ADECEES: Anomaly DEtection of CO_2 Emissions via Ensemble Segmentation

Andrianirina Rakotoharisoa
DataLearning Group, Department of Earth Science & Engineering
Imperial College London, London, UK
andrianirina.rakotoharisoa19@imperial.ac.uk

Simone Cenci
Institute for Sustainable Resources
University College London, London, UK

Rossella Arcucci
DataLearning Group, Department of Earth Science & Engineering, London, UK
r.arcucci@imperial.ac.uk

Abstract

Latest studies show that we are not on track to limit global warming below 1.5°C compared to pre-industrial levels. Reaching Net Zero is an essential target to reduce global warming and requires accurate and global monitoring of global emissions. In this paper, we introduce our Anomaly DEtection of CO2 Emissions via Ensemble Segmentation (ADECEES) system for the identification of consequences of CO_2 emissions on the atmosphere relying on partial diffusion and ensemble segmentation. We apply our system on a global XCO_2 dataset and illustrate that it can be used both for the detection of point sources and the detection of variation of emissions.

1 Introduction

There are growing concerns whether we will be able to keep global warming below 1.5° compared to pre-industrial levels with current policies [1, 2]. Countries who ratified the Paris Agreement [3] submitted Nationally Determined Contributions (NDCs) [4] to commit to reduce their emissions and implement adaptive measures. Anthropogenic emissions of greenhouse gases (GHGs) indeed play a significant role in global warming and CO_2 specifically is responsible for a 0.8°C (0.5°C–1.2°C) historical increase [5]. To track emissions and make sure they are reduced effectively, countries have implemented Mandatory Reporting Programs [6, 7]. However, these reporting programs present two issues: they are inconsistent between countries, as there is no universal protocol [8], and inconsistent between companies as the level of disclosure is voluntary [9]. An investigated alternative is to directly monitor CO_2 emissions from space using remote sensing data [10, 11]. This solution suffers from limitations as multiple satellite overpasses are necessary to constrain emissions accurately [12]. More recently, works take advantage of deep learning to infer and predict emissions based on other variables such as vapor plumes or energy production [13, 14]. While most of these studies focus on the quantification of CO_2 emissions from known sources, the topic of CO_2 point source detection and discovery remains relatively unexplored. In this paper, we present a multi-purpose CO_2 emissions monitoring system, ADECEES (Anomaly DEtection of CO2 Emissions via Ensemble Segmentation) for the detection of point sources and the study of change in emissions. It relies on a daily global XCO_2 (column-average dry-air mole fraction of atmospheric CO2 [15]) dataset to perform anomaly

Tackling Climate Change with Machine Learning: workshop at NeurIPS 2025.

detection using partial diffusion [16] and segment high CO_2 concentration zones attributable to point sources. The paper is organised as follows: Section 2 presents the system, Section 3 the dataset used for anomaly detection along the preprocessing steps and Section 4 regroups the tests conducted on our system. Finally, our findings are summarized in Section 5.

2 Anomaly detection system

Point Source Detection with Partial Diffusion Denoising Diffusion Probabilistic Models (DDPM) are a class of latent variable models that consist of two phases and can be used to produce samples from a target distribution $q(x_0)$. Following the work of Wyatt et al. [17] on AnoDDPM ("Anomaly Detection with Denoising Diffusion Probabilistic Models using Simplex Noise") and recent progress made in diffusion-based anomaly detection [18], we adapt the standard DDPM framework to our task by employing a *partial diffusion* process driven by *simplex noise*, rather than the full Markov chain of Gaussian steps. Compared to the standard DDPM framework, our partial simplex diffusion approach leverages multi-scale simplex perturbations that better preserve structural information and enhance sensitivity to subtle anomalies. Specifically, given a healthy reference image $x_0 \in \mathbf{D}_{healthy}$, we add simplex noise via a truncated forward diffusion pass to obtain a noisy sample x_t using:

$$q_{\text{simplex}}(x_t \mid x_0) = \text{SimplexNoiseScale}(t) \cdot \varepsilon_{\text{simplex}}, \quad \varepsilon_{\text{simplex}} \sim \text{SimplexNoise}, \tag{1}$$

where $\text{SimplexNoiseScale}(t)$ controls anomaly-size sensitivity. The model is then trained to denoise this partially simplex-noised input and reconstruct the original x_0 via

$$\hat{x}_0 = f_\theta(x_t, t), \tag{2}$$

by minimizing the reconstruction loss

$$\mathcal{L}_{\text{simplex}} = \mathbb{E}_{x_0 \sim \mathbf{D}_{healthy},\, t}\left[\|f_\theta(x_t, t) - x_0\|^2\right]. \tag{3}$$

This *partial simplex diffusion* strategy allows the model to scale to high-resolution imagery and provides control over anomaly granularity, as shown in Wyatt et al. [17]. During this phase, we feed the model CO_2 maps without suspected unreported point sources or variations, ensuring that the denoising process learns a faithful distribution of healthy references against which anomalies can be detected. A more detailed explanation and illustration of the diffusion process, along with the training procedure can be found in Annex A. In our context, each location possesses an iteration of the model with specific weights as XCO_2 profiles vastly differ by region, making it impractical to train a single global model.

Inference During inference, we feed the model CO_2 profiles with suspected emission or concentration anomalies x_0. These maps are diffused before being reconstructed in the reverse process into healthy samples $\hat{x}_0$. As the reconstructed maps follow the same data distribution as the training dataset, we note $\hat{x}_0 \sim \mathbf{D}_{healthy}$. Furthermore, generating only one output introduces excessive randomness into the anomaly detection process. To increase the robustness and minimize the uncertainty in our reconstructions, we generate an ensemble of n=50 reconstructions which we then average and segment. Once we have generated the reconstructed map $\hat{x}_0$, we perform a difference x_{diff} with the real XCO_2 map x_0. A segmented array $x_{seg.}$, with the same shape as x_0, is then computed. If the values of x_{diff} are higher than a threshold t, they are then segmented as high anomalies of XCO_2.

3 Dataset

XCO_2 global dataset We employed a global daily XCO_2 dataset for CO_2 monitoring with a $0.03° \times 0.04°$(<5km) spatial resolution from 1 January 2015 until 28 February 2022 generated by Rakotoharisoa et al. [19]. This dataset results from the downscaling of a L3 dataset from the OCO-2 mission from NASA [20]. Samples of the dataset over the span of four years are presented in the Annex C for visualisation.

Preprocessing Some preprocessing steps are needed in order to generate the input data for our system. First, from the daily global maps $x_{1\to7}^{global}$, we extract a 256x256 array $x_{1\to7}^{local}$ centered on the point source we wish to investigate. We then normalize each daily subarray between -1 and 1 before generating a weekly average x_{weekly}^{local}. The averaging of our daily local maps is performed

in order to limit the impact of weekly and exceptional events on the predictive performance of our model [21, 22]. Examples of significant changes are presented in Annex B. The whole preprocessing pipeline is represented in Figure 1.

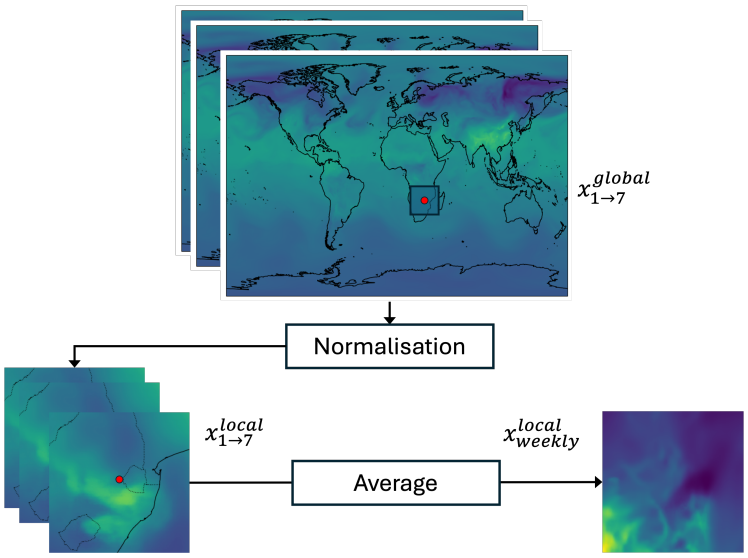

Figure 1: Preprocessing steps of our input data. The red dots represent the location of the point source.

4 Test cases

In this section, we present two applications highlighting the potential of our system for CO_2 emissions monitoring: firstly the identification of point sources and secondly the detection of changes in CO_2 profiles resulting from variations in CO_2 emissions.

Point source detection To assess whether ADECEES is able to detect emission sources, we select locations with coal mines from the Global Coal Mine Tracker Dataset [1] where coal mines are not continuously active between 2015 and early 2022 (see Annex D). with the objective to see if the model is able to detect their activity. We evaluate the accuracy, precision and recall of the system in these four locations with a 50/50 ratio of operational and inactive period and a varying threshold t. High accuracy indicates that the model can effectively differentiate profiles according to the presence of an emission source.

The metrics of the system's performance are reported in Table 1. Our first observation is that its performance depends on the location. The Invierno mine exhibits poorer metrics compared to other locations, with accuracy and precision values never exceeding 0.5. This suggests that the system struggles to distinguish between CO_2 profiles corresponding to an operational and inactive point source, even with the tuning of the threshold. Conversely, except for the Invierno mine, we can observe that both the accuracy and F1-score of the system tend to improve as the threshold increases. This trend is particularly evident for the Otvialny and Hazelwood mines. For these facilities, high accuracy and F1-score (>70%) indicate the system's ability to reliably detect the presence of a point source while keeping false positives low.

Variation of emissions Secondly, we evaluate our system's capability to detect emission variations from known power plants. Emissions for the first half of 2022 have been quantified [11] and the Global Energy Monitor (GEM) [2] reports events that may correspond to the variation we aim to visualise. The power plants we investigate are described in Annex D.

[1] Global Coal Mine Tracker, Global Energy Monitor, May 2025 release.

[2] `https://globalenergymonitor.org` Accessed on 20/05/2025.

Table 1: Detection rate of various point sources using ADECEES depending on the threshold t. Best for each site in **bold**.

	t=0.3				t=0.5				t=0.7			
	Bg.	Ot.	Iv.	Hw.	Bg.	Ot.	Iv.	Hw.	Bg.	Ot.	Iv.	Hw.
Acc.	0.50	0.60	**0.50**	0.50	**0.60**	0.80	**0.50**	0.50	**0.60**	**0.90**	0.40	**0.70**
Prec.	0.50	0.56	**0.50**	0.50	0.57	0.71	**0.50**	0.50	**0.60**	**1.00**	0.43	**0.67**
Recall	**1.00**	**1.00**	**1.00**	**1.00**	0.80	**1.00**	**1.00**	**1.00**	0.60	0.80	0.60	0.80
F1-Sc.	**0.67**	0.71	**0.67**	0.67	**0.67**	0.83	**0.67**	0.67	0.60	**0.89**	0.50	**0.73**

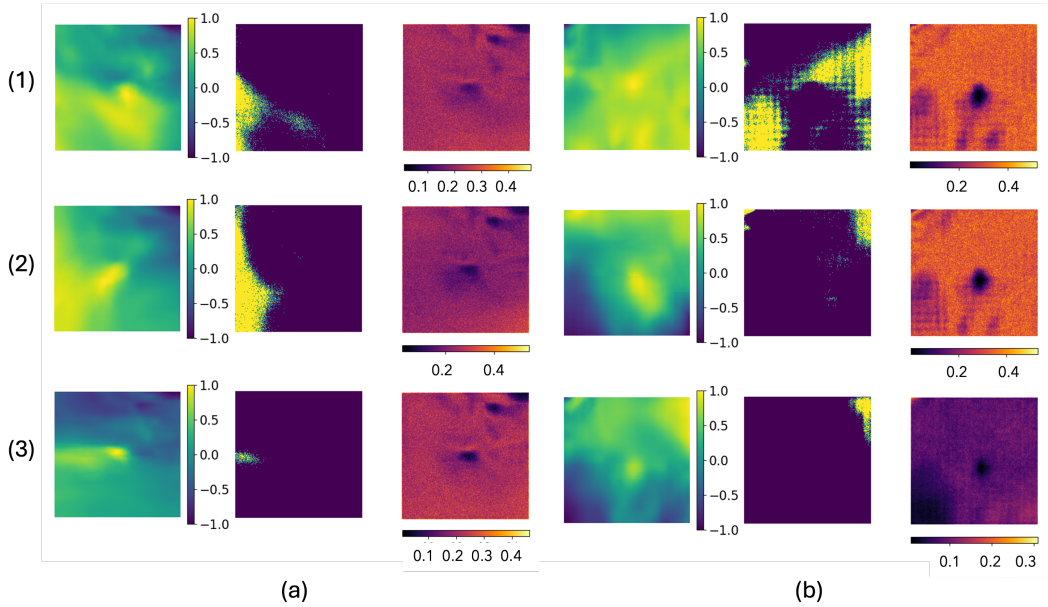

Figure 2: Visualization of CO_2 profiles, predicted anomalies and their variance as detected by ADECEES for the Niederaussem plant (a-1, a-2 and a-3 being from first week of May 2017, May 2018 and May 2021 resp.) and Parish plant (b-1, b-2 and b-3 being from last week of April 2016, April 2018 and April 2019 resp.).

Figure 2-a highlights the potential impact of two units' closure in 2020 and 2021 on the Niederaussem plant. The regional CO_2 profiles have significantly changed between 2017 and 2021: the rows associated with 2017 and 2018 (a-1 and a-2) present a pic in pollution in the area surrounding the plant, which disappears in 2021 (a-3). Similarly, Figure 2-b presents a significant decrease in flagged anomaly areas after the introduction of the carbon capture upgrade in 2017 (b-2 and b-3). This provides a complementary approach to exisiting emissions quantification studies, by considering the temporal aspect of emissions and their impact on the regional CO_2 landscape.

5 Discussion and conclusion

In this paper we introduce ADECEES, a flexible system for detecting CO_2 anomalies globally, aming to support the adoption of efficient measures to reduce global warming. We begin by evaluating the system's ability to identify emission sources before applying it to the detection of shifts in CO_2 emissions. Our results suggests the system's ability to detect point sources and variations in emission but location significantly influences its performance. We argue that ADECEES can provide valuable insights when combined with existing research on CO_2 monitoring [11, 23]. Following this first iteration, future works will focus on improving the partial diffusion process by considering auxiliary variables such as wind or vegetation and further validating the results presented in this work.

References

[1] Nick J. Dunstone, Doug M. Smith, Chris Atkinson, Andrew Colman, Chris Folland, Leon Hermanson, Sarah Ineson, Rachel Killick, Colin Morice, Nick Rayner, Melissa Seabrook, and Adam A. Scaife. Will 2024 be the first year that global temperature exceeds 1.5°c? *Atmospheric Science Letters*, 25:e1254, 9 2024. ISSN 1530-261X. doi: 10.1002/ASL.1254.

[2] Joeri Rogelj, Taryn Fransen, Michel G.J. den Elzen, Robin D. Lamboll, Clea Schumer, Takeshi Kuramochi, Frederic Hans, Silke Mooldijk, and Joana Portugal-Pereira. Credibility gap in net-zero climate targets leaves world at high risk. *Science*, 380:1014–1016, 6 2023. ISSN 10959203. doi: 10.1126/SCIENCE.ADG6248/SUPPL_FILE/SCIENCE.ADG6248_SM.PDF. URL `https://www.science.org/doi/10.1126/science.adg6248`.

[3] United Nations. Paris agreement, 2015. URL `https://unfccc.int/sites/default/files/english_paris_agreement.pdf`.

[4] Frauke Röser, Oscar Widerberg, Niklas Höhne, and Thomas Day. Ambition in the making: analysing the preparation and implementation process of the nationally determined contributions under the paris agreement. *Climate Policy*, 20(4):415–429, 2020.

[5] Yang Ou, Gokul Iyer, Allen Fawcett, Nathan Hultman, Haewon McJeon, Shaun Ragnauth, Steven J Smith, and James Edmonds. Role of non-CO2 greenhouse gas emissions in limiting global warming. *One Earth*, 5(12):1312–1315, 2022.

[6] Lavender Yang, Nicholas Z Muller, and Pierre Jinghong Liang. The real effects of mandatory csr disclosure on emissions: Evidence from the greenhouse gas reporting program. Technical report, National Bureau of Economic Research, 2021.

[7] EEA EEA. EMEP/EEA air pollutant emission inventory guidebook–2016. *European Environment Agency, Copenhagen*, 2016.

[8] Chris J van Staden and Jill Hooks. A comprehensive comparison of corporate environmental reporting and responsiveness. *The British accounting review*, 39(3):197–210, 2007.

[9] Florence Depoers, Thomas Jeanjean, and Tiphaine Jérôme. Voluntary disclosure of greenhouse gas emissions: Contrasting the carbon disclosure project and corporate reports. *Journal of Business Ethics*, 134:445–461, 2016.

[10] 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, 2019.

[11] Ge Han, Yiyang Huang, Tianqi Shi, Hongyuan Zhang, Siwei Li, Haowei Zhang, Weibiao Chen, Jiqiao Liu, and Wei Gong. Quantifying co2 emissions of power plants with aerosols and carbon dioxide lidar onboard dq-1. *Remote Sensing of Environment*, 313:114368, 2024.

[12] Ray Nassar, Jon-Paul Mastrogiacomo, William Bateman-Hemphill, Callum McCracken, Cameron G MacDonald, Tim Hill, Christopher W O'Dell, Matthäus Kiel, and David Crisp. Advances in quantifying power plant co2 emissions with oco-2. *Remote Sensing of Environment*, 264:112579, 2021.

[13] Joëlle Hanna, Michael Mommert, Linus Mathias Scheibenreif, and Damian Borth. Multitask learning for estimating power plant greenhouse gas emissions from satellite imagery. Tackling Climate Change with Machine Learning workshop at NeurIPS., 2021.

[14] Heather D Couture, Madison Alvara, Jeremy Freeman, Aaron Davitt, Hannes Koenig, Ali Rouzbeh Kargar, Joseph O'Connor, Isabella Söldner-Rembold, André Ferreira, Jeyavinoth Jeyaratnam, et al. Estimating carbon dioxide emissions from power plant water vapor plumes using satellite imagery and machine learning. *Remote Sensing*, 16(7):1290, 2024.

[15] Ryu Saito, PK Patra, Nicholas Deutscher, D Wunch, K Ishijima, V Sherlock, T Blumenstock, S Dohe, David Griffith, F Hase, et al. Latitude-time variations of atmospheric column-average dry air mole fractions of co 2, ch 4 and n 2 o. *Atmospheric Chemistry and Physics*, 12(16): 7767–7777, 2012.

[16] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*, 56(4):1–39, 2023.

[17] Julian Wyatt, Adam Leach, Sebastian M Schmon, and Chris G Willcocks. Anoddpm: Anomaly detection with denoising diffusion probabilistic models using simplex noise. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 650–656, 2022.

[18] Julia Wolleb, Florentin Bieder, Robin Sandkühler, and Philippe C Cattin. Diffusion models for medical anomaly detection. In *International Conference on Medical image computing and computer-assisted intervention*, pages 35–45. Springer, 2022.

[19] Andrianirina Rakotoharisoa, Simone Cenci, and Rossella Arcucci. A high resolution spatially consistent global dataset for co2 monitoring. *Remote Sensing*, 17(9):1617, 2025.

[20] Brad Weir, Lesley Ott, OCO-2 Science Team, et al. Oco-2 geos level 3 daily, 0.5 x0. 625 assimilated co2 v10r. *Goddard Earth Sci. Data Inf. Serv. Cent. (GES DISC)*, 2022.

[21] Ye Yuan, Ludwig Ries, Hannes Petermeier, Thomas Trickl, Michael Leuchner, Cédric Couret, Ralf Sohmer, Frank Meinhardt, and Annette Menzel. On the diurnal, weekly, and seasonal cycles and annual trends in atmospheric co 2 at mount zugspitze, germany, during 1981–2016. *Atmospheric Chemistry and Physics*, 19(2):999–1012, 2019.

[22] Masahito Ueyama and Tomoya Ando. Diurnal, weekly, seasonal, and spatial variabilities in carbon dioxide flux in different urban landscapes in sakai, japan. *Atmospheric Chemistry and Physics*, 16(22):14727–14740, 2016.

[23] Xiaojuan Lin, Ronald Van Der A, Jos De Laat, Henk Eskes, Frédéric Chevallier, Philippe Ciais, Zhu Deng, Yuanhao Geng, Xuanren Song, Xiliang Ni, et al. Monitoring and quantifying co 2 emissions of isolated power plants from space. *Atmospheric Chemistry and Physics*, 23(11): 6599–6611, 2023.

[24] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

[25] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015.

[26] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. pmlr, 2015.

[27] Qinsheng Zhang and Yongxin Chen. Fast sampling of diffusion models with exponential integrator. *arXiv preprint arXiv:2204.13902*, 2022.

[28] Jeffrey R Scherffius, Satish Reddy, John P Klumpyan, and Anthony Armpriester. Large-scale co2 capture demonstration plant using fluor's econamine fg plussm technology at nrg's wa parish electric generating station. *Energy Procedia*, 37:6553–6561, 2013.

Annex

A Partial Diffusion

During the forward or diffusion process, an original image $x_0 \sim q(x_0)$ is gradually corrupted following $q(x_t|x_{t-1})$(see Eq. 4) while the reverse or backward process is learned(see Eq. 5) and starts from a noise distribution (commonly gaussian) $x_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ to progressively generate samples from $q(x_0)$(see Eq. 5) [24].

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t|x_{t-1}\sqrt{1-\beta_t}, \beta_t\mathbf{I}) \tag{4}$$

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}|\mu_\theta(x_t, t), \hat{\beta}_t\mathbf{I}) \tag{5}$$

where $\mu_\theta(x_t, t)$ can be implemented or learned [25], $\beta_0, ..., \beta_T$ is a variance schedule and $\hat{\beta}_t = \frac{1-\bar{\alpha}_t}{\bar{\alpha}_t}$ [26] with the notation $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{i=0}^{t} \alpha_i$. In practice, to speed up the slow sampling process [27], Ho et al. [24] show that $\mu_\theta(x_t, t)$ can be parametrized as:

$$\mu_\theta(x_t, t) = \frac{1}{\sqrt{\alpha_t}}(x_t - \frac{\beta_t}{\sqrt{1-\bar{\alpha}_t}}\epsilon_\theta(x_t, t)) \tag{6}$$

while x_t can be sampled from x_0 without the whole sequence as such:

$$q(x_t|x_0) = \mathcal{N}(x_t|x_0\sqrt{\bar{\alpha}_t}, (1-\hat{\alpha}_t)\mathbf{I}) \tag{7}$$

$$x_t = x_0\sqrt{\bar{\alpha}_t} + \epsilon_t\sqrt{1-\bar{\alpha}_t}, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \tag{8}$$

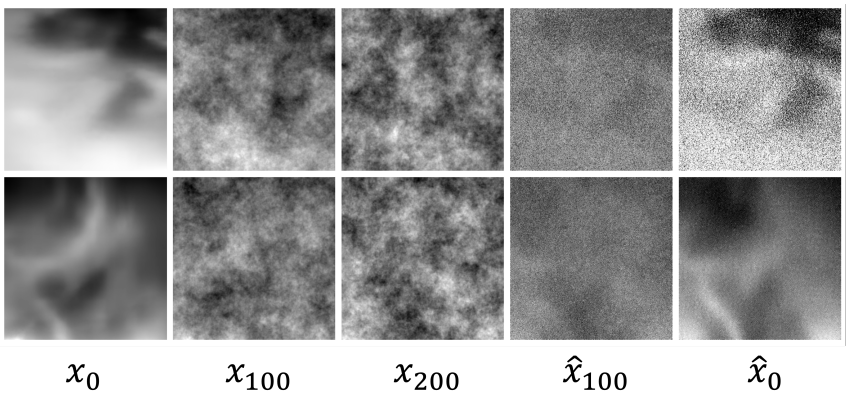

Figure 3: Examples of forward and backward pass of the partial diffusion cycle for T=200.

B Training

In our application, $\mathbf{D}_{healthy}$ represents the dataset containing the average local maps we consider as healthy (eg. when a point source is closed). During the training phase, the model learns to reconstruct them. Eq. 6, gives us that training $\mu_\theta(x_t, t)$ is equivalent to training the estimator $\epsilon_\theta(x_t, t)$ and we set the depth of the forward diffusion pass to T=300.

Figure 4 represents the relationship between the size n of the ensemble and the anomalies segmented by the model. We can see that a smaller size tends to generate loosely bounded segmented areas while a larger ensemble (n=100) will have more intense anomalies with reduced uncertainty.

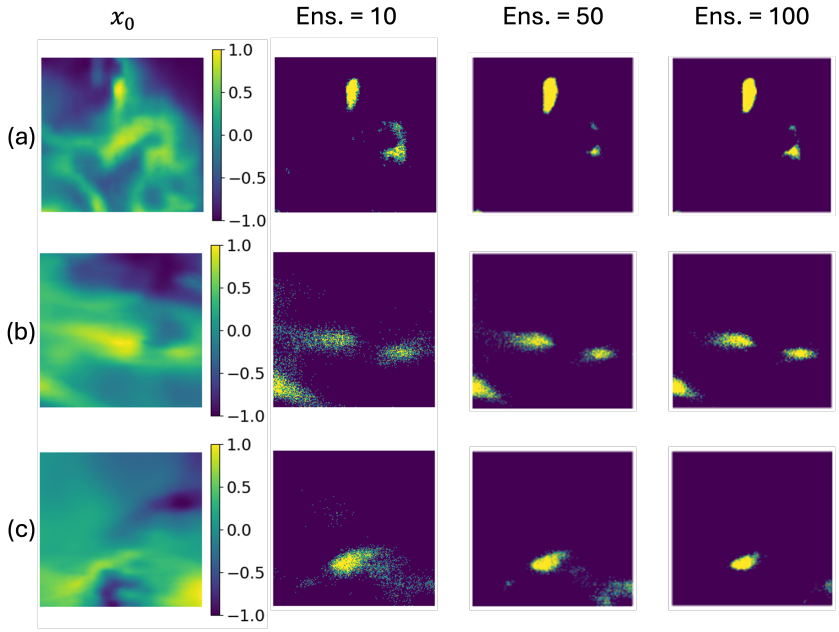

Figure 4: Influence of the ensemble size on the segmentation task. x_0 is the input, the next three columns are the predicted anomalies, influenced by the ensemble size

C Dataset and preprocessing steps

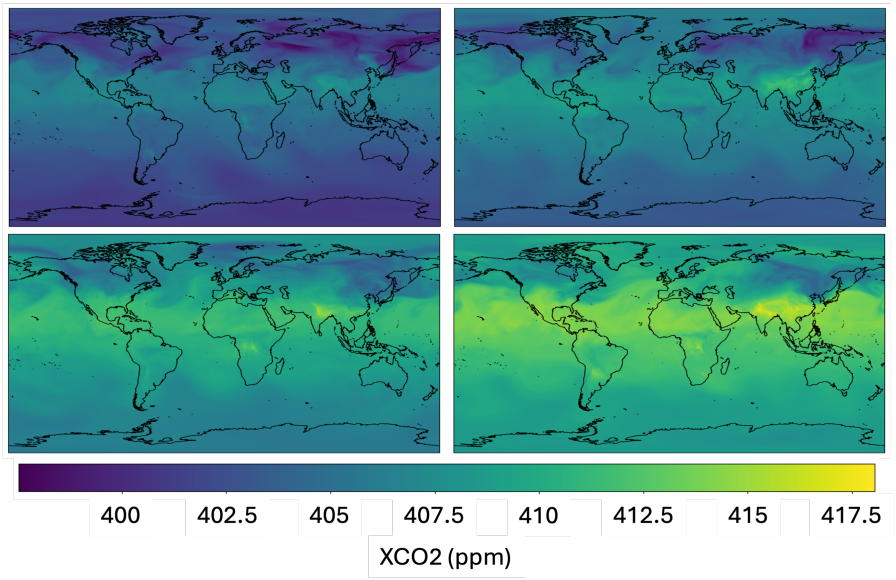

Figure 5: Sample of the high resolution dataset from [19] on 7 July from 2017 to 2020.

On Figure 6, we can see examples of significant changes in relative XCO_2 concentration throughout the week. On x_1 we see a zone of very high CO_2 concentration on the top of the feature map that isn't appearing on other days. Similarly, x_2 and x_3 present higher peaks of XCO_2 in the center left part of the feature map as opposed to other days. By averaging the daily maps over a week, we generate an input x_{weekly} that only retains the common characteristics of each day.

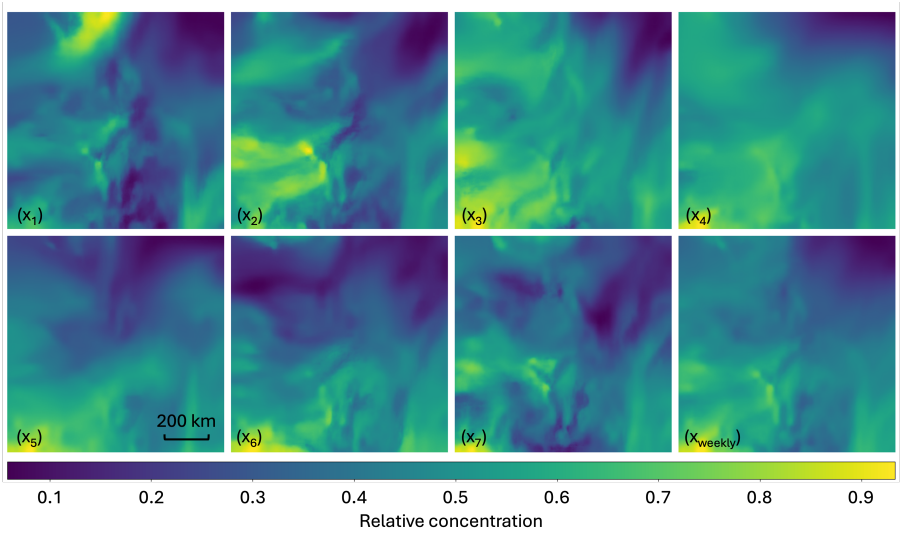

Figure 6: Daily fluctuations in XCO_2 on normalized feature maps. x_1 to x_7 are the XCO_2 daily maps and x_{weekly} is the weekly average. Areas in yellow represent areas of high CO_2 concentration in the region.

D Coal mines and powerplants investigated in this paper

Table 2: Coal mines and their characteristics for emissions detection

Mine (Abbrv.)	Loc. (Lat,Lon)	Prod. (Mtpa)	Active	Closed	Reason
Benga (Bg.)	-16.17°, 33.66°	1.24	07/19–01/21	07/16–01/18	Extension
Otvalny (Ot.)	54.15°, 87.13°	2.50	01/17–12/17	01/21–12/21	Rehabilitation
Invierno (Iv.)	-53.00°, -72.42°	2.30	06/18–12/19	06/20–12/21	Closed
Hazelwood (Hw.)	-38.25°, 146.38°	15.30	06/18–12/19	06/20–12/21	Rehabilitation

Table 3: Powerplants and the possible reason for emission change

Power plant (Abbrv.)	Loc. (Lat,Lon)	Reason for emission variation
Niederaussem (Nd.)	50.99°, 6.67°	Units got retired in 2020 and 2021
Parish (Pa.)	29.48°, -95.63°	Unit upgraded with carbon capture in 2017 [28]