a reasonable assumption to apply this error value to all observations from all sites? I know this choice is currently under discussion in some inverse modelling groups, so I don't know the 'correct' answer for this myself, but it might be useful to comment on why these observational uncertainties were chosen for this work in particular.

Yes, we agree. The definition of uncertainties (and particularly the role of model error) is indeed an active area of research, and different choices are still being discussed within the inversion community. We are happy to discuss this, as it is also an aspect that we would like to investigate further in future work.

In general, allowing the observational error to vary in space and time based on model information or observational variability certainly has the potential to improve inversion results, as demonstrated by recent studies such as Steiner et al. (2024). Different inversion groups have proposed different approaches for estimating such errors. However, given the complexity of the problem, a simple estimate may in some cases be preferable, particularly when the aim is to isolate and investigate the general behaviour of an inversion system.

In our case, the main purpose of considering a broad range of uncertainties was to test the algorithm under different levels of observational constraint. For this purpose, we considered a single uncertainty value to be the most appropriate. The range of uncertainties tested in this study is largely within the range of values traditionally used or explored in previous inversion studies. To clearly demonstrate the artefacts that can arise under weak observational constraints, however, some combinations of relatively small prior uncertainties and large observational uncertainties result in particularly weak observational constraints and thus rather represent the lower end of the range considered. Similarly, somewhat larger prior uncertainties could have been explored (looking at the Gain metric), but an important objective was to show the behaviour under weak observational constraints. For the real-data inversion, the selected observational uncertainty was primarily guided by tests in previous studies (e.g., Vojta et al., 2025). We, however, also tested other uncertainty values, although these tests were limited to 40 iterations, and found no substantial differences in the general results.

We added/changed:

To evaluate the performance of the implementation of the lognormal distribution, we perform multiple inversions testing various prior uncertainties of the prior state variables (0.18, 0,34, 0.47, 0.59, 0.69) and observation uncertainties (0.02, 0.04, 0.06, 0.08, 0.1 ppt). These values largely fall within the range explored in previous inversion studies (e.g. Vojta et al., 2025) and allow us to investigate the inversion behaviour under different levels of observational constraint.

Line 253: what inversion period are you using? Just one inversion over the whole of 2020, or multiple shorter inversion periods across 2020?

We performed, one inversion over the whole of 2020 - we included:

For all performed inversions, the temporal interval was set to one year.

in Section 3.2.5

Line 254: EDGAR's distribution seems to be closer to a lognormal, so why was GAINS chosen as the prior for these tests, instead of EDGAR?

Indeed, looking at the lognormal fit, EDGARv8 appears somewhat closer to the underlying distribution. On the other hand, GAINS provides a broad range of emission values with fewer extreme high-emission values, which motivated our initial choice of GAINS for our experiments. We do not, however, claim that there is a clear criterion for preferring one inventory over the other.

Line 273: "Synthetic data experiment setup" is confusing as a subheading under the "Results" heading, should "setup" be removed?

Yes, indeed, thank you !

Line 313: missing space in "Fig.5j".

Yes, thank you – done!

Line 320: you note that the results are influenced by the specific random realization. Did you try rerunning these tests with different random perturbations, to see how the results varied?

Yes, we also tested a few additional random realizations (although not in the same detail) to ensure that the observed behaviour was not specific to a single realization. While the specific values of the Gain metrics were sensitive to the realization, mainly depending on whether larger perturbations occurred in regions that were better or less well constrained by the observations, the overall behaviour and the artefacts observed under weak observational constraints remained consistent. There were also realizations in which these artefacts were less apparent. however, the main purpose of these tests was to raise awareness that such artefacts can arise from the additional terms in the cost function associated with the mean and mode.

Line 360: Does this claim that negative emission increments favours error reduction apply in all areas or just these 4 cases and in the two representative grid cells discussed in the next paragraph? Did you investigate this for all grid cells?

Yes, yes, we also investigated other grid cells and found that this behaviour generally applies across the domain, although not always in the same magnitude. I think this can be best seen when comparing Figure 11a to 11c, where the pattern of small or even negative error reduction is largely consistent with regions of positive increments, and vice versa.

Figures 9 and 10: these figures may be easier to interpret if fewer histogram bins are used.

Thank you for this suggestion. We tested different numbers of histogram bins, but unfortunately reducing the number of bins did not substantially improve the readability of the figures. We therefore retained the original binning

Line 476: should "hx" be "H(x)" here?

No, in this case we intentionally used hx to denote a linear transport function in one dimension. To make this clear, we added:

To be consistent with the approximation around x_{MAP} let's consider now the cost function for the mode (see Eq. 14) in one dimension, assuming a linear transport function h

References:

Harnisch, J. and Eisenhauer, A.: Natural CF4 and SF6 on Earth, 25, 2401–2404, https://doi.org/10.1029/98GL01779, 1998

Kovács, T., Feng, W., Totterdill, A., Plane, J. M. C., Dhomse, S., Gómez-Martín, J. C., Stiller, G. P., Haenel, F. J., Smith, C., Forster, P. M., García, R. R., Marsh, D. R., and Chipperfield, M. P.: Determination of the atmospheric lifetime and global warming potential of sulfur hexafluoride using a three-dimensional model, Atmospheric Chemistry and Physics, 17, 883–898, https://doi.org/10.5194/acp-17-883-2017, 2017.

Super, I., Dellaert, S. N. C., Visschedijk, A. J. H., and Denier Van Der Gon, H. A. C.: Uncertainty Analysis of a European High-Resolution Emission Inventory of CO2 and CO to Support Inverse Modelling and Network Design, 20, 1795–1816, https://doi.org/10.5194/acp-20-1795-2020, 2020.

Steiner, M., Cantarello, L., Henne, S., and Brunner, D.: Flow-dependent observation errors for greenhouse gas inversions in an ensemble Kalman smoother, Atmos. Chem. Phys., 24, 12447–12463, https://doi.org/10.5194/acp-24-12447-2024, 2024.

Vojta, M., Plach, A., Thompson, R. L., Purohit, P., Stanley, K., O'Doherty, S., Young, D., Pitt, J., Arduini, J., Lan, X., et al.: Quantifying European SF6 emissions from 2005 to 2021 using a large inversion ensemble, Atmospheric Chemistry and Physics, 25, 15 197–15 243, https://doi.org/10.5194/acp-25-15197-2025, 2025