

# CONDITIONAL DIFFUSION-BASED RETRIEVAL OF ATMOSPHERIC CO<sub>2</sub> FROM EARTH OBSERVING SPECTROSCOPY

**William R. Keely**

The University of Oklahoma  
Norman, OK 73019, USA  
william.r.keely@ou.edu

**Otto Lamminpää**

Jet Propulsion Laboratory  
California Institute of Technology  
Pasadena, CA 91109, USA  
otto.m.lamminpaa@jpl.nasa.gov

**Steffen Mauceri**

Jet Propulsion Laboratory  
California Institute of Technology  
Pasadena, CA 91109, USA  
steffen.mauceri@jpl.nasa.gov

**Sean M. R. Crowell**

LumenUs Scientific  
Oklahoma City, OK 73102, USA  
sean@belumenus.com

**Christopher W. O'Dell**

Cooperative Institute for Atmospheric Research  
Colorado State University  
Fort Collins, CO 80521, USA  
christopher.odell@colostate.edu

**Gregory, R. McGarragh**

Cooperative Institute for Atmospheric Research  
Colorado State University  
Fort Collins, CO 80521, USA  
greg.mcgarragh@colostate.edu

## ABSTRACT

Satellite-based estimates of greenhouse gas (GHG) properties from observations of reflected solar spectra are integral for understanding and monitoring complex terrestrial systems and their impact on the carbon cycle due to their near global coverage. Known as *retrieval*, making GHG concentration estimations from these observations is a non-linear Bayesian inverse problem, which is operationally solved using a computationally expensive algorithm called Optimal Estimation (OE), providing a Gaussian approximation to a non-Gaussian posterior. This leads to issues in solver algorithm convergence, and to unrealistically confident uncertainty estimates for the retrieved quantities. Upcoming satellite missions will provide orders of magnitude more data than the current constellation of GHG observers. Development of fast and accurate retrieval algorithms with robust uncertainty quantification is critical. Doing so stands to provide substantial climate impact of moving towards the goal of near continuous real-time global monitoring of carbon sources and sinks which is essential for policy making. To achieve this goal, we propose a diffusion-based approach to flexibly retrieve a Gaussian or non-Gaussian posterior, for NASA's Orbiting Carbon Observatory-2 spectrometer, while providing a substantial computational speed-up over the current operational state-of-the-art.

## 1 INTRODUCTION

Satellite observations of greenhouse gases (GHGs) are essential for monitoring and quantifying the impact of anthropogenic activities (e.g., fossil fuel combustion, deforestation, urbanization) on the climate of the planet particularly in rapidly changing regions like the high latitudes and Tropics which are difficult to measure *in situ*. Satellite-borne spectrometers such as NASA's Orbiting Carbon Observatory-2 and 3 (OCO-2/3; Crisp (2015)), and JAXA's Greenhouse Gases Observing Satellite (GOSAT; Yokota et al. (2009)) make high spectral resolution measurements of reflected solar spectra in order to infer changes in GHGs. A physics-based algorithm is used to estimate or

retrieve the atmospheric state including the column averaged dry air-mole fraction of atmospheric carbon dioxide ( $\text{XCO}_2$ ) from these spectra.  $\text{XCO}_2$  observations are used to constrain surface fluxes of GHGs, which are further used to study the Earth’s complex carbon cycle. (Crowell et al., 2019).

Estimating GHG concentrations from spectral observations is typically posed as an inverse problem, where we aim to recover the state of the atmosphere, denoted  $\mathbf{x}$ , from observed spectra  $\mathbf{y}$  with instrument viewing geometries  $\mathbf{b}$ . The relationship between state and observation can be written as:  $\mathbf{y} = F(\mathbf{x}, \mathbf{b}) + \nu$ , where  $F$  is the forward operator, and noise  $\nu \sim \mathcal{N}(0, \mathbf{R})$  which is scene and wavelength-dependent. OCO-2 uses an Optimal Estimation (OE) based algorithm called Atmospheric Carbon Observations from Space (ACOS) detailed in Connor et al. (2008); (O’Dell et al., 2012). OE methods solve the inverse problem using Bayes Rule:  $p(\mathbf{x}|\mathbf{y}) = p(\mathbf{y}|\mathbf{x})p(\mathbf{x})$ , where estimates of  $\mathbf{x}$  are obtained as a *maximum a posteriori* (MAP) point estimate which maximizes the posterior probability  $p(\mathbf{x}|\mathbf{y})$  (Rodgers, 2000) with an associated Gaussian approximation to the posterior (co-)variance as an uncertainty estimate. The ACOS retrieval employs a Radiative Transfer Model forward model (FM) coupled with an instrument model for the OCO-2 spectrometer as  $F$  to produce simulated spectra,  $\mathbf{y}_{\text{sim}}$ , and partial derivatives for a given  $\mathbf{x}$ . The algorithm iteratively seeks the MAP estimate that minimizes a cost function (Appendix C) that represents the sum of squared errors between the observations  $\mathbf{y}_{\text{Obs}}$  and  $\mathbf{y}_{\text{sim}}$ .

Hobbs et al. (2017), Lamminpää et al. (2019) and Braverman et al. (2021) illustrate how realizations of the true conditional posterior for  $\text{XCO}_2$  can be non-Gaussian or bimodal due to the nonlinear nature of the forward model, arising from varying aerosol and surface conditions. The classical formulation does not apply in these cases. The Gaussian assumed uncertainty estimate produced by ACOS is often over confident, e.g., the prediction interval often does not contain the true value of  $\text{XCO}_2$ . Applications of these data to constrain  $\text{CO}_2$  emissions can be sensitive to the assumed uncertainties and often incorporate ad hoc adjustments to the uncertainties prior to their use. This introduces additional uncertainties in the  $\text{CO}_2$  emissions estimates themselves (Peiro et al., 2022). The uncertainty in predicted emissions is critical for inference at policy relevant (e.g., national) scales necessary for treaties and enforcement (Byrne et al., 2022).

OE requires repeated evaluations of the computationally expensive forward FM per observation. OCO-2 currently makes  $10^5$  cloud-free observations per day and has a record that spans from 2014 to present. The decade-long catalog has undergone several updates, with reprocessing taking around a year to complete with the current ACOS version and modern HPC architectures. This computational bottleneck will be prohibitive for future missions such as ESA’s upcoming Carbon Dioxide Monitoring Mission (Sierk et al., 2021), which will generate tens of millions of observations daily.

It is critical to develop new retrieval techniques that are fast and accurate while also providing robust uncertainties. In this work, we apply a diffusion-based approach for accurate atmospheric  $\text{CO}_2$  retrieval that avoids the need for repeated FM evaluations after training is complete, while also providing a robust uncertainty for NASA’s OCO-2 instrument. Diffusion is a generative machine learning method that allows for efficient sampling of potentially complex posteriors and has seen application in solving inverse problems in the natural sciences. Diffusion’s success comes from the ability to define a prior distribution which is typically a standard Gaussian  $\mathcal{N}(0, \mathbf{I})$  and then training a *denoising* model that iteratively learns to shift samples drawn from the prior to the true data distribution using a Markov process. This allows for efficient sampling to recover either a Gaussian uncertainty by taking the first two statistical moments of the draws or to explore non-Gaussianity in the retrieved posterior.

## 2 METHODS

To train the retrieval we use the L2 product (Payne et al., 2023) of version 11 ACOS FM simulated spectra denoted as  $\mathbf{y}_{\text{sim}}$ . OCO-2’s wavelength ranges are the near-infrared Oxygen A band near  $0.76 \mu\text{m}$ , the shortwave infrared ‘weak’  $\text{CO}_2$  band near  $1.6 \mu\text{m}$ , and the shortwave infrared ‘strong’  $\text{CO}_2$  band near  $2.05 \mu\text{m}$ . Additional retrieval covariates include the solar zenith angle ( $SZA$ ) and a weather model predicted surface pressure term  $P_{\text{surf}}$  that comes from analysis produced by the GEOS5-FP-IT data assimilation system (Lucchesi, 2013). The concatenated covariate vector for training point  $i$  is denoted as  $\bar{\mathbf{y}}_i = \{\mathbf{y}_i, SZA_i, P_{\text{surf}_i}\}$ . The  $\text{XCO}_2$  scalars from the L2 state vector  $\mathbf{x}$ , in parts per million (ppm), serve as the training label  $x_i$ . In total, we construct a training dataset of  $3 \times 10^6 \{\bar{\mathbf{y}}_i, x_i\}$  pairs. Dependence on the FM for generating training data can introduce errors

into the retrieval due to differences between simulated training radiance and real observations seen at inference time:  $(\mathbf{y}_{\text{Sim}} - \mathbf{y}_{\text{Obs}})$ . We account for this difference by careful perturbation of the simulated radiance prior to model training, as shown in Appendix B.1 and resample both  $\mathbf{y}_{\text{Sim}}$  and  $\mathbf{y}_{\text{Obs}}$  to a common wavelength grid. We use the FM for training despite the abundance of OCO-2 observations with an eye to current and future missions that may lack sufficient temporal or spatial coverage of observations for an ML-based retrieval. These missions cannot adopt ML algorithms trained directly on observations to generalize well across the globe, thus requiring learning on simulated training data.

Given realizations of the covariates, we aim to retrieve the posterior distribution  $p(x|\bar{\mathbf{y}})$ , using the conditional diffusion method described in Han et al. (2022). This paper suggests the use of a conditional mean prior,  $x_{\text{prior}}$ , which is a pre-trained neural network  $f_\phi$  with weights  $\phi$  to parameterize the forward diffusion distribution as  $\mathcal{N}(x_{\text{prior}}, 1)$  instead of the more standard choice of  $\mathcal{N}(0, 1)$ . This introduces conditioning information into the forward process and potentially speeds up inference by reducing the number of diffusion steps required. We explore different methods to specify  $x_{\text{prior}}$  including a neural network retrieval of  $\text{XCO}_2$  as well as  $\text{XCO}_2$  values from GEOS-FP-IT in Appendix B.2. The denoising model,  $\epsilon_\theta(\bar{\mathbf{y}}_i, x_{\text{prior}_i}, x_{i,t}, t)$  is then trained as described in Ho et al. (2020) using the mean squared error loss, where  $x_t$  is the noised label at diffusion step  $t$ .

The OCO-2 operational  $\text{XCO}_2$  product includes an extensive post hoc bias correction to reduce systematic errors or *biases* in the raw OE retrieved  $\text{XCO}_2$  (Keely et al., 2023; O’dell et al., 2018). This process involves fitting a regression model on errors in  $\text{XCO}_2$  with predictors selected from the retrieved atmospheric state.  $\text{XCO}_2$  errors are calculated using ground truth labels of  $\text{XCO}_2$  from the Total Carbon Column Observing Network (TCCON; (Wunch et al., 2011)) that are collocated with OCO-2 measurements in space and time. In this spirit, we evaluate the potential of reducing bias in the trained diffusion model using a similar dataset of collocated TCCON  $\text{XCO}_2$  labels to finetune the retrieved value. We collocate OCO-2 and TCCON samples within 100km spatially and 1 hour temporally, for which we also have FM produced radiance for years 2015 to 2021. This results in a finetuning dataset of  $25 \times 10^4$  pairs. To finetune the retrieval, we first unfreeze the last layers of  $f_\phi$  and train an additional 10 epochs. We then continue training  $\epsilon_\theta$  using the finetuned conditional mean backbone until the loss for a small validation set stabilizes.

To evaluate the performance of the proposed diffusion retrieval relative to ACOS, we use a hold-out set of data of OCO-2 observations collocated with TCCON sites and retrieve on the OCO-2 measured reflected solar radiance ( $\mathbf{y}_{\text{Obs}}$ ) for the year of 2022. This test dataset contains 20,000 collocated pairs, and contains examples from all operational TCCON sites for the year. We assess the quality of the point estimates of  $\text{XCO}_2$  for ACOS and the diffusion retrieval with and without finetuning/bias-correction, as well as the quality of their uncertainty estimates.

### 3 RESULTS

#### 3.1 EFFICIENT POSTERIOR SAMPLING:

The reverse diffusion process allows us to generate samples from  $p(x|\bar{\mathbf{y}})$  by first drawing each  $x_T \sim \mathcal{N}(x_{\text{prior}}, 1)$  and then iteratively denoising with  $\epsilon_\theta$  from  $t = T$  to  $t = 0$  to recover the denoised  $\text{XCO}_2$  samples, thus enabling efficient density estimation of potentially complex posteriors as shown in Figure 1. The diffusion retrieval takes on average 0.0003 seconds to draw 10 conditional samples for a single observation on a 2017 Apple MacBook Pro. Lamminpää et al. (2024) showed that FM evaluations on similar hardware take approximately 33-55 seconds.

#### 3.2 COMPARISON TO THE OPERATIONAL OCO-2 RETRIEVAL:

In Figure 2, we evaluate the  $\text{XCO}_2$  point estimate and uncertainty for both the operational ACOS retrieval and our proposed diffusion-based retrieval on a holdout year of OCO-2 observations collocated with TCCON sites for 2022. The diffusion-based retrieval without finetuning shows a 3% improvement to ACOS prior to bias correction. When we compare to the operational product with additional finetuning for the diffusion retrieval, the finetuned diffusion retrieval shows a 10% improvement in RMSE over bias-corrected ACOS retrieved  $\text{XCO}_2$  for the test year. We also compare the uncertainty estimate of both methods. Since the ACOS retrieval enforces a Gaussian form, we also assume a Gaussian prediction interval for the diffusion retrieval by taking the first two moments




Figure 1: Example of bi-modal distribution recovered through 10 conditional samples from the Diffusion posterior (blue) without finetuning compared to the operational OE posterior which is derived from the point estimate  $\pm$  an estimate of the variance (black) without bias correction. The Diffusion posterior covers the ground truth value from TCCON (red) within the first mode which is not captured by the Gaussian operational posterior.

over the diffusion samples. The performance of the uncertainties is assessed by calculating the Miscalibration Area which is the integrated region between the calibration curve and the  $x = y$  line. The smaller the area, the better the coverage of the prediction interval for a given percentile. The assumed Gaussian prediction interval from the diffusion retrieval contains the collocated TCCON  $\text{XCO}_2$  value more frequently than the operational retrieval. Finetuning provides an added benefit of further calibrating the uncertainty and reducing the Miscalibration Area on the holdout year. In B.3, we train several additional types of deep learning models that can provide uncertainty quantification of their estimate and compare them to diffusion and ACOS.




Figure 2: Comparison of ACOS and Diffusion with and without bias correction/finetuning for the holdout year of 2022. We compare the point estimate of each method (top row) by evaluating the RMSE. We compare the Gaussian uncertainties of each method (bottom row) by evaluating the Miscalibration Area.

## 4 CONCLUSIONS AND FUTURE WORK

We present a diffusion-based approach for retrieving atmospheric  $\text{XCO}_2$  from satellite observations that is competitive with the current state-of-the-art approach while offering the for an increase of several orders of magnitude in speed during inference. The ability to quickly sample the conditional posterior also allows for improved uncertainty quantification compared to the operational product, thus allowing for an improved constraint of flux estimates by inversion models. In future work, we will further benchmark the computational speed up against the operational retrieval as well as derive an averaging kernel which is essential for assimilation of  $\text{XCO}_2$  observations into carbon flux inversion models. It will also be important to assess the potential for the finetuning process to introduce biases in regions where TCCON is not available, since regionally coherent systematic errors in  $\text{XCO}_2$  propagate directly into inferred  $\text{CO}_2$  surface fluxes. The approach is highly generalizable to retrieval of other important GHGs such as methane ( $\text{CH}_4$ ) and to future GHG missions such as Carbon-I (Frankenberg et al., 2024) and  $\text{CO}_2\text{M}$ , which will produce large data sets directly impacting future climate policy.

### ACKNOWLEDGMENTS

We would like to deeply thank the OCO-2/3 Science and Algorithm Teams.

We appreciate the TCCON teams, who measure and provide ground-based  $\text{XCO}_2$  validation to the carbon cycle research community ((De Mazière et al., 2014); (Té et al., 2022), (Liu et al., 2023); (Hase et al., 2022); (Kivi et al., 2014); (Weidmann et al.); (Zhou et al., 2022); (Deutscher et al.); (Dubey et al., 2022); (Wunch et al., 2022); (Garcia et al., 2022); (Shiomi et al., 2022); (Strong et al., 2022); (Notholt et al., 2022); (Morino et al., 2022a); (Morino et al., 2022b); (Morino et al., 2022c); (Sussmann & Rettinger, 2023); (Buschmann et al., 2022); (Sherlock et al., 2022a); (Wennberg et al., 2022a); (Wennberg et al., 2022b))

This work was funded by NASA grant: **80NSSC24K0763**.

### DATA AVAILABILITY

OCO-2 data including the forward model spectra are available at: [https://disc.gsfc.nasa.gov/datasets/OCO2\\_L2\\_Diagnostic\\_11r/summary?keywords=oco2%20diagnostic](https://disc.gsfc.nasa.gov/datasets/OCO2_L2_Diagnostic_11r/summary?keywords=oco2%20diagnostic). TCCON data is available at: <https://tccon-wiki.caltech.edu/Main/DataDescriptionGGG2020>. We also provide the predictions and Gaussian uncertainty from the diffusion model with collocated TCCON and ACOS estimates here: [10.5281/zenodo.15115961](https://doi.org/10.5281/zenodo.15115961)

## REFERENCES

- Alexander Amini, Wilko Schwarting, Ava Soleimany, and Daniela Rus. Deep evidential regression. *Advances in neural information processing systems*, 33:14927–14937, 2020.
- Amy Braverman, Jonathan Hobbs, Joaquim Teixeira, and Michael Gunson. Post hoc uncertainty quantification for remote sensing observing systems. *SIAM/ASA Journal on Uncertainty Quantification*, 9(3):1064–1093, 2021.
- François-Marie Bréon, Leslie David, Pierre Chatelanaz, and Frédéric Chevallier. On the potential of a neural-network-based approach for estimating  $\text{xco}_2$  from oco-2 measurements. *Atmospheric Measurement Techniques*, 15(18):5219–5234, 2022.
- J. Brynjarsdóttir, J. Hobbs, A. Braverman, and L. Mandrake. Optimal estimation versus mcmc for  $\text{co}_2$  retrievals. *Journal of Agricultural, Biological and Environmental Statistics*, 23:297–316, 2018. doi: 10.1007/s13253-018-0319-8.
- M Buschmann, C Petri, M Palm, T Warneke, and J Notholt. Tccon data from ny-ålesund, svalbard (no), release ggg2020. r0. (*No Title*), 2022.
- Brendan Byrne, David F Baker, Sourish Basu, Michael Bertolacci, Kevin W Bowman, Dustin Carroll, Abhishek Chatterjee, Frédéric Chevallier, Philippe Ciais, Noel Cressie, et al. National co 2

- budgets (2015–2020) inferred from atmospheric co<sub>2</sub> observations in support of the global stock-take. *Earth System Science Data Discussions*, 2022:1–59, 2022.
- Brian J Connor, Hartmut Boesch, Geoffrey Toon, Bhaswar Sen, Charles Miller, and David Crisp. Orbiting carbon observatory: Inverse method and prospective error analysis. *Journal of Geophysical Research: Atmospheres*, 113(D5), 2008.
- David Crisp. Measuring atmospheric carbon dioxide from space with the orbiting carbon observatory-2 (oco-2). In *Earth observing systems xx*, volume 9607, pp. 960702. SPIE, 2015.
- Sean Crowell, David Baker, Andrew Schuh, Sourish Basu, Andrew R Jacobson, Frederic Chevallier, Junjie Liu, Feng Deng, Liang Feng, Kathryn McKain, et al. The 2015–2016 carbon cycle as seen from oco-2 and the global in situ network. *Atmospheric Chemistry and Physics*, 19(15):9797–9831, 2019.
- Leslie David, Francois Marie Breon, and Frederic Chevallier. Xco<sub>2</sub> estimates from the oco-2 measurements using a neural network approach. *Atmospheric Measurement Techniques*, 14:117–132, 1 2021. ISSN 18678548. doi: 10.5194/AMT-14-117-2021.
- M De Mazière, MK Sha, F Desmet, C Hermans, F Scolas, N Kumps, J-M Metzger, V Dufлот, and J-P Cammas. Tccon data from réunion island (re), release ggg2014. r0. (*No Title*), 2014.
- NM Deutscher, DWT Griffith, C Paton-Walsh, NB Jones, VA Velazco, SR Wilson, RC Macatangay, GC Kettlewell, RR Buchholz, MO Riggensbach, et al. Tccon data from wollongong, australia, release ggg2020r0, tccon data archive, hosted by caltechdata [data set], california institute of technology, pasadena, ca, usa.
- MK Dubey, BG Henderson, NT Allen, J-F Blavier, CM Roehl, and D Wunch. Tccon data from manaus (br), release ggg2020. r0. (*No Title*), 2022.
- Christian Frankenberg, Anna Michalak, Daniel J Jacob, Andrew K Thorpe, Yi Yin, Lori Bruhwiler, Ermias Kebeab, Alison Hoyt, Alex J Turner, Paul O Wennberg, et al. Carbon-i, a nasa earth system explorer mission concept for greenhouse gas observations. *AGU24*, 2024.
- Omaira Garcia, Matthias Schneider, Benedikt Herkommer, Jochen Groß, Frank Hase, Thomas Blumenstock, and Eliezer Sepulveda. Tccon data from izana (es), release ggg2020. r1. 2022.
- Xizewen Han, Huangjie Zheng, and Mingyuan Zhou. Card: Classification and regression diffusion models. *Advances in Neural Information Processing Systems*, 35:18100–18115, 2022.
- Frank Hase, Benedikt Herkommer, Jochen Groß, Thomas Blumenstock, Matthäus Kiel, and Susanne Dohe. Tccon data from karlsruhe (de), release ggg2020. r0. 2022.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Jonathan Hobbs, Amy Braverman, Noel Cressie, Robert Granat, and Michael Gunson. Simulation-based uncertainty quantification for estimating atmospheric co<sub>2</sub> from satellite data. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1):956–985, 2017.
- William Keely and Steffen Mauceri. Towards a machine learning retrieval of xco<sub>2</sub> from oco-2 radiances with uncertainty quantification. In *AGU Fall Meeting Abstracts*, volume 2022, pp. A14E–06, 2022.
- William R Keely, Steffen Mauceri, Sean Crowell, and Christopher W O’Dell. A nonlinear data-driven approach to bias correction of xco<sub>2</sub> for nasa’s oco-2 acos version 10. *Atmospheric Measurement Techniques*, 16(23):5725–5748, 2023.
- R Kivi, P Heikkinen, and E Kyrö. Tccon data from sodankylä (fi), release ggg2014. r0. (*No Title*), 2014.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.

- Otto M Lamminpää, Jouni I Susiluoto, Jonathan M Hobbs, James L McDuffie, Amy J Braverman, and Houman Owhadi. Forward model emulator for atmospheric radiative transfer using gaussian processes and cross validation. *Atmospheric Measurement Techniques Discussions*, 2024:1–32, 2024.
- Otto Lamminpää, Jonathan Hobbs, Jenný Brynjarsdóttir, Marko Laine, Amy Braverman, Hannakaisa Lindqvist, and Johanna Tamminen. Accelerated mcmc for satellite-based measurements of atmospheric  $\text{CO}_2$ . *Remote Sensing*, 11(17), 2019. ISSN 2072-4292. doi: 10.3390/rs11172061. URL <https://www.mdpi.com/2072-4292/11/17/2061>.
- C Liu, W Wang, Y Sun, and C Shan. Tccon data from hefei (prc), release ggg2020. r1, 2023.
- R Lucchesi. File specification for geos-5 fp-it (forward processing for instrument teams). Technical report, NASA, 2013.
- Wesley J Maddox, Pavel Izmailov, Timur Garipov, Dmitry P Vetrov, and Andrew Gordon Wilson. A simple baseline for bayesian uncertainty in deep learning. *Advances in neural information processing systems*, 32, 2019.
- I Morino, H Ohyama, A Hori, and H Ikegami. Tccon data from tsukuba (jp), 125hr, release ggg2020. r0. (No Title), 2022a.
- I Morino, H Ohyama, A Hori, and H Ikegami. Tccon data from rikubetsu (jp), release ggg2020. r0. (No Title), 2022b.
- Isamu Morino, Voltaire A Velazco, Akihiro Hori, Osamu Uchino, and David WT Griffith. Tccon data from burgos, ilocos norte (ph), release ggg2020. r0. (No Title), 2022c.
- Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pp. 8162–8171. PMLR, 2021.
- J Notholt, C Petri, T Warneke, and M Buschmann. Tccon data from bremen (de), release ggg2020. r0. (No Title), 2022.
- Christopher W O’dell, Annmarie Eldering, Paul O Wennberg, David Crisp, Michael R Gunson, Brendan Fisher, Christian Frankenberg, Matthias Kiel, Hannakaisa Lindqvist, Lukas Mandrake, et al. Improved retrievals of carbon dioxide from orbiting carbon observatory-2 with the version 8 acos algorithm. *Atmospheric Measurement Techniques*, 11(12):6539–6576, 2018.
- CW O’Dell, B Connor, H Bösch, D O’Brien, C Frankenberg, R Castano, M Christi, D Eldering, B Fisher, M Gunson, et al. The acos  $\text{CO}_2$  retrieval algorithm—part 1: Description and validation against synthetic observations. *Atmospheric Measurement Techniques*, 5(1):99–121, 2012.
- V. H. Payne, A. Chatterjee, R. Rosenberg, M. Kiel, B. Fisher, and L. Dang. Carbon observatory-2 & 3 (oco-2 & oco-3) data product user’s guide, operational level 2 lite files. Technical report, 2023. URL [https://docserver.gesdisc.eosdis.nasa.gov/public/project/OCO/OCO2\\_V11.1\\_OCO3\\_V10\\_L2\\_Data\\_Users\\_Guide\\_v3.0\\_RevA.pdf](https://docserver.gesdisc.eosdis.nasa.gov/public/project/OCO/OCO2_V11.1_OCO3_V10_L2_Data_Users_Guide_v3.0_RevA.pdf).
- Hélène Peiro, Sean Crowell, Andrew Schuh, David F Baker, Chris O’Dell, Andrew R Jacobson, Frédéric Chevallier, Junjie Liu, Annmarie Eldering, David Crisp, et al. Four years of global carbon cycle observed from the orbiting carbon observatory 2 (oco-2) version 9 and in situ data and comparison to oco-2 version 7. *Atmospheric Chemistry and Physics*, 22(2):1097–1130, 2022.
- Clive D Rodgers. *Inverse methods for atmospheric sounding: theory and practice*, volume 2. World scientific, 2000.
- V Sherlock, B Connor, J Robinson, H Shiona, D Smale, and DF Pollard. Tccon data from lauder (nz), 125hr, release ggg2020. r0 (version r0)[data set]. caltechdata, 2022a.
- K Shiomi, S Kawakami, H Ohyama, K Arai, H Okumura, H Ikegami, and M Usami. Tccon data from saga (jp), release ggg2020. r0. (No Title), 2022.

- Bernd Sierk, Valerie Fernandez, J-L Bézy, Y Meijer, Y Durand, G Bazalgette Courrèges-Lacoste, C Pachot, A Löscher, H Nett, K Minoglou, et al. The copernicus co2m mission for monitoring anthropogenic carbon dioxide emissions from space. In *International Conference on Space Optics—ICSO 2020*, volume 11852, pp. 1563–1580. SPIE, 2021.
- K Strong, S Roche, JE Franklin, J Mendonca, E Lutsch, D Weaver, PF Fogal, JR Drummond, R Batchelor, R Lindenmaier, et al. Tccon data from eureka (ca), release ggg2020. r0 (version r0)[data set]. caltechdata, 2022.
- R Sussmann and M Rettinger. Tccon data from garmisch (de), release ggg2020. r0 (version r0)[data set]. caltechdata, 2023.
- Yao Té, Pascal Jeseck, and Christof Janssen. Tccon data from paris (fr), release ggg2020. r0. (*No Title*), 2022.
- Michael Turmon and Amy Braverman. Uncertainty quantification for JPL retrievals. Technical report, 2019. URL <https://hdl.handle.net/2014/45978>.
- D Weidmann, R Brownsword, and S Doniki. Tccon data from harwell, oxfordshire (uk), release ggg2020. r0, caltechdata[data set].
- PO Wennberg, CM Roehl, D Wunch, GC Toon, J-F Blavier, R Washenfelter, G Keppel-Aleks, and NT Allen. Tccon data from park falls (us), release ggg2020. r0. (*No Title*), 2022a.
- PO Wennberg, D Wunch, Y Yavin, GC Toon, J-F Blavier, NT Allen, G Keppel-Aleks, and CM Roehl. Tccon data from jet propulsion laboratory (us), 2007, release ggg2020. r0. (*No Title*), 2022b.
- Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P Xing. Deep kernel learning. In *Artificial intelligence and statistics*, pp. 370–378. PMLR, 2016.
- D Wunch, J Mendonca, O Colebatch, NT Allen, JF Blavier, K Kunz, S Roche, J Hedelius, G Neufeld, S Springett, et al. Tccon data from east trout lake, sk (ca), release ggg2020. r0 (version r0)[data set]. caltechdata, 2022.
- Debra Wunch, Geoffrey C Toon, Jean-François L Blavier, Rebecca A Washenfelter, Justus Notholt, Brian J Connor, David WT Griffith, Vanessa Sherlock, and Paul O Wennberg. The total carbon column observing network. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 369(1643):2011–2112, 2011.
- Fengxin Xie, Tao Ren, Changying Zhao, Yuan Wen, Yilei Gu, Minqiang Zhou, Pucai Wang, Kei Shiomi, and Isamu Morino. Fast retrieval of xco2 over east asia based on orbiting carbon observatory-2 (oco-2) spectral measurements. *Atmospheric Measurement Techniques*, 17(13): 3949–3967, 2024.
- T Yokota, Y Yoshida, N Eguchi, Y Ota, T Tanaka, H Watanabe, and S Maksyutov. Global concentrations of co2 and ch4 retrieved from gosat: First preliminary results. *Sola*, 5:160–163, 2009.
- M Zhou, P Wang, N Kumps, C Hermans, and W Nan. Tccon data from xianghe, china, release ggg2020. r0 (version r0)[dataset]. caltechdata, 2022.

## A RELATED WORK

Machine learning approaches to greenhouse gas retrieval from Earth observing spectroscopy have made positive impact in reducing the computational cost over full physics based approaches. For OCO-2, David et al. (2021) and Bréon et al. (2022) offered a first attempt at directly retrieving column CO<sub>2</sub> by pairing XCO<sub>2</sub> labels from the carbon inversion model CAMS directly with observations. Their approach improved computational speed but did not improve on precision due to the coarse scale of the CAMS labels. Keely & Mauzeri (2022) introduced an initial attempt at providing a robust uncertainty quantification for a deep learning retrieval through ensembling but on a limited simulated training data set. Xie et al. (2024) illustrated the limitations of training an

ML retrieval for OCO-2 using observations only, and enhanced the training data set using a fast FM emulator. Lamminpää et al. (2024) proposes emulating the FM with a Gaussian Process to speed-up the iterative forward model evaluations in OE. The non-Gaussianity of posterior distribution of OCO retrieval has been explored using MCMC (Brynjarsdóttir et al. (2018), Lamminpää et al. (2019)) and Simulation Based UQ (Turmon & Braverman (2019), Braverman et al. (2021)). In this work, we propose a diffusion-based approach to efficiently sample the conditional posterior and train directly on radiance produced by the high fidelity FM used by the current operational retrieval algorithm for OCO-2. We also show how finetuning the retrieval using co-sampled ground truth labels from TCCON (Wunch et al., 2011) can make ML retrievals competitive with the current state-of-the-art.

## B EXPERIMENTS

### B.1 EMPIRICAL ORTHOGONAL FUNCTIONS

A matrix of training residuals  $M_j$  is given for each band  $j$ .  $M_j$  is decomposed into its eigenvectors using the singular value decomposition  $M_j = U_j W_j V_j^T$ . The  $k = 1$  eigen vector explains the largest fraction of the total variance. For the OCO-2 operational retrieval, coefficients of these fixed EOFs are retrieved to minimize the effect of errors from forward model discrepancies during retrieval. Hence, they serve as a proxy for the forward model error in the simulated radiances. To reduce similar errors in the ML retrieval we retrain using the EOFs to perturb the simulated training set ideally making the model more robust to sim-to-real differences. Before application to each  $i^{th}$  training point, each EOF  $e_{j,k}$  is randomly scaled by a coefficient  $c_{i,j,k}$  drawn from the distributions formed by the operational scaling terms. Thus, the training radiance for band  $y_{sim}$  become:

$$y'_{sim_{i,j}} = y_{sim_{i,j}} + \sum_{k=1}^K c_{i,j,k} e_{j,k} \quad (1)$$

Figure 3 shows the mean difference between the simulated radiance and observation pairs withheld from fitting of the EOFs. ACOS includes the coefficients of the EOFs in the OE state vector. In future work we will look at retrieving the scaling term in a similar fashion using the diffusion-based retrieval as part of an expanded retrieved state. In Table 1, we perform a simple ablation to quantify the effects of including the Sim-to-Real mitigation during training of an MLP retrieval. By multiplying the EOFs by a random scaler drawn from the scaling distributions of those from the OE retrieval we see a 8% reduction in error versus the 2022 TCCON evaluation set.




Figure 3: The top figure shows an example of an observed radiance (blue) and the FM modeled radiance (orange) for a single sounding. The bottom figure shows the mean and variance of remaining error before and after application of the randomly scaled EOFs.

Table 1: Ablation of retrieval with EOF perturbed training radiances and without. RMSE in parts per million [ppm] is calculated for the 2022 TCCON set.

| Retrieval                 | RMSE [ppm]  |
|---------------------------|-------------|
| MLP + EOF                 | 3.15        |
| MLP + Randomly Scaled EOF | <b>3.11</b> |
| MLP Baseline              | 3.37        |

## B.2 THE CONDITIONAL MEAN PRIOR

In Table 2, we evaluate several choices for the conditional mean prior,  $x_{prior}$ , in the forward and reverse diffusion processes. First, we follow the method outlined in Han et al. (2022), where  $x_{prior} = f(\mathbf{y})_\phi$  and  $f_\phi$  is a pre-trained MLP model. For the second approach we modify the MLP to use 1-dimensional convolution layers with max-pooling. Lastly, we evaluate using the GEOS-FP-IT (Lucchesi, 2013) XCO<sub>2</sub> labels which are also used as priors in the operational OE-based retrieval. We find that the convolutional model offers the best performance, where the other two priors provide a poorer RMSE as compared to the ACOS retrieval (shown in Table 3).

Table 2: Comparison of diffusion retrievals with different conditional mean priors. Mean and Standard Deviation of bias in parts per million [ppm] is calculated for the 2022 TCCON set.

| Cond. Mean Prior | Bias [ppm]                         |
|------------------|------------------------------------|
| MLP              | -0.70 $\pm$ 2.88                   |
| Conv1D           | <b>-0.23 <math>\pm</math> 2.77</b> |
| GEOS-FP-IT       | -1.03 $\pm$ 2.95                   |

## B.3 COMPARISON OF DIFFUSION WITH ALTERNATIVE ML RETRIEVALS

In Table 3, we compare the conditional diffusion retrieval to several other deep learning-based methods that provide an uncertainty of their point estimates. The choice of methods are state of the art in several different paradigms of UQ for deep learning models: ensemble, kernel learning, and hierarchical Bayesian.

Table 3: Mean and Standard Deviation of bias in parts per million [ppm] and Miscalibration Area (Miscal. Area) for ML retrieval calculated for the 2022 TCCON set.

| Method      | Bias [ppm]                         | Miscal. Area |
|-------------|------------------------------------|--------------|
| ACOS        | -0.62 $\pm$ 2.82                   | 0.31         |
| MLP         | -0.8 $\pm$ 3.33                    | N/A          |
| SWAG        | -0.56 $\pm$ 3.14                   | 0.23         |
| Deep Kernel | -1.54 $\pm$ 4.50                   | 0.38         |
| Evidential  | -0.49 $\pm$ 2.80                   | 0.26         |
| Diffusion   | <b>-0.23 <math>\pm</math> 2.72</b> | <b>0.21</b>  |

**SWAG:** Stochastic Weighted Averaging-Gaussian or SWAG is a Bayesian extension of Deep Ensembles (Lakshminarayanan et al., 2017) by approximating a Bayesian Model Average through a learned distribution over multiple Stochastic Gradient Descent iterates (Maddox et al., 2019).

**Evidential Regression:** Deep Evidential Regression (Amini et al., 2020) captures both aleatoric and epistemic sources of uncertainty with a single model by placing hyper-priors over the Gaussian likelihood and employing a Inverse-Gamma distribution over the prior to form a Normal Inverse-Gamma conjugate.

**Deep Kernel Learning:** Deep Kernel Learning (Wilson et al., 2016) offers a highly scalable approach to Gaussian Processes by parameterizing a kernel such as a Matern with a neural network and provides an epistemic uncertainty of its point estimates.

The diffusion model is implemented using MLP architectures with *Conv1D* layers for both  $f_\phi$  and  $\epsilon_\theta$ . We use a  $T = 50$  with a cosine noise schedule (Nichol & Dhariwal, 2021). The remaining parameters for the start and end values for the schedule and EMA parameters are left as suggested in Han et al. (2022). The retrieved XCO<sub>2</sub> from the diffusion model is closer to a mean bias of 0.0 parts-per-million than the operational ACOS retrieval without any finetuning or bias correction with a similar error variance. We also note that both the Evidential and SWAG retrievals provide an improved comparison to the TCCON ground truth than ACOS. These methods do not require a conditional mean prior, as observed in Table 2. which can have a substantial impact on the diffusion-based retrieval.

## C THE OPTIMAL ESTIMATION COST FUNCTION

The ACOS retrieval employs a Radiative Transfer Model forward model (FM) coupled with an instrument model for the OCO-2 spectrometer as  $F$  to produce simulated spectra and partial derivatives for a given  $\mathbf{x}$ . The algorithm iteratively seeks the MAP estimate that minimizes the cost function that represents the sum of squared errors with the observations  $\mathbf{y}$  together with a prior state vector  $\mathbf{x}_a$ , both scaled by their uncertainties expressed as covariance matrices:

$$J(\mathbf{x}) = (\mathbf{y} - F(\mathbf{x}, \mathbf{b}))^T \mathbf{R}^{-1} (\mathbf{y} - F(\mathbf{x}, \mathbf{b})) + (\mathbf{x} - \mathbf{x}_a)^T \mathbf{B}^{-1} (\mathbf{x} - \mathbf{x}_a), \quad (2)$$

where  $\mathbf{R}$  is the covariance matrix of the noise  $\nu \sim \mathcal{N}(0, \mathbf{R})$ , the prior is Gaussian with mean  $\mathbf{x}_a$  and covariance  $\mathbf{B}$ , i.e.  $\mathbf{x} \sim \mathcal{N}(\mathbf{x}_a, \mathbf{B})$ . The minimizer  $\hat{\mathbf{x}}$  of  $J$  is the MAP estimate  $\arg \max_{\mathbf{x}} \exp(-\frac{1}{2}J(\mathbf{x}))$ , and under the assumption of Gaussianity the posterior uncertainty has an analytic form, namely,  $\left( \left( \frac{d}{dx} F(\hat{\mathbf{x}}, \mathbf{b}) \right)^T \mathbf{R}^{-1} \left( \frac{d}{dx} F(\hat{\mathbf{x}}, \mathbf{b}) \right) + \mathbf{B}^{-1} \right)^{-1}$  (Rodgers, 2000).