

Author's Response:
**GHGPSE-Net: A method towards spaceborne
automated extraction of greenhouse-gas point
sources using point-object-detection deep neural
network**

Yiguo Pang^{1,2}, Denghui Hu¹, Longfei Tian¹, Shuang Gao¹, and Guohua Liu^{1,2}

¹Innovation Academy for Microsatellites of Chinese Academy of Sciences, Shanghai, China

²University of Chinese Academy of Sciences, Beijing, China

Contact author: Yiguo Pang (pangyiguo21@mailsucas.ac.cn)

Corresponding author: Guohua Liu (liugh@microstate.com)

February 18, 2026

We sincerely thank the editor and reviewers for the valuable comments and suggestions. We are glad to receive the positive feedback from the reviewers. We have carefully considered all the comments and revised the manuscript accordingly.

Our detailed point-by-point responses to reviews are given as follows, where the reviewer's comments are in *blue italics* and our responses are in normal text.

Point-by-point Response to Reviewer #1

The authors have adequately addressed the major concerns raised in the previous round of review, and the overall quality of the manuscript has been significantly improved. I have no further major comments on the scientific content or the methodology. However, the manuscript requires a final round of careful proofreading to ensure professional presentation. There are several grammatical inconsistencies and typographical errors that need to be addressed before publication.

We appreciate the reviewer’s valuable comments, which have greatly helped improve the readability of our manuscript. We are grateful for the recognition of the improvements we have made. We apologize for the grammatical errors and typos, despite our proofreading efforts. Thank you for pointing out the technical issues in such detail. We have addressed all the points raised and have also corrected additional grammatical errors and typos.

1.Line 45: "...equipped with a Onboard intelligent..." should be corrected to "an Onboard".

We’ve corrected this point accordingly.

2.Line 66: "...model referred GHGPSE-Net..." should be corrected to "...model referred to as GHGPSE-Net..."

We’ve corrected this point accordingly.

3.Line 92: "...spatial-temporal limited dataset..." would be better phrased as "spatiotemporally limited dataset" or "spatially and temporally limited dataset".

We’ve corrected this point accordingly.

4.Line 155: "Here, $\langle I_{bg} \rangle$ denote the mean concentration..." Subject-verb agreement error. It should be "denotes".

We’ve corrected this point accordingly.

5.Line 190: "Deep neural network is a class..." should be "A deep neural network is a class..." or "Deep neural networks are a class..."

We’ve corrected this point accordingly.

6.Line 209 (Eq. 7): Ensure the formatting of the condition " $a_i \geq 0$ " is consistent with the text.

Here, a_i in Eq. 7 is a scalar, which denotes any random member from the vector \mathbf{a} . This expression is consistent within the text. To make this point clearer, we revised Eq.7 from

$$\mathcal{L}(\boldsymbol{\theta}; \mathbf{I}, \hat{\mathbf{I}}) = \sum_{p=1}^P \left(I(\mathbf{x}_p) - \hat{I}(\mathbf{x}_p; \boldsymbol{\theta}) \right)^2 \quad s.t. \quad a_i \geq 0. \quad (1)$$

to

$$\mathcal{L}(\boldsymbol{\theta}; \mathbf{I}, \hat{\mathbf{I}}) = \sum_{p=1}^P \left(I(\mathbf{x}_p) - \hat{I}(\mathbf{x}_p; \boldsymbol{\theta}) \right)^2 \quad s.t. \quad a_i \geq 0, \forall a_i \in \mathbf{a}. \quad (2)$$

7.Line 260: "...which is generated..." (Check tense consistency with the rest of the paragraph).

Thank you for the comment. We have revised the tense to past tense to better describe that the SMARTCARB dataset was previously generated by other studies. The revised sentence is "...which was generated by..."

8.Line 289: "...metrics reach maximum with..." suggest changing to "...reach a maximum..." or "...reach their peak...".

We've corrected this point accordingly.

9.Line 365 (Figure 6 Caption): "The observation timestamps are provided in UTC." – This is clear, but ensure the x-axis labels in Figure 6 are legible.

We've revised the x-axis labels in Fig. 6 to make them more clear and readable.

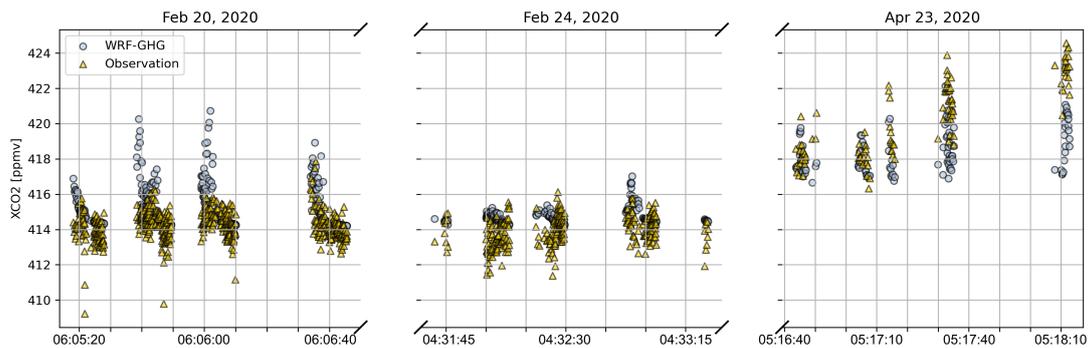


Figure 1: Revised Fig. 6 with sparser x-axis labels.

Line 358: "...extract emission source." should be plural: "...extract emission sources."

We've corrected this point accordingly.

10.Line 378: "...increases the localization performance..." – phrasing is acceptable, but "improves" might be more natural than "increases" for performance.

We've improved this point accordingly.

11.Figure 5 (Y-axis label): The label "Temperautre error[K]" contains a spelling error. It should be "Temperature".

Thank you for pointing out this typo with such detail. We've corrected this typo accordingly.

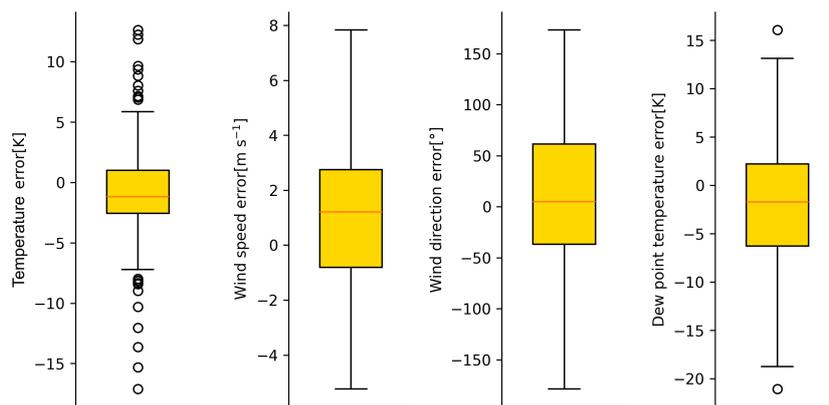


Figure 2: Revised Fig. 5 with corrected y-axis label.

12.Line 322: "SMARTCARAB" is a typo. It should be "SMARTCARB".

We've corrected this point accordingly.

13.Line 334: "1.2 to 1.6 6 ktCO₂ hr⁻¹". There appears to be a stray digit "6" before the unit

Thank you for your kind suggestion. However, after carefully checking the original text, we found that there is no stray digit "6" before the unit. I supposed this issue may be arised by the PDF renderer. The original text is "1.2 to 1.6 ktCO₂ hr⁻¹".

14.Line 383: "...more spial context..." is a typo. It should be "spatial".

We've corrected this point accordingly.

15.Please ensure that the citation style for "et al." is consistent throughout the manuscript (e.g., Line 51 vs Line 136).

Thank you for your suggestion. We assume the reviewer refers to the mixed use of "et al." in in-text citations v.s. parenthetical citations, e.g., "Author et al. (2020)" v.s. "(Author et al., 2020)". However, after careful checking of the journal's submission guidelines, we find that both forms appear acceptable. Therefore, to ensure full consistency with the GMD's style policy, we prefer to leave any necessary adjustments to the proofreading stage, where the editorial office can apply the final formatting.

Point-by-point Response to Reviewer #2

Dear authors and GMD team, Most of my comments have been addressed satisfactorily. There are a few items that are still left to clarify in my view

Thank you for your comments. We are glad to receive the positive feedback from the reviewer. We have revised the manuscript according to the comments, and we hope that the

revised manuscript is now clearer and more readable.

According to the authors, the augmented datasets are split "into training, validation, and testing sets in a 3:1:1 ratio". The authors do not specify if this split is random, and if so, any potential consequences of this - i.e. whether the model achieves high scores due to skill, or due to the test samples being very similar to the training ones. Testing on unseen OCO-3 and SmartCarb samples is a significantly better test of skill - for example comparing performance of the model on the test set of synthetic OCO-3 augmented samples, vs true OCO-3 observations shows a decrease in performance. There is no assessment of the skill of the model on the augmented dataset with SMARTCAB specific noise, so this comparison is not possible.

The dataset was randomly split into training, validation, and testing sets. We acknowledge that a purely random split may introduce statistical similarity among subsets, potentially resulting in optimistic performance estimates on the synthetic test set. We agree that evaluation on independent datasets provides a more rigorous assessment of model skill. The performance degradation on real observations from Orbiting Carbon Observatory-3 (OCO-3) and SMART-CARB indicates limited generalization and demonstrates differences among simulations and real-world observations. This issue was discussed in the previous version of the manuscript (L395–407). In this version, we have now further emphasized the potentially overestimated performance on the synthetic dataset at this point in the Discussion section — "The decreased performance on the two independent datasets also suggests that evaluations conducted exclusively on a single synthetic dataset may systematically overestimate model performance, due to the inherent statistical homogeneity within such datasets."

Regarding training with SMARTCAB-specific noise, such an experiment could help analyze contributing factors. However, we have performed the experiment and evaluated the performance on noise-free SMARTCAB case, and the results show that quantification performance remains largely unchanged compared to the noisy-SMARTCAB case. This suggests that the model's performance on the SMARTCAB dataset is not significantly affected by the noise patterns specific to SMARTCAB, and that other factors may be more influential in constraining model generalization. This result was described in the previous version of the manuscript (L334-L337).

Moreover, at the current stage, several limitations remain. (1) All these experiments are based on simulations and lack large-scale real-world data for validation, making it less meaningful to perform comparisons involving synthetic noise patterns. (2) Multiple additional factors may constrain model generalization, including plume morphology, background pattern, etc. (3) A comprehensive investigation of these factors seems beyond the scope of this work, which focuses on proposing a new deep learning framework trained on synthetic data and providing an initial evaluation using independent datasets. Thus, we will leave such an analysis as future work, with large-scale evaluations on factors limiting the generalization of the model.

2) I would like clarification on why, from the OCO-3 dataset, only eight observations with a clear background are selected.

Thank you for the comment. Due to the limited spatial coverage of OCO-3, only a small number of snapshots exhibit clear & spatially isolated plume structures. A large fraction of the observations are affected by one or more of the following issues: (1) spatial discontinuities caused by cloud contamination; (2) weak plume signals with low emission strength; (3) strong interference from background anthropogenic emissions. Consequently, very few observations are selected for robust evaluation. Similar limitations have also been reported in large-scale studies. For example, in Lin et al. (2023), a heuristic-based filter is implemented to filter suitable power plant targets for quantification, resulting in only 50 targets from 4 years of combined OCO-2 and OCO-3 observations.

We further elaborate on this point in Section 2.3.3 — "Due to the relatively limited spatial coverage of OCO-3, most of the observations are faced with issues such as discontinuities by cloud contamination, weak plume signals, and background anthropogenic emissions, we only select 8 observations covering 4 major power plants ($> 1 \text{ ktCO}_2 \text{ hr}^{-1}$) with relatively clear and spatially isolated plumes from the OCO-3 Level 2 v10.4r dataset (OCO-2/OCO-3 Science Team et al., 2022) for reliable evaluation."

3) The distributions you use for data augmentation, described in 2.1.3, are shown in the Supplement, which is very helpful. However this is not mentioned in the text.

Thank you for the comment. We have now added a reference to the Supplement in the main text, specifically in Section 2.1.3. — "The Supplement provides further details on the relevant distributions."

4) Kernel size is a parameter investigated through the paper, but it is not defined explicitly. Is it the sigma in equation 5, labelled below "the spatial extent"? Refer to the parameter consistently as "kernel size" for clarity.

You are right that the kernel size refers to the σ in Eq. 5. In the previous version of the manuscript, the term "kernel size" was defined at the beginning of Section 3.2 of Results. We agree that it would be clearer to define it explicitly and to refer to it consistently as "kernel size". Thus, we have now revised the manuscript accordingly. We have added a clear definition of "kernel size" right after Eq. 5 in Section 2.2.1. — "... where σ defines the spatial extent (hereafter intuitively referred to as the kernel size)."

5) You evaluate the models with and without GKF, but don't address explicitly how the locations and emissions rate is calculated without GKF.

Thank you for the comment. You are right that we evaluated the models with and without Gaussian kernel fitting (GKF). In the previous version of the manuscript, we described how

the locations and emission rates are calculated without GKF in L285-286 in Section 3.2: "... the sources are extracted from the local peaks ..." The "local peaks" is defined as the local maximum in the heatmap. The method of using local peaks to represent the detected objects is also named as "non-maximum suppression (NMS)" in the field of object detection. We have now revised the manuscript to further clarify this point. Specifically, we have added a clear description why we perform GKF in Section 2. — "To extract point objects from predicted heatmaps, traditional non-maximal suppression (NMS) uses the local peak pixels to represent the object (Law and Deng, 2019), which limits the accuracy in subpixel localization. In this regard, we ..."

6) typos and minor corrections: line 161 - "Each mean $e_{ps,i}$ is sampled [...]" I assume you mean each individual $e_{ps,i}$ is sampled from the inventory, and the mean calculated from this sample

Thank you for the careful reading. You are right that in "Each $\bar{e}_{ps,i}$ is sampled from ...", $\bar{e}_{ps,i}$ should be $e_{ps,i}$. We would also like to clarify that $\bar{e}_{ps,i}$ denotes the reference emission rate, i.e., the mean emission rate in the training dataset, for point source i . To improve clarity, we have now revised the notation $\bar{e}_{ps,i}$ by $e_{ps,i}^{\text{ref}}$, and \bar{e}_{anth} by e_{anth}^{ref} .

line 314 - remove "on the generalization"

line 322 — SMARTCARB instead of SMARTCARAB

Thank you for pointing this out. We've corrected these points accordingly.

List of Changes

We have made the following changes in the revised manuscript:

1. We clarified the definition of "kernel size" in Section 2.2.2.
2. We clarified the dataset split strategy in Section 2.2.3.
3. We clarified the reasons for selecting only 8 OCO-3 observations in Section 2.3.3.
4. We revised the typos in Fig. 5 and adjusted the x-axis labels in Fig. 6 for better readability.
5. We clarified the motivation for performing Gaussian kernel fitting (GKF) in Section 2.2.2 and further clarified the method of extracting source locations and emission rates without GKF in Section 3.2.
6. We extended the discussion on the potentially overestimated performance on the synthetic dataset in Discussions and Conclusion.

7. We revised the notation in Eq. 3 and Eq.7 for better clarity. The corresponding notations in the Supplement have also been revised accordingly.
8. We corrected the typos, grammatical errors and other technical issues throughout the manuscript, including but not limited to the points raised by the reviewers.

References

- H. Law and J. Deng. CornerNet: Detecting Objects as Paired Keypoints, Mar. 2019.
- X. Lin, R. van der A, J. de Laat, H. Eskes, F. Chevallier, P. Ciais, Z. Deng, Y. Geng, X. Song, X. Ni, D. Huo, X. Dou, and Z. Liu. Monitoring and quantifying CO₂ emissions of isolated power plants from space. *Atmospheric Chemistry and Physics*, 23(11):6599–6611, June 2023. ISSN 1680-7316. doi: 10.5194/acp-23-6599-2023.
- OCO-2/OCO-3 Science Team, A. Chatterjee, and V. Payne. OCO-3 Level 2 bias-corrected XCO₂ and other select fields from the full-physics retrieval aggregated as daily files, Retrospective processing v10.4r, 2022.