Egusphere-2025-5086

Response to Reviewer 1

Yamanouchi et al.

We thank the reviewer for the constructive comments. This document summarizes our responses and documents the changes made to the manuscript. All reviewer comments are written in black font and were copied directly. All responses are written in red font, and all references to changes in the revised manuscript are written in **bold red font highlighted**.

General comments

This paper presents two methods for improving the Gaussian plume atmospheric dispersion model with the ultimate goal of improving atmospheric inversions of emission source and rate at the industrial facility scale. First, they propose a method of weighting neighboring atmospheric stability classes based on surface roughness and Obukhov length that eliminates the need to choose one specific class. Second, they propose a method for updating the emission source location used to run the Gaussian plume model by minimizing mismatch between simulated concentrations and observed concentrations. They find that the blended stability classes result in model performance roughly between the two stability classes when used independently. Furthermore, they find that updating the source location often results in closer alignment of modeled and observed concentrations.

Overall, I think the manuscript is well written and proposes two useful updates to the Gaussian plume model. However, the manuscript would benefit from a more explicit definition of the methods. As currently written, the methods can be understood at a high- level but not enough details are provided to fully understand the nuances or to reproduce them. Some examples are provided below in the “specific comments” section. Furthermore, I believe that the manuscript would benefit from a clearer description of how the proposed methods would handle situations where multiple sources are emitting simultaneously (and are separated enough to result in two distinguishable plumes). It is not clear to me if multiple sources were emitting simultaneously during the Petrolia controlled releases, and in Figure 2, it appears as if there could be two Gaussian distributions in the observed data (indicating two sources). More clarity about how this situation is handled would benefit the manuscript.

Specific comments

- Paragraph starting at L71: Several studies have explored alternative approaches that partially address the limitations discussed in this paragraph. Briefly mentioning how this study relates to these alternative approaches would help highlight its contributions. For example, Cartwright et al. 2019 (https://doi.org/10.5194/amt-12- 4659-2019) and Newman et al. 2025 (https://doi.org/10.1214/25-AOAS2101) learn the dispersion coefficients of the Gaussian plume model within their inversion, avoiding the use of atmospheric stability classes. Studies like Jia et al. 2025 (https://doi.org/10.1038/s41598-025-99491-x) and Daniels et al. 2025 (preprint - https://doi.org/10.48550/arXiv.2506.03395) use the Gaussian puff dispersion model to capture time-varying wind characteristics and an event-based or moving time window approach to account for time-varying emission characteristics. A short discussion of how the proposed method differs from or builds on these approaches would benefit the manuscript.

  We thank the reviewer for this very good point of discussion; **a summary of what is discussed below was added to the manuscript in Section 1 (introduction).**

  First, compared with approaches that learn dispersion coefficients within the inversion, such as Cartwright et al. (2019) and Newman et al. (2025), our methodology preserves a clear and physically interpretable link to boundary-layer physics. In those studies, the horizontal and vertical dispersion parameters are treated as free (or at least semi-parametric) variables/functions to be estimated jointly with emissions; this increases flexibility and can also potentially reduce bias associated with incorrect stability classification, but it also shifts part of the physical atmospheric description into the statistical inference (inversion) layer. In contrast, our blending framework remains within the classical Gaussian plume formulation while replacing discrete stability classes with a continuous weighting based on physically meaningful quantities (e.g., Obukhov length). Rather than learning σ_y and σ_z from concentration data alone, we provide a method to regularize dispersion using micrometeorological information, thereby reducing equifinality between emissions and transport and improving interpretability and transferability across sites.

  Second, relative to Gaussian puff implementations such as those used by Jia et al. (2025) and Daniels et al. (2025), our method addresses a different structural limitation of the steady-state plume model. Puff models resolve temporal variability in wind and emissions by superposing time-evolving puffs, often within event-based or moving-window inversions. This enhances realism under nonstationary conditions but typically retains conventional parameterizations of turbulence (often still tied to stability categories or empirical σ-parameterizations). Our approach, by contrast, targets the discretization error inherent in the stability-class framework itself. By continuously blending the two most representative classes, we reduce artificial discontinuities in σ_y and σ_z across stability transitions without incurring the computational overhead of a full puff simulation.

Importantly, our framework can be viewed as complementary to both lines of work. It could (i) serve as a physically informed prior or structural constraint within dispersion-learning inversions, limiting overfitting of σ parameters, or (ii) be embedded within a Gaussian puff model to provide a continuous, meteorology-driven parameterization of subgrid turbulence. In this sense, the method offers a middle ground: more physically grounded than purely data-driven dispersion learning, yet computationally simpler and more transparent than fully time-resolved puff systems.

- Equations 1 and 2: It would make it easier to interpret the model if the dimensions of each variable were provided. This would help differentiate between variables in concentration space versus variables in emission/flux space. It would also help the reader understand the size of the modeling problem if you provided the number of sources and the number of observations used in a typical application of this model.

  Dimensions (units) have been added in text.

- Equations 1 and 2: If I understand the model correctly, then this formulation assumes that the errors between the observed, Y, and simulated concentrations, M*F_a, are Gaussian. Similarly, the prior distribution on the emission rates, F_a, is also Gaussian. It may be worth noting these distributional assumptions.

  Assumption of normally distributed observational and emission uncertainties have been explicitly added in text.

- Equations 1 and 2: Please also state if the observation operator is formed explicitly by, e.g., forward simulating a unitary rate from the a priori specified source locations. If this is the case, please discuss how source locations are identified a priori.

  The observation matrix M is formed explicitly **(added in the paper at Section 2.2).** A priori source locations were taken to roughly be the center of the emission site. Care was taken not to be overly presumptuous about the a priori to test the source translating methodology. While developing the methodology, for some oil and gas facilities, we also tried examining the sites via satellite imagery to see where industrial activity may be taking place, as a rough guess.

- L120-122: How are sources of different diameters modeling using the gaussian plume model? Many implementations of this model assume a point source.

  The Gaussian plume model implementation in Polyphemus allows specifying diameters of sources. As the primary purpose of this study was to develop and test the stability class blending and the source locating algorithms, source diameters were chosen to roughly match the overall sizes of the source sites. We also took steps to not overfit the sizes, to better study effects of the stability classes chosen (e.g., more stable stability classes like E or F will usually result in narrower plumes, while unstable classes will result in wider plumes; to better test the stability class choice, the diameter of the source was kept constant).

- L128-129: Please specify how the 48% uncertainty value is incorporated into the model. Does this value make up the diagonals of B?

  The 48% value is the overall uncertainty. The values included on the diagonal of B represent the uncertainty associated with the prior emission estimates. In our case, we prescribe a fixed relative uncertainty for the prior. I usually assume values on the order of 80 to 100%, which allows sufficient flexibility for the inversion system to adjust the emissions based on the observations (although I am not sure which value you used in your inversion, Shoma).

  The 48% uncertainty refers instead to the overall uncertainty of the posterior emission estimates. This value reflects the combined uncertainties in the inversion results, including the limitations of the Gaussian plume approach, which relies on several assumptions regarding wind conditions, atmospheric stability, and other meteorological factors.

Therefore, the 48% value is not used to populate the diagonal elements of B, but rather represents the estimated uncertainty associated with the final posterior emissions. Discussions on inversion uncertainty, and other details of this analysis are presented in Ars et al. (2017), as referenced in the paper.

- L144-146: Please provide more details about the plume rise methodology or provide a reference that discusses it in more detail.

  The core model used in this study is an “off-the-shelf” Gaussian plume dispersion model in the Polyphemus suite. The reference to a paper discussing this model suite is given in text. **The reference to this paper was also added in this section for clarity.**

- Section 2.2.1: please state exactly how the weights are calculated for blending the resulting concentration predictions from the two selected stability classes.

  The weights are calculated by the distances to the nearest stability curves. The distances were calculated holding z_0 (y-axis) constant (i.e., the horizontal distance), as surface roughness is assumed not to change (at least not in the time span relevant to these analyses). **This was clarified in the paper.**

- L229: “Here, the a priori source location taken to be several hundred meters north of the actual source.” Please discuss if this was an intention modeling choice to test the source-adjustment procedure, or if this is your interpretation of the model output. That is, your a prior source locations were truly your best guess before running the inversion.

  This was intended to test the source-adjustment procedure. Indeed, the development of this methodology was done blind, meaning the true location/emission rates were not known while developing this algorithm.

- Figure 2 and 3: It is not exactly clear to me which dimension the “distance” axis refers to. Is this the distance along the dimension perpendicular to the wind vector?

This distance refers to the distance along the path of the transect (along the path driven by the car).

- Table 1 and 2: Indicating which releases were controlled releases vs. real emissions would make it easier to interpret this table. For Table 2, it would be very useful to list the true emission rate and the true source location (or perhaps the distance between the “optimized” source estimate and the true source location).

  The tag to indicate controlled release transects have been **added/made more apparent**. The true emission rates of the controlled release experiments are given and discussed in text.

- Figure 6: There is, in my opinion, weak justification for omitting the 11,000 kg/day release from diagnostics / statistics. Could the authors please provide additional justification for omitting this release?

  This point was dropped as it was a clear outlier, being an order of magnitude higher than other data points. Furthermore, when we examined the transect and model, the plume lacked a Gaussian shape and it was clearly not a good data point. Had this been an observation that we took outside of a controlled release experiment, it would not have met our quality control standards and would have been dropped. We kept this in the analysis and mentioned it here as a point of discussion. However, it was included in the figure, and as discussed in text, correlation was not severely affected with or without the inclusion of this point. However, it was included in the figure, and as discussed in text, correlation was not severely affected with or without the inclusion of this point.

- More details about the controlled release experiment would benefit the manuscript. For example, where the source locations known a priori, and if so, where they used as your initial guess in your inversion algorithm? Was there just one source, or multiple sources emitting at the same time?

  A priori source locations were taken to be the center of the emission site. Care was taken not to be overly presumptuous about the a priori to test the source translating methodology. Multiple sources were emitting at the same time; we used our algorithm to find and isolate the source that we observed in our transects. If multiple plumes were observed within one transect, then the algorithm would shift over the source to the closest peak. This meant that with multiple plumes, we needed to manually nudge the sources over to and rerun the model to try and figure out the source locations for each of the plumes.

Additionally, if the sources are very close to each other, we design these experiments such that we drive past and transect the plumes at distances where the signals of different sources blend into one plume so the Gaussian plume modelling approach is valid and usable. This approach is especially useful in real-world situations, e.g., when quantifying facility emissions (in such cases, whether the emission was from e.g., one chimney stack or the one next to it is not important).

- Table 2: the direction of translation is not very meaningful without seeing a site diagram. As currently presented, it gives the impression that source locations can be moved in a 2D plane, but they can only be moved in the direction of the transect. This fact is mentioned in the text article, but also stating it in Table 2 would help interpret this table.

  Fixed.

- Figure 4: Do the pins show the true controlled release source locations, the original guess used in the inversion algorithm, or the nudged source location obtained after running the inversion the first time. Clarifying this point would greatly improve the interpretability of this figure.

  Fixed.

- Section 2.2.2 and Section 2.3: the manuscript would benefit from a more rigorous explanation of the methods used to adjust the source location and average repeat transects. As currently written, the procedures can be understood at a high-level, but enough details to fully understand and reproduce the methods are missing.

  Clarifications (minor wording on 2.2.2 and major additions to 2.3) were added to the manuscript.

- Table 2: Using RMSE and R^2 as an evaluation metric is a good way of showing how well the simulated concentrations line up with the observed concentrations. However, the source optimization procedure is essentially adding an extra parameter to your model (a distance correction along the dimension of the transect). The authors are then evaluating this more flexible model on the same data used to learn the value of the added parameter (if I understand correctly). This point is worth noting in the manuscript, as this procedure will always increase R^2 and decrease RMSE according to statistical theory and could potentially lead to overfitting.

  This discussion **has been added to Section 3.2.**

Technical corrections

- L39: Typo, “off” instead of “of”.

  Fixed.

- L50-53: This sentence would benefit from references supporting the 25-90% numbers.

  Several references have been added.

- L65-67: It is worth noting that this steady-state assumption applies to the Gaussian plume model. Another variant of the Gaussian dispersion model, the Gaussian puff model, relaxes the steady-state assumption. See, for example, https://doi.org/10.1038/s41598-025-99491-x for a comparison of the two models. In general, differentiating between the “Gaussian plume model,” “gaussian puff model,” and the “Gaussian dispersion model” throughout the manuscript would improve clarity.

  Discussion explicitly mentioning the Gaussian puff model, which uses time-varying fields, **have been added**, along with a statement clarifying that the Gaussian plume model used in this study assumes the steady-state case.

- L167-168: “if the skewness in the transects were sufficiently close to each other.” Please clarify if you are coming the skewness between the observed concentrations and the simulated concentrations from the Gaussian plume.

  Wording was fixed to clarify this.

- L170: why is an upper bound on the autocorrelation used?

  If autocorrelation was higher than 0.955, then no correction was done. **This clarification was added in text.**

- Table 1: The [3] label does not appear within the table.

  Fixed.

- While the manuscript is overall well written, the language is somewhat informal in places. For example, the authors occasionally use quotation marks to describe or define an idea loosely, rather than using more formal, precise language.

  Many of these were fixed (parenthesis dropped in cases where it was deemed not necessary, and in other instances, more precise wording/explanations were added).

- There are a few instances of unclosed or unopened parenthesis.

  Fixed.