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 Hu[1], Longfei Tian[1], Shuang Gao[1], and Guohua Liu[1,2]

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

[2]University of Chinese Academy of Sciences, Beijing, China

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

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

December 9, 2025

We sincerely thank the editor and reviewers for the valuable comments and suggestions. We have carefully considered all the comments and revised the manuscript accordingly. As requested by the topic editor, we've added the corresponding manuscript corrections in this response.

Before presenting our point-by-point responses to the reviewers' comments, we would like to report a mistake in the original manuscript. The spatial resolution of OCO-3 is incrorrectly set as "2.25×0.7 km^2", which should be "2.2×1.6 km^2" according to Eldering et al. [2019]. The reason for this mistake is that we transcribed the spatial resolution from ESA's eoPortal `https://www.eoportal.org/satellite-missions/iss-oco-3`, which is not accurate, at the early stage of this work. This mistake does not affect the main conclusions of the manuscript. As we interpolated the OCO-3 data to a uniform grid, the slightly different grid size does not affect the spatial distribution pattern of the XCO_2, or make significant differences in the model predictions. We have corrected this mistake in the revised manuscript and re-performed the experiments related to OCO-3. The performance of GHGPSE-Net on the OCO-3

datasets only slightly changed. The main results and conclusions hold unchanged.

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 manuscript "GHGPSE-Net: A method towards spaceborne automated extraction of greenhouse-gas point sources using point-object-detection deep neural network" presents a novel deep learning framework and a large dataset for detecting and quantifying greenhouse gas (GHG) point sources from satellite imagery.

We appreciate your positive feedback on our manuscript. We are glad that you found our deep learning framework and dataset to be novel and valuable.

The authors are the first to introduce the point object detection approach in this domain, which is very valuable and insightful for the GHG point source monitoring community, as it has the potential to integrate and simplify the processing complexity largely. The authors also demonstrate the feasibility of the model by evaluating it on two datasets, including authentic satellite observations.

That is a very precise summary of our work. We are grateful for your recognition of our contribution to GHG point source monitoring.

Though I anticipate more evaluations may be required on the upcoming moderate-resolution carbon monitoring satellites (e.g., CO2M and TanSat-2) to fully explore the potential. This work marks an important step towards automated GHG point source monitoring and has the potential to make a significant contribution to the GHG remote sensing community.

Thank you for your valuable suggestion. We agree that further evaluations on real satellite data are required. Currently, there are limited available observations from such satellites with moderate resolution specifically designed for GHG monitoring. We are currently testing our methods on more high-resolution satellite data, such as EMIT, and we plan to evaluate our model on the upcoming missions. However, in this study, we focused on developing and validating the methodology using synthetic data and available satellite observations to demonstrate its feasibility. Further evaluations may be out of the scope of this initial study. Future work will continue to fully explore the potential of our approach on both moderate-resolution and high-resolution satellite data.

Minor suggestions:
(1-1) The dataset construction process, including WRF-GHG simulation, XCO2 construction and data augmentation, involves multiple scenarios, especially it seems that the model is

trained on the synthetic dataset and evaluated using independent datasets. It may be better clarified using a diagram.

We appreciate your suggestion. We agree that the current description of the dataset construction process is rather complex. Here is a summary of the dataset construction process:

- **WRF-GHG Simulation:** We use the WRF-GHG model to simulate greenhouse gas plumes.
- **Pseudo-observation:** We generate pseudo-observations of XCO2 for each instrument scenarios (HGET, HGET-3ppm, UCPI, etc.) based on the WRF-GHG outputs.
- **Data augmentation:** Based on the pseudo-observations, we apply various data augmentation techniques (e.g., rotation, scaling, noise addition) to enhance the diversity of the dataset. The data augmentation is performed separately for each instrument scenario. Although the augmentation processes can be applied in the training process, we perform them beforehand for simplicity and efficiency. There are 3,144 snapshots generated from WRF-GHG simulations and there are 24,000 samples for each instrument scenario after augmentation. If not specifically mentioned, the data augmentation is applied to all the datasets.

We've added a new figure to Section 2.1 to illustrate the whole process of dataset construction with more clarity:

> "The overall workflow for dataset construction is illustrated in Fig.1. (1) WRF-GHG is used to simulate the three-dimensional CO_2 concentration for multiple tracers over Shanghai from January to April 2020, with hourly snapshots (Section 2.1.1). (2) The column-averaged dry air mole fraction, denoted as XCO_2, is derived from the three-dimensional CO_2 concentration, and the pseudo-observation images are then synthesized for different instrument configurations (Section 2.1.2). (3) A data augmentation method is proposed to improve the generalization of the model trained on the spatial-temporal limited dataset (Section 2.1.3). "

(1-2) The authors summarized GHGPSE-Net in Figure 3. However, the overall methodology, including simulation, simulation evaluation, training dataset preparation, and deep learning evaluation, is quite complex and somewhat difficult to follow. The authors may consider summarizing the entire methodology in Section 2.

We appreciate the detailed and constructive suggestion. We will consider better summarizing the entire methodology in Section 2. The construction of the datasets is described earlier in our response to comment (1-1). Here, we summarize the evaluation process as follows:

Experiment	Trainning	Evaluation	GKF?	σ [km]
Sect. 3.2 (Fig. 6)	HGET	HGET	•	0.25
	HGET	HGET	•	0.50
	HGET	HGET	•	0.75
	HGET	HGET	•	1.50
	HGET	HGET		0.25
	HGET	HGET		0.50
	HGET	HGET		0.75
	HGET	HGET		1.50
Sect. 3.2 (Fig. 7)	HGET	HGET	•	0.25
	HGET	HGET	•	0.50
	HGET	HGET	•	0.75
	HGET	HGET	•	1.50
	HGET	HGET	•	3.00
	HGET	HGET		0.25
	HGET	HGET		0.50
	HGET	HGET		0.75
	HGET	HGET		1.50
	HGET	HGET		3.00
Sect. 3.3 (Tab. 3)	HGET	HGET	•	0.50
	HGET-3ppm	HGET-3ppm	•	0.50
	UCPI	UCPI	•	2.00
	UCPI-3ppm	UCPI-3ppm	•	2.00
	OCO-3[1]	OCO-3	•	2.20
Sect. 3.4 (paragraph 1)	HGET	HGET	•	0.50
	HGET (no-aug)	HGET	•	0.50
Sect. 3.4 (SMARTCARB)	UCPI-0.7ppm	SMARTCARB[2]	•	2.00
Sect. 3.4 (OCO-3)	OCO-3	OCO-3 (obs)[3]	•	2.20

[1] It's worth noting that the OCO-3 dataset is constructed based on synthetic data.
[2] The SMARTCARB dataset is constructed based on Kuhlmann et al. [2019b].
[3] The OCO-3 (obs) dataset is constructed based on real OCO-3 satellite observations [OCO-2/OCO-3 Science Team et al., 2022].

Table 1: Summary of the experiment configurations conducted in the "Results" section.

Figure 1: Workflow for generating the synthetic XCO_2 observation dataset.

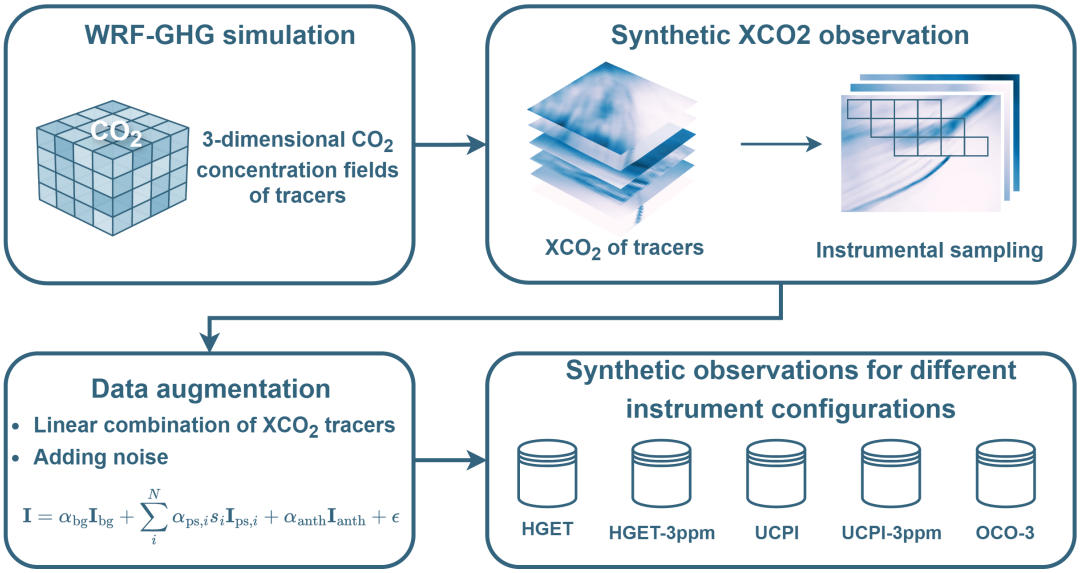

(1-3) In some related GHG plume detection studies, deep learning models usually require wind as an input. Does GHGPSE-Net not require the 2D wind field as input?

Currently, GHGPSE-Net does not require the 2D wind field as input and performs source extraction solely using XCO2 observations. Indeed, additional input to the model, such as wind field, NO_2 concentration, will enhance the model performance, as demonstrated by Kuhlmann et al. [2019a], Kuhlmann et al. [2020] and Dumont Le Brazidec et al. [2024]. Additional input, such as cloud cover, surface type and land-ocean mask, may be helpful to perform quality control. However, the inclusion of additional input will increase the model complexity and make the analysis more difficult. In this initial study, we focus on developing a basic framework for GHG point source extraction using only XCO2 observations. Future work will explore the inclusion of additional inputs and improve its real-world applications. We added the following text to the "Discussions and Conclusion" section to deliver this point:

> "Future work could . . . investigate robustness to real-world discontinuities such as cloud cover and water bodies; and explore integrating auxiliary information, such as NO_2 observations [Dumont Le Brazidec et al., 2024] to help alleviate the low contrast issue in CO_2 plumes. "

(1-4) According to the result (e.g., Table 3), it seems the "2 km × 2 km" in L10 should be 0.5 km × 0.5km.

Thank you for pointing out this inconsistency. You are correct that the case referred in the abstract should be HGET, which is 0.5 km × 0.5km. We've revised this mistake in the abstract (L10).

Technical comments:

(2-1) Typo in L59 and L137.

Thank you for your careful review. It should be “further” in L59 and L137 should be “... the concentration map can be expressed as: ...”.

(2-2) It should be "mean squared error" instead of "mean square error" in L10, L208, and L223.

Thank you for your suggestion and we’ve corrected this mistake in L208. We have corrected this mistake at L208 (now L216 in the revised version). Regarding the instances at L10 and L223 where “mean squared error” is mentioned, the metric used and described in the manuscript is the “root mean square error (RMSE)”; therefore, these parts have been kept unchanged.

Point-by-point Response to Reviewer #2

This work applies deep learning to satellite snapshots of CO2, performing point source detection, localisation and quantification. As the authors identify, most work in this area focuses on plume identification or source location, without performing quantification directly.

We appreciate your recognition of our work. Indeed, most existing studies focus on plume detection and the source localization are often performed mannully based on the detected plumes. The quantification process and plume detection are often treated as separate processes. This issue becomes significant when considering large volume of moderate satellite observations, where each pixel covers a broad area. Such data requires both precise localization and efficient processing.

The authors present an innovative solution, where a CNN is trained to identify sources as hotspots, which then get detected and quantified using a statistical Gaussian Kernel fitting method.

We thank the reviewer for the summary of our work. You pointed out the key components of our approach, including the CNN-based hotspot identification and the Gaussian Kernel fitting for quantification. Though, it is worth noting that the Gaussian kernel fitting is based on least squares fitting rather than statistical fitting. This offers better precision than traditional Non-Maximum Suppression (NMS) method, as demonstrate in Fig. 6. We estimate the source parameters by minimizing the residuals between the UNet-predicted heatmap and the Gaussian kernel model.

The choice of Gaussian kernel is rather empirical and is selected for several reasons. (1) To represent point objects in a CNN-generated heatmap, a smooth 2D function is necessary. Directly using delta functions, such as one-hot binary masks, results in sparse gradients and

poor convergence. (2) The Gaussian kernel function has a simple mathematical form and is used by previous studies to represent object size and location [Law and Deng, 2019]. (3) The form of Gaussian kernel is aligned with the point source parameters, where the center and peak value have physical meanings.

The authors also introduce a greenhouse-specific augmentation method, varying the linear coefficients for each source, that presents a valuable contribution to the area of GHG+deep learning.

We appreciate the reviewer's commnets. It is worth noting that the augmentation method is directly inspired by that used in Dumont Le Brazidec et al. [2024]. Here, we generalize the problem in a linear combination and write it into a clear mathematical form. Furthermore, we also use statistical analysis to decide the linear coefficient ranges, instead of empirically defined. We suppose this form of GHG-specific data augmentation method can be intuitive and useful for future studies. For example, as we mentioned in the manuscript, the noise term can be more complex than naive Gaussian noise. Future studies may consider adding instrument-specific/scene-specific noise patterns into the augmentation process.

Clarifications required:
The datasets used should be summarised more cohesively, as at the moment it is hard to follow 1) what regions, areas and time periods are covered for each dataset 2) the number of samples in each dataset 3) whether the test datasets were augmented. The number of samples added in the augmentation needs to be specified.

Thank you for your constructive suggestions. We agree that a clear summary of the dataset generation process is important for better clarification, as the entire generation process is rather complicated. We would answer your questions as follows:

(1) The $X\mathrm{CO}_2$ data before augmentation is generated by WRF-GHG simulations. We generated a $X\mathrm{CO}_2$ snapshot for each tracer channel every hour. Based on the synthetic $X\mathrm{CO}_2$ snapshots, we further generated instrument-specific datasets (HGET, HGET-3ppm, UCPI, etc.). So, in a nutshell, the regions and time periods covered by each dataset are the same as those of the WRF-GHG simulations, which are described in detail in Section 2.1.1. Further details, including source codes, configuration files, initial conditions, boundary conditions and emissions files, are provided in a Zenodo repository (`https://zenodo.org/records/17337441`).

(2) The model is still running. For this work, we use snapshots spanning from 2020-01-01 to 2020-04-30, recorded at 1-hour intervals. For each instrument scenario, we generated 24,000 samples using the method described in Section 2.1.2, with data augmentation (Section 2.1.3) from the original snapshots. Each dataset is divided into training, validation, and testing sets in a 3:1:1 ratio (Section 2.2.3).

(3) The data augmentation is performed before the construction of datasets, due to efficiency consideration and simplicity. Therefore, the test datasets are also augmented. Actually, all the datasets used in this study are generated with data augmentation, except for the control group's training stage at the beginning of Section "Generalization evaluation using SMART-CARB dataset and OCO-3 observation".

We've summarized the data construction process (please see the following responses) and added the following text to clarify the simulation time periods coverage (Section 2.1.1):

> "The model runs continuously from January 1, 2020, to April 30, 2020, and produces snapshots at 1-hour intervals . . . "

The architecture of the model needs to be outlined more clearly, including how it was tuned, the learning rate, loss functions etc.

Thank you for your comments. The general architecture of the model is described in Section 2.2.1 and the training setups are described in Section 2.2.3. We will further improve the description of the model training process. Specifically, we used the Adam optimizer, with an initial learning rate of 0.001, which is a classical choice in deep learning tasks. The loss function is the Mean Squared Error (MSE) between the predicted heatmap and the ground truth heatmap, as described in Section 2.2.3. Further details of the model can be referred to a Zenodo repository (`https://zenodo.org/records/16751293`). We added the following text to Section 2.2.3:

> ". . . the model weights are updated by minimizing the mean squared error (MSE) between the CNN-predicted and true heatmaps. The Adam optimizer is used for training with an initial learning rate of 0.001 and a batch size of 16. The models generally converge whitin 30 epochs."

It is not clear to me if the resolution of the model was kept constant across the different instruments, or adapted to each instrument.

Thank you for your question. As we described in Section 2.1.2, the resolution of the XCO_2 data follows the instrument pixel size and image size. We currently focus on building a single-instrument-specific model, instead of a general model that can adapt to different instruments. For each instrument scenario, the models are trained the evaluated using different datasets. We suppose a further clarification of the dataset generation process will help to understand this point. We added a new figure and better summarized the workflow of dataset construction in Section 2.1:

> "The overall workflow for dataset construction is illustrated in Fig.1. (1) WRF-GHG is used to simulate the three-dimensional CO_2 concentration for multiple

tracers over Shanghai from January to April 2020, with hourly snapshots (Section 2.1.1). (2) The column-averaged dry air mole fraction, denoted as XCO_2, is derived from the three-dimensional CO_2 concentration, and the pseudo-observation images are then synthesized for different instrument configurations (Section 2.1.2). (3) A data augmentation method is proposed to improve the generalization of the model trained on the spatial-temporal limited dataset (Section 2.1.3). ”

We also described the application of the instrument effect with more clarity in Section 2.1.2:

> “The downsampling process simulates the ground pixel size of different instruments . . . Here, we consider five different instrument configurations, including low-retrieval-noise (1.5 ppm) and high-retrieval-noise (3 ppm) scenarios for both HGET and UCPI of TanSat-2, as well as the OCO-3 instrument. The instrument configurations, including ground pixel size, image size and noise level, are shown in Table 2.”

Is it possible that the significantly lower skill on the SMARTCARB dataset is due to a different transport model?

That’s an insightful point. The SMARTCARB dataset is generated using the COSMO-GHG model (Kuhlmann et al. [2019b]), which is different from the WRF-GHG model used in our training dataset. The different transport models may lead to different plume shapes and distributions as discussed by Brunner et al. [2023]. We treat the emission sources as ground sources, and the plumes are mostly confined near the ground. This is slightly different from the elevated sources used in the SMARTCARB dataset [Brunner et al., 2019, Kuhlmann et al., 2019b]. Our model may underestimate the plumes entering the higher altitudes, as we discussed in the “Discussions and Conclusion” section. We also found a false positive case where a plume structure that may have escaped the boundary layer was identified as an independent plume (shown in Fig. S6). Similarly, other biases among different transport models may also contribute to the performance degradation. However, further analysis is difficult and seems beyond the scope of this study. Further investigations may be needed to fully understand these biases, and they are currently less discussed in the GHG remote sensing community. We complemented the following discussions to the “Discussions and Conclusion” section:

> “Other factors, such as differences in meteorological conditions, terrain effects, and the use of different transport models (COSMO-GHG v.s. WRF-GHG), may also contribute to the discrepancies.”

We also added more discussions on other challenges reflected by the evaluation experiments to the “Discussions and Conclusion” section:

"The complex and heterogeneous noise characteristics, diverse plume morphologies, and varying meteorological conditions in real observations also pose challenges to model generalization."

Would other inputs to the CNN, like the wind direction or the location of known sources, improve predictive skill?

Thank you for your suggestions. This point is also raised by Reviewer #1, and thus we share the same response here. Currently, GHGPSE-Net does not require the 2D wind field as input and performs source extraction solely using XCO2 observations. Indeed, additional input to the model, such as wind field, NO_2 concentration, will enhance the model performance, as demonstrated by Kuhlmann et al. [2020] and Dumont Le Brazidec et al. [2024]. Additional input, such as cloud cover, surface type and land-ocean mask, may be helpful to perform quality control. However, the inclusion of additional input will increase the model complexity and make the analysis more difficult. In this initial study, we focus on developing a basic framework for GHG point source extraction using only XCO2 observations. Future work will explore the inclusion of additional inputs and improve its real-world applications. We added the following text to better clarify this point in the "Discussions and Conclusion" section:

"Future work could ... investigate robustness to real-world discontinuities such as cloud cover and water bodies; and explore integrating auxiliary information, such as NO_2 observations [Dumont Le Brazidec et al., 2024] to help alleviate the low contrast issue in CO_2 plumes. "

The Gaussian kernel fitting is demonstrated to improve when the kernel size is the same as the instrument pixel size, for the one instrument tested. Do you expect performance to improve for other instruments when the same process is applied? It is not clear what kernel size you applied for evaluation for each instrument.

That's an accurate understanding. The kernel size achieves optimal performance in a moderate size. As the kernel that is too wide tends to blur local features, making the kernel fitting less accurate. In contrast, a kernel that is too narrow tends to bring in overly localized gradients, hindering parameter updates in the training. The initial experiment using HGET shows that the optimal kernel size is close to the instrument pixel size. We suppose this conclusion can be generalized to other instruments, as similar forms and reasons also apply to other instruments. Moreover, it is worth noting that the optimal kernel size may not exactly match the instrument pixel size, as other factors such as instrument noise, scene complexity may also affect the optimal kernel size. The complexity of the deep learning model further makes theoretical analysis difficult. However, as the core goal of this work is to demonstrate the feasibility of using a point-detection neural network for GHG point source extraction, further large-scale controlled

experiments for the choice of kernel size can be rather petty and make the literature redundant.

Regarding the kernel size (σ value) used for other instruments, we used a kernel size of 0.5 km for the baseline HGET scenario and HGET-3ppm scenario. For the UCPI and UCPI-3ppm scenarios, we used a kernel size of 2 km. For the OCO-3 scenario, we used a kernel size of 2.2 km. A controlled experiment using varying kernel sizes is only applied to the HGET scenario.

We added the following text to Section 3.2 for better clarity:

> "The kernel size of GKF is set to match each instrument's spatial resolution, i.e., 500 m for HGET, 2 km for UCPI, and 2.2 km for OCO-3, in the following text."

Typos and minor corrections:

Figure 5 shows, from the text, two days with large absolute errors. Clarify in the caption/figure labelling that these two days are not representative of the usual error distribution.

Thank you for your constructive suggestion. We will add it to the caption in the next revision. It is worth noting that these errors are mostly observed downwind of the city center of Shanghai and heavy industrial areas near the Yangtze River. We've attached a spatial map of the errors in Fig. S13, which may help to understand the error distribution. We added the following text to the caption of Figure 6 (formerly Figure 5):

> "The large deviations in Feb 20, 2020 and Apr 23, 2020 are mostly observed downwind of urban or heavy-industry sources (Supplement)."

Line 39 ("and The new generation…")

Thank you for pointing out this typo. We've corrected it in the revised manuscript (L39):

> "…, and the new generation …"

Line 151: areas (not area)

Thank you for pointing out this typo. We've corrected this point in the revised manuscript (L160).

L342

Thank you for pointing out this typo. "The" should be "the", and we've corrected it (L358).

Supplement S3.1, title: Definition (not defination)

Thank you for pointing out this typo. We've corrected this point in the revised Supplement.

List of Changes

We have made the following changes in the revised manuscript:

1. We improved the description of dataset construction with more clarity, including a new figure summarizing the workflow and further clarification of the simulation time, procedures, and instrumental configurations.

2. We improved the clarity of the experiment setups.

3. We further described the outliers in Figure 6 (Comparison between XCO2 simulated by WRF-GHG and that retrieved from OCO-3 observations).

4. We clarified the kernel size setting in Section 3.3.

5. We corrected the OCO-3 instrument settings and re-performed the experiments. The new results are updated in Table 3, Section 3.4 and Supplement. We also moved the figure of the OCO-3 observations and predicted heatmaps from the supplement into the main text and provided additional explanations.

6. We expanded the discussion to better explain the performance deterioration in SMART-CARB simulation, current limitations, future work, and other related aspects.

7. We modified the "Code and data availability" section as required by the editor-in-chief to meet the journal's polices.

8. We made additional minor modifications to improve clarity and conciseness.

9. We corrected typos and other textual errors.

References

Dominik Brunner, Gerrit Kuhlmann, Julia Marshall, Valentin Clément, Oliver Fuhrer, Grégoire Broquet, Armin Löscher, and Yasjka Meijer. Accounting for the vertical distribution of emissions in atmospheric CO_2 simulations. *Atmospheric Chemistry and Physics*, 19(7): 4541–4559, April 2019. ISSN 1680-7316. doi: 10.5194/acp-19-4541-2019.

Dominik Brunner, Gerrit Kuhlmann, Stephan Henne, Erik Koene, Bastian Kern, Sebastian Wolff, Christiane Voigt, Patrick Jöckel, Christoph Kiemle, Anke Roiger, Alina Fiehn, Sven Krautwurst, Konstantin Gerilowski, Heinrich Bovensmann, Jakob Borchardt, Michal Galkowski, Christoph Gerbig, Julia Marshall, Andrzej Klonecki, Pascal Prunet, Robert Hanfland, Margit Pattantyús-Ábrahám, Andrzej Wyszogrodzki, and Andreas Fix. Evaluation of simulated CO_2 power plant plumes from six high-resolution atmospheric transport models.

Atmospheric Chemistry and Physics, 23(4):2699–2728, February 2023. ISSN 1680-7316. doi: 10.5194/acp-23-2699-2023.

Joffrey Dumont Le Brazidec, Pierre Vanderbecken, Alban Farchi, Grégoire Broquet, Gerrit Kuhlmann, and Marc Bocquet. Deep learning applied to CO_2 power plant emissions quantification using simulated satellite images. *Geoscientific Model Development*, 17(5):1995–2014, March 2024. ISSN 1991-959X. doi: 10.5194/gmd-17-1995-2024.

Annmarie Eldering, Thomas E. Taylor, Christopher W. O'Dell, and Ryan Pavlick. The OCO-3 mission: Measurement objectives and expected performance based on 1 year of simulated data. *Atmospheric Measurement Techniques*, 12(4):2341–2370, April 2019. ISSN 1867-1381. doi: 10.5194/amt-12-2341-2019.

Gerrit Kuhlmann, Grégoire Broquet, Julia Marshall, Valentin Clément, Armin Löscher, Yasjka Meijer, and Dominik Brunner. Detectability of CO_2 emission plumes of cities and power plants with the Copernicus Anthropogenic CO_2 Monitoring (CO2M) mission. *Atmospheric Measurement Techniques*, 12(12):6695–6719, December 2019a. ISSN 1867-1381. doi: 10.5194/amt-12-6695-2019.

Gerrit Kuhlmann, Valentin Clément, Julia Marshall, Oliver Fuhrer, Grégoire Broquet, Christina Schnadt-Poberaj, Armin Löscher, Yasjka Meijer, and Dominik Brunner. SMARTCARB – Use of satellite measurements of auxiliary reactive trace gases for fossil fuel carbon dioxide emission estimation. Technical report, Zenodo, January 2019b.

Gerrit Kuhlmann, Dominik Brunner, Gregoire Broquet, and Yasjka Meijer. Quantifying CO2 emissions of a city with the Copernicus Anthropogenic CO2 Monitoring satellite mission. *Atmospheric Measurement Techniques*, 13(12):6733–6754, December 2020. ISSN 1867-1381. doi: 10.5194/amt-13-6733-2020.

Hei Law and Jia Deng. CornerNet: Detecting Objects as Paired Keypoints, March 2019.

OCO-2/OCO-3 Science Team, Abhishek Chatterjee, and Vivienne 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.