Response to Reviewer 1

We would like to sincerely thank the reviewer for the thorough evaluation of our manuscript and for the constructive comments and suggestions provided. We deeply appreciate the time and effort devoted to reviewing our work. The feedback has been invaluable and has contributed significantly to improving the quality and clarity of the manuscript. Below, we address each of the comments in detail.

Major Comments

- **Reviewer Comment:** ***First, while the authors provide general background on radiative transfer models and cite CRTM and previous versions of RTTOV, there is no direct comparison or benchmarking of their approach against existing fast RTMs. This omission limits the reader's ability to evaluate the benefits or drawbacks of the proposed method relative to established techniques.***

  Authors Response: The referee is right in noting that a benchmark comparison between established fast RT models—such as CRTM and RTTOV with LASSO-induced sparsity—is of interest. However, this falls outside the intended scope of the present study. Such a comparison would require a dedicated and thorough analysis, given the fundamental differences in their core parameterization strategies: RTTOV employs an additive gas-by-gas optical depth parameterization, whereas CRTM relies on a joint, global parameterization of gas absorptions.

  The main focus of this work is a methodological proposal to evaluate and improve the computational efficiency of the optical depth parameterization in the RTTOV v13 model. This is achieved by replacing the standard OLS regression with LASSO regression to induce sparsity, using inferential statistical techniques to discard gases that are not relevant for the numerical approximation of transmittances, and assessing the advantages and limitations of this approach within the RTTOV scheme.

  Nevertheless, we acknowledge the value of such a comparative study and plan to extend our approach in the future to develop sparsity-driven parameterizations for joint gas absorption schemes, which may be applicable to CRTM or similar frameworks. This represents a promising direction for future research.

- **Reviewer Comment:** ***Second, the method relies on threshold parameters to determine the relevance of gases, yet there is no guidance or sensitivity***

***analysis provided on their selection. Since the method's validity depends on safely discarding certain absorbers, this is a critical omission.*.**

Authors Response: We agree that performing a sensitivity analysis on the inclusion or exclusion of individual gases can provide additional insight into their relevance, as has been done in other studies that assess gas importance based on brightness temperature variability.

Currently, in our approach, the decision to discard a gas is not based on arbitrarily chosen statistical thresholds. Instead, it is grounded in the structure of RTTOV's optical depth parameterization and is implemented automatically through a statistical framework based on confidence intervals, as detailed in Subsection 4.1.

Specifically, we compare the results of RTTOV using all gases that exhibit absorption lines in a given channel against those obtained using only the subset selected via our method. In our statistical thresholding approach, we construct confidence intervals for the transmittance at each pressure level for each gas, aiming to contain the true transmittance value with a high confidence level of $p = 1 - \alpha = 1 - 10^{-6}$. This is a deliberately strict threshold, ensuring reliable inferences.

A gas is considered negligible at a given pressure level only if two conditions are simultaneously satisfied: (1) the confidence interval length (or standard deviation of transmittance) is smaller than a relative tolerance $\epsilon_1 = 10^{-6}$, which is stringent considering transmittance lies within the range $(0, 1]$; and (2) the corresponding optical depth is smaller than $\epsilon_2 = 10^{-6}$, allowing us to approximate the transmittance by 1 with a relative error less than ϵ_2, as explained in Case III (2.2.5). Both ϵ_1 and ϵ_2 are fixed relative tolerances, and any future adjustment would likely involve even stricter values.

For a gas to be entirely discarded, both conditions must be satisfied across all pressure levels, meaning that the gas must exhibit transmittance values close to 1 throughout the atmospheric column with a confidence probability of $(1-\alpha)^{100} \approx 0.9999$ (assuming independence between layers, consistent with the RTTOV parameterization scheme). This constitutes a qualitative assessment informed by rigorous statistical inference; it is not a direct exclusion based solely on the numerical values of threshold parameters. That said, the practical guideline for choosing these parameters, as inferred from their construction and intended purpose, is that they should be sufficiently small.

In the revised version of the manuscript, we extend our numerical results using smaller values for these tolerances, which allow for better fits, and we provide the corresponding analysis.

Reviewer Comment: ***Additionally, the manuscript does not discuss scenarios where the assumptions of the method might break down — for example, in unusual atmospheric compositions, extreme pollution events, or volcanic emissions.***

Authors Response: This is an important and insightful point for the development of a robust Fast-RT framework. The proposed approach could indeed be extended to handle more complex scenarios, such as treating SO2 as a variable gas as in RTTOV for volcanic environments (rather than a fixed gas as in our current setting), incorporating additional aerosol types to account for extreme pollution events, or refining the treatment of the water vapor continuum, which is not included in our current configuration. Nonetheless, the core idea of our methodology—based on a separation between six variable gases and 22 fixed gases identified as the dominant absorbers—remains valid as the foundation for the proposed absorption parameterization scheme. These potential extensions represent promising directions for future development within the same methodological framework.

However, as stated in our response to the first comment, the present work is primarily a methodological proposal aimed at improving the computational performance of RTTOV by introducing a sparse regression-based optical depth parameterization for gas absorption and enabling automatic gas selection. The study is not intended to enhance the classical RTTOV performance under diverse or extreme atmospheric conditions. For such considerations, we refer to the official RTTOV v13 Science and Validation Report by Saunders (2020), which defines the core assumptions and validation framework of the model.

- **Reviewer Comment:** ***Third, LASSO parameter tuning is conducted via a basic grid search, but the authors do not provide any justification for this choice nor discuss why alternative standard methods (such as cross-validation) were not pursued.***

  Authors Response: We appreciate the reviewer's insightful comment regarding our use of grid search to tune the LASSO regularization parameter. Our initial approach aimed to provide a baseline method that is straightforward and widely understood. We acknowledge that alternative techniques, such as cross-validation, are commonly used for parameter tuning; however, in our specific context, cross-validation is not directly applicable. A brief discussion of this point has been added starting at line 243 in the revised manuscript.

  To address this limitation, we will incorporate in a revised version of the manuscript alternative advanced methods, including model selection criteria such as the Bayesian Information Criterion (BIC) and parametric bilevel optimization frameworks. One of the authors has worked extensively on these approaches in recent years, and they offer promising avenues for efficiently and reliably selecting the optimal regularization parameters.

- **Reviewer Comment:** ***Finally, while the authors claim that their approach leads to a "substantial reduction in computational cost," no quantitative analysis is provided to support this. There are no measurements or estimates of runtime or memory savings, nor is there any discussion of what constitutes an acceptable or unacceptable reduction in accuracy for practical applications. The results show mixed performance — with improved RMSE in some VIIRS bands but increased errors in others — but there is no clear guidance on when the method is expected to perform well or***

poorly.

Author Response: The qualitative support for this claim is grounded in the observed sparsity level achieved in the transmittance parameterization. Specifically, the reduction in runtime for evaluating the parameterized transmittance function is directly proportional to the reduction in the number of active parameters, as fewer predictor evaluations are required. Similarly, memory usage is also reduced proportionally; however, in this context, memory savings are of limited practical relevance due to the inherently low memory requirements of a full parameterization. In a revised version of the manuscript, we will refine Section 5.2 to improve clarity by incorporating a quantitative comparison between the percentage reduction in the number of parameters and the corresponding decrease in runtime.

Regarding model performance, we agree that the RMSE of transmittance and BT alone are insufficient to determine whether a Fast-RT model performs well or poorly. For this reason, a second level of validation was included in the manuscript, consisting of the computation of relative errors in BT estimation: average between $\mathcal{O}(10^{-5})$ and $\mathcal{O}(10^{-3})$, and maximum relative errors which are between $\mathcal{O}(10^{-2})$ and $\mathcal{O}(10^{-4})$. From this error, M7 performs poorly for all methods due to an unacceptably large error; the rest of the methods perform comparably to standard RTTOV. These errors are obtained by comparing the BTs predicted by the Fast-RT models against those derived from radiances computed using a Line-by-Line model, which serves as the reference or 'ground truth'. This approach aligns with the core objective of Fast-RT methods: approximating the output of LBL models.

In the revised version of the manuscript, the water vapor continuum absorption was disabled in the LBLRTM output in order to clarify the comparison and improve the consistency between simulations, since this absorption component is not explicitly parameterized in the Fast Radiative Transfer (Fast-RT) models under evaluation. This modification enables a more accurate benchmarking of the line-by-line reference against the Fast-RT approximations for a baseline model, whose methodology can be extended to more comprehensive models in future work. Additionally, the analysis includes a comparison between the brightness temperature residuals and the Noise-Equivalent Differential Temperature (NEdT) of the VIIRS sensor. This evaluation framework serves as an effective diagnostic for validating the performance of the Fast-RT schemes prior to their implementation in satellite radiance assimilation workflows.

Structural Comment

Reviewer Comment: ***The manuscript currently devotes substantial space in Sections 2 and 3 to background material on radiative transfer theory, line-by-line modeling, and general Fast-RT model formulations. While this content is clearly written and technically accurate, much of it summarizes well-established concepts that are not essential for understanding the specific methodological contribution of this paper. The level of detail presented here feels more appropriate for a thesis or tutorial-style document rather than a journal article***

focused on a specific methodological advance. To improve readability and focus, I recommend substantially condensing these sections in the main text or moving parts of them to an appendix. This would allow the reader to reach the core methodological development (Section 4) more efficiently, without sacrificing completeness for readers who may need additional background.

Author Response: We agree with this observation. In the revised manuscript, we will present the relevant theoretical background on radiative transfer and Fast RT models in a more concise manner, in order to improve readability without sacrificing the necessary foundations to understand the methodological development.

Minor Comments

- L.10: *retrieval* → *retrievals*
  Done.
- L.28: *model* → *modeling*
  Done.
- L.38: add PCRTM reference:
  Liu, X., Smith, W. L., Zhou, D. K., Larar, A. M., Huang, H.-L., Ma, X., & Strow, L. L. (2006). Retrieval of atmospheric profiles and cloud properties from IASI spectra using super-channels. *Atmospheric Chemistry and Physics*, **6**, 255–265. https://doi.org/10.5194/acp-6-255-2006
  Done.
- L.39: Even though RTTOV is more efficient than line-by-line models, it remains prohibitively expensive for operational ~~use in small to medium-sized agency~~ use cases.
  Done.
- L.42: RT model, ~~similar to models based on neural networks~~.
  Done.
- L.43: model ~~further~~ less computationally
  Done.
- L.44: These decisions must account for the multitude of possible combinations and trade-offs, ~~which is why large meteorological agencies rely on~~ **and are typically made by** expert teams ~~to identify an optimal configuration of parameters and gases for the Fast RT model~~.
  Done.
- L.50: cite or remove 'various large-scale applications'
  Response: The references in lines 51–57 pertain to large-scale applications of LASSO for variable selection in the context of radiative transfer.
- L.51–57: mentions multiple papers that perform 'variable selection' without much context. Not sure what to do with this information.
  Response: Since the paragraph begins with 'In the context of radiative transfer,' the

citations, as mentioned in the previous comment, refer to large-scale applications of LASSO.

- L.58: specify what those 'parameters' are in Fast RT models.
  Changes: ...we target the automatic selection of gases and **optical depth predictors** ~~parameters~~ in Fast RT models by inducing sparsity in the **weight predictors** ~~parameters~~ using LASSO regression.

- L.64: Has LASSO been applied to other RTM models?
  Response: To the best of the authors' knowledge, the use of LASSO regression specifically for modeling optical depth or transmittance within radiative transfer models has not been previously documented in the literature.

- L.65: Remove section 1.1 title.
  Done.

- L.82: what is 'carbon powder'?
  Response: "Carbon powder" is a fine particulate form of elemental carbon, typically produced by incomplete combustion or pyrolysis, and includes substances like soot and black carbon that contribute to atmospheric aerosols.

- L.135: ml is ~~de~~ **the** number
  Done.

- L.245: how was that value chosen? $mse(\lambda) < \mathbf{2}\, mse(0)$
  Response: This represents a relative tolerance that specifies how close the mean squared error (MSE) of the LASSO solution should be to that of the ordinary least squares solution. Since $mse(\lambda) > mse(0) > 0$, the condition can be rewritten as $\frac{mse(\lambda)-mse(0)}{mse(0)} < 1$. We then select the largest value of λ among the candidates that satisfies this inequality. This was replaced by a model selector based on the Bayesian Information Criterion.

- L.291–299: move to appendix
  Not done.

- L.375: "This suggests that the inclusion of statistical thresholds and LASSO regression in RTTOV v13 slightly affects the accuracy of the transmittance approximation, either improving or worsening it, but the overall variation in error remains negligible."
  The error in M7 and M8 increases by about 40% with the LASSO method. Why is that negligible?
  Response: We acknowledge the increase in error for M7 and M8; however, the absolute values of the errors remain small. The overall accuracy of the transmittance approximation is not significantly affected, so we consider the variation to be minor in practical terms. To clarify this in the manuscript, we made the following correction in line 376: ~~remains negligible~~ **does not significantly impact the quality of the transmittance approximation**.

- L.395: Don't understand the 'order of magnitude' comparison. For M10 the error increases from 0.89 to 1.4

Response: We agree with the observation. The error difference in M10 does not represent an order of magnitude. Therefore, this channel has been excluded from the statement in the revised manuscript.

- L.406: "These findings suggest that while the proposed methods are generally comparable to RTTOV v13 in terms of accuracy, there are specific channels where improvements or further adjustments in the statistical threshold parameters may be necessary to enhance precision if needed."
  This work should have been part of this study.
  Response: We agree with the observation. We have added the numerical results for different values of ϵ_1 in the revised manuscript.