

Dear Reviewer 1,

We would like to offer our sincere thanks for spending your time in the review of our work and helping us to significantly improve the quality and clarity of the manuscript with your very detailed comments and suggestions. Each of your comments is listed below in black text, followed by our response and edits in blue text.

Review 1 Comments:

Overview: In Casey et al., the authors investigated the performance of calibration models developed for ambient O₃ and CO₂ across time and space using field deployments spanning 2014 – 2017 as case studies. Specifically, they looked at the impact of post-deployment calibration vs pre- and post- calibration, and the impact of applying a calibration model developed in one location on U-Pods deployed in other locations. Calibration models investigated included linear models and artificial neural networks. The size and scope of the study is impressive, and I believe there is a significant quantity of insightful information within this paper.

However, in general, I found the narrative of the paper to be confusing (it is hard to effectively distill such a breadth of research) and the take home points could be made considerably clearer.

Additionally, I think this paper would benefit with a few more analyses of general model performance implications and a closer look at the impact of relative humidity. Following these corrections to comments identified below, I believe the publication is suitable to be published in Atmospheric Measurement Techniques.

Response: We have carefully addressed each of the comments below and carried out the analyses suggested by the reviewer to investigate the implications and impact of relative humidity and sensor drift in time. We have also worked to significantly improve and clarify the narrative of the paper as well as clarified and outlined take home points.

GENERAL COMMENTS

Comment: In general, I found this manuscript a little hard to read, because I felt like a cohesive narrative was missing. A lot of information is presented in a rapid-fire manner such that the results and discussion section reads more like a results section, with limited discussion. Although there is no straight forward solution to this problem, I would suggest that the authors think about the three to five key messages they wish to convey in the manuscript and that they tune and streamline the text to support this narrative.

Response: Thanks very much to the Reviewer for helping point out a way for us to improve a cohesive and clear narrative in this work.

Edits: We have added the following five key points to the conclusion as well as supporting edits in the abstract and throughout the results and discussion section:

“The following findings from this work, and associated recommendations, are made to help inform the logistics of future studies that employ field calibration methods of low-cost gas sensors.

1. **Finding:** For O₃ models, LMs perform better than ANNs when the chemical composition of local emissions sources is significantly different in the model-training location relative to the model-application location. We found that when models were trained in an urban area with significant mobile sources, then tested in a peri-urban area, more strongly influenced by oil and gas emissions, the differences in local sources of pollution were significantly different enough that LMs outperformed ANNs. Alternatively, when models were trained in one oil and gas production region and tested in another the different composition of local emissions (lighter vs. heavier hydrocarbons) was not significant enough for LM performance to surpass the performance of ANNs, though some positive bias was evident in predicted O₃ mole fractions.

Explanation: ANNs are very effective at compensating for the influence of interfering gas species through pattern recognition of a training dataset. However, if different patterns, in terms of the relative abundance of various oxidizing and reducing compounds in the air, are present in the testing location relative to the training location, ANNs may not be able to compensate for the influence of interfering gas species as effectively. The relative abundance of interfering oxidizing and reducing compounds are not included as model parameters, but ANN performance is challenged by these circumstances.

Recommendation: When measuring O₃ or other gas species with a metal oxide type sensor, if the nature of dominant emissions sources at the model training location is significantly different than the nature of dominant emissions sources in the model application location, use an LM instead of an ANN. For the best performance, try to train models in locations with similar emissions sources to a desired sampling location. If the nature of dominant emissions sources at the model training and application locations are similar, signals from an array of multiple unique metal oxide sensors will likely augment model performance.

2. **Finding:** For CO₂ models, LMs perform better than ANNs when model training occurs significantly (more than several months) prior to or subsequent to the model application period.

Explanation: CO₂ sensors drift over time in terms of sensitivity and baseline response. When models are extrapolated in time (when training takes place more than several months prior or subsequent to the model application period), ANN performance can be compromised to a greater extent than LM performance because ANNs are able to represent relationships during training very effectively, and with significant more complexity and nonlinear relationships among time and other model inputs than LMs. The more complex the model, the less likely it can be extrapolated effectively. LMs, with no interaction terms like we employ in this work, are not able to fit data and potentially complex patterns inherent in sensor drift over time during training as closely as an ANN, but the simple linear relationships they represent between the time input and the target gas mole fraction over the course of training are more likely to hold prior or subsequent to the training period.

Recommendation: When measuring CO₂ with a NDIR sensor, if model-training data is only available more than several months prior or subsequent to the model application period, use a LM instead of an ANN. For the best model performance, use training data that is collected directly pre or post of the model application period, and preferably data from both pre and post of the model application period. Training models using data from both pre and post of a given model application period helps models to encompass sensor drift over time as well as increases the likelihood of covering the full range of environmental parameter space that occurs during the model application period so that extrapolation of these parameters is avoided.

3. **Finding:** Extrapolation of an O₃ or CO₂ model in time, and especially significant extrapolation in time, can change both the type of model that is most effective, as well as the specific model input signals that are most effective.

Explanation: Low-cost sensors change over time, both in terms of their baseline response and in terms of their sensitivity to target and interfering gas species. Different sensor types drift due to different physical phenomena so further a generalization across sensor types is difficult.

Recommendation: Use training data collected directly pre and post of the model application period in order to implement a 'best performing model' for each gas species that can be applied using data from different model training and application pairs.

4. **Finding:** ANNs yield less bias and more accurate gas mole fraction quantification than LMs, even when transferred to a new location under the following circumstances: when extrapolation of training parameters is avoided during the model application period, when training takes place for several weeks to a month prior and subsequent to the model

application period, and when the dominant local emissions sources are similar in the model training and application locations.

Explanation: Our previous study and multiple other ambient and laboratory based experiments have shown, arrays of low-cost sensors in combination with ANN regression models can support useful quantification of gases in mixtures and in the ambient environment because ANNs can more effectively represent complex nonlinear relationships among environmental variables and signals in a sensor system like a U-Pod than LMs. With this work, we have explored limitations associated with these methods when challenged in different ways, as we present with a number of case studies.

Recommendation: If minimizing error and bias in measurements of gas mole fractions using low-cost sensors systems is a primary goal, design sensor system training and field deployment experiments so that extrapolation of model training parameters is avoided during the model application period, so that training takes place for several weeks to a month directly prior and directly subsequent to the model application period, and so that the dominant local emissions sources are similar in the model training and application locations. When these conditions are satisfied, ANNs can be robustly implemented, with better performance than LMs.

It is also imperative that sensor users keep in mind the primary importance of minimizing extrapolation of temperature, humidity and sensor signal from model training to application.”

Comment: Many of the figures are needlessly complicated by an overload of case studies, unintelligible sensor signal labels, and colours. If there is any way of summarizing this data more cohesively, it would significantly improve the paper.

Response: Thank you very much for the feedback and helping us to simplify and clarify figures.

Edits: According to the specific comments below, we have split what was previously Figure 8 into two figures (now Figure 10 and 11) in order to simplify the graphics and highlight the content of each and simplified and clarified Figure 9 (now Figure 12). We have added definitions for the sensor inputs in the Figure captions for what were Figures 8 and 9 (now Figures 10, 11, and 12). We have also updated Figure 1 (now Figure 2) in order to clarify model training and test periods for each case study, as well as how many U-Pods were included in each case study.

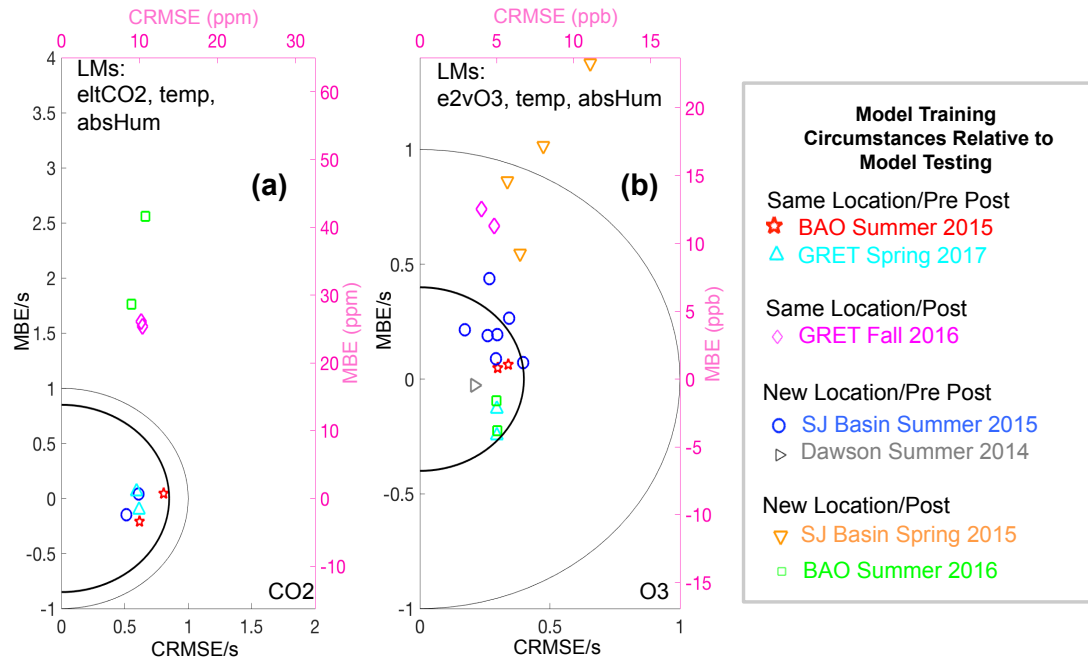


Figure 10: Target diagrams demonstrating performance of a previously determined best-performing model across all new test datasets. (a) CO₂ and (b) O₃ LM performance when only the primary gas sensor, temperature and humidity are inputs. (c) CO₂ and (d) O₃ ANN performance with inputs that were found to perform best at the GRET site in the spring of 2017 (Casey et al., 2017). Model input definitions: eltCO₂ (ELT S300 CO₂ sensor), e2vO₃ (e2v MiCs-2611 sensor), temp (temperature), and absHum (absolute humidity).

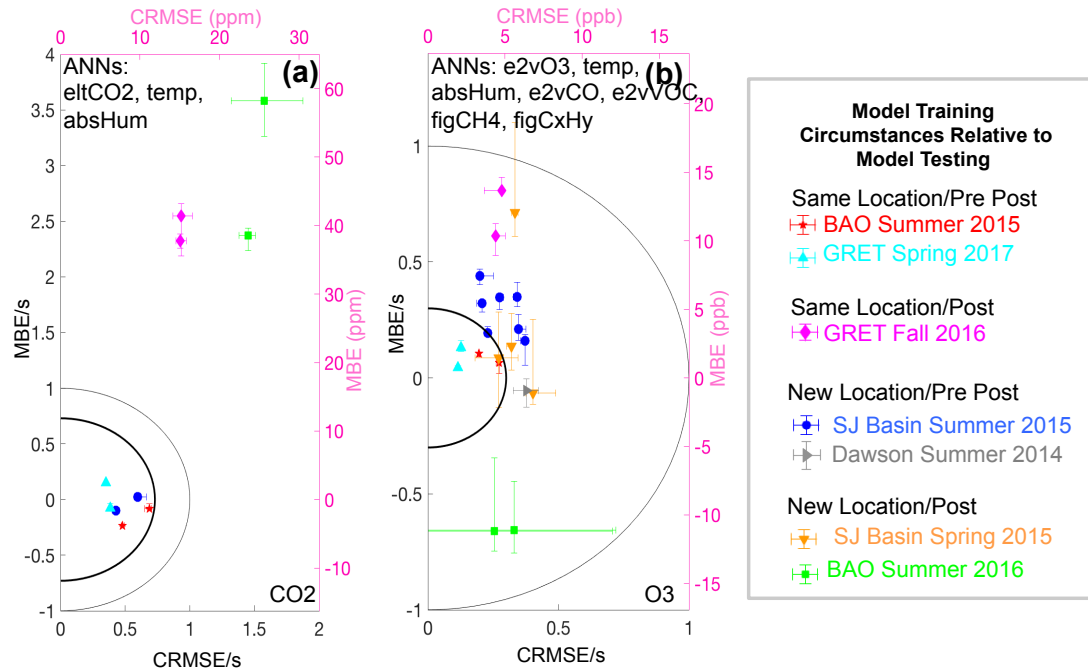
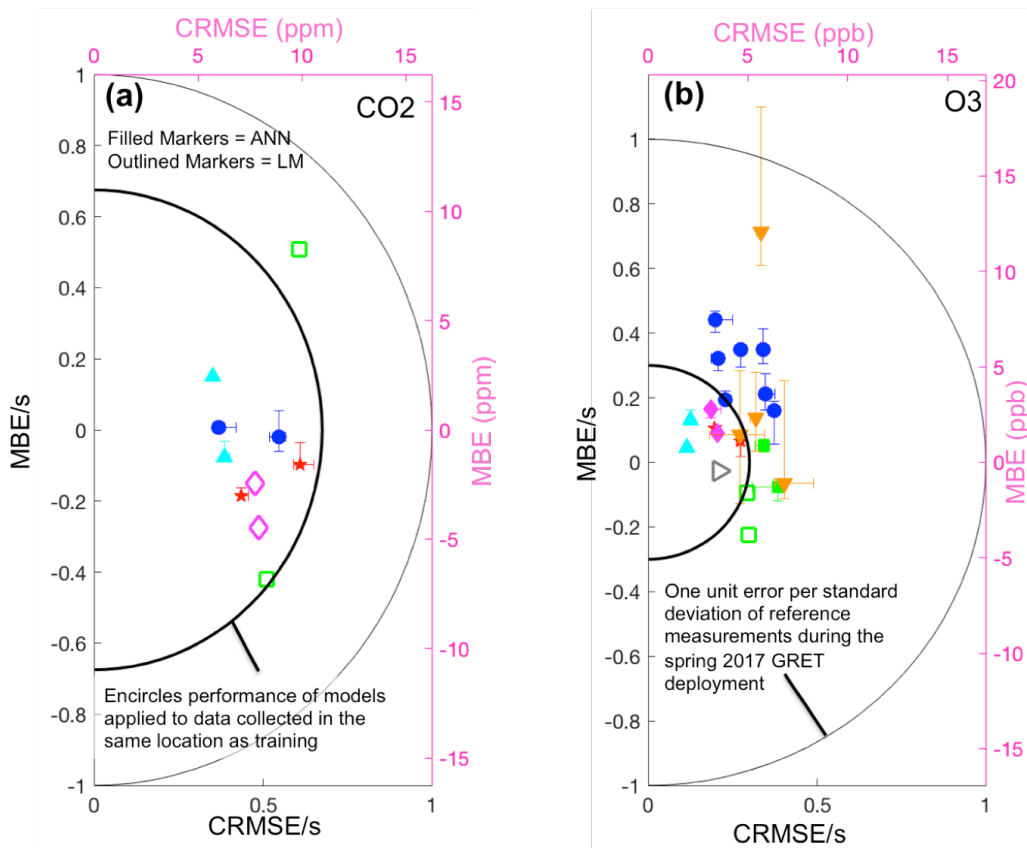


Figure 11: Target diagrams demonstrating performance of a previously determined best-performing model across all new test datasets (a) CO₂ and (b) O₃ ANN performance with inputs that were found to perform best at the GRET site in the spring of 2017 (Casey et al., 2017). Model input definitions: eltCO₂ (ELT S300 CO₂ sensor), e2vCO (e2v MiCs-5525 sensor), e2vVOC (e2v MiCs-5521 sensor), e2vO₃ (e2v MiCs-2611 sensor), figCH₄ (Figaro TGS 2600 sensor), figCxHy (Figaro TGS 2602 sensor), temp (temperature), and absHum (absolute humidity).



Case Study	CO2 Model Inputs	O3 Model Inputs	Description
BAO Summer 2015	eltCO2 temp rh e2vVOC alphaCO	e3vO3 temp absHum e2vVOC e2vCO figCH4 figCxHy	Same Location/Pre Post
SJ Basin Summer 2015	eltCO2 temp rh e2vVOC alphaCO	e3vO3 temp absHum e2vVOC e2vCO figCH4 figCxHy	New Location/Pre Post
SJ Basin Spring 2015		e3vO3 temp absHum e2vVOC e2vCO figCH4 figCxHy	New Location/Post
GRET Spring 2017	eltCO2 temp absHum	e3vO3 temp absHum e2vVOC e2vCO figCH4 figCxHy	Same Location/Pre Post
BAO Summer 2016	eltCO2 temp absHum time	e3vO3 temp absHum	New Location/Post
GRET Fall 2016	eltCO2 temp absHum rh e2vVOC alphaCO figCH4 figCxHy e2vCO time	e3vO3 temp e2vVOC e2vCO figCH4 figCxHy	Same Location/Post
Dawson Summer 2014		e3vO3 temp absHum	New Location/Pre Post

Figure 12: Target diagrams for (a) CO₂ and (b) O₃ calibration model performance for the best performing model for each particular case when tested on data from a number of field deployments. Model input definitions: eltCO2 (ELT S300 CO2 sensor), e2vCO (e2v MiCs-5525 sensor), e2vVOC (e2v MiCs-5521 sensor), e2vO3 (e2v MiCs-2611 sensor), figCH4 (Figaro TGS 2600 sensor), figCxHy (Figaro TGS 2602 sensor), alphaCO (Alphasense CO-B4 sensor) temp (temperature), absHum (absolute humidity), rh (relative humidity), and time (absolute time).

Comment: It would be good if the authors could elaborate on which U-Pods were where over all these campaigns. Given that temporal degradation / time was investigated in detail in this paper, some assessment of UPod changes over the three years of campaigns would be helpful if possible.

Response: Thanks very much to the reviewer for this helpful comment.

Edits: Figure 2 (previously Figure 1) has been updated to clearly state the number of U-Pods included in each case study, for both O₃ and CO₂, as well as names of the specific U-Pods that were used during each case study. We have also performed an assessment of U-Pod sensor drift from the summer of 2015 through the summer of 2017, shown below in Figure s26, that we have added to the Supplemental Materials. Figure s26 had been sited in section 3.2.2 of the manuscript, and described in the following text: "While we did not measure and record metal oxide sensor heater resistance for sensors included in U-Pods, we have investigated eltCO2 and e2vO3 sensor signal drift from the summer of 2015 through the summer of 2017. These data are presented in Fig. S26. Systematic

downward drift in all eltCO₂ sensor signals is apparent over this time frame. A clear and consistent pattern of systematic drift over this time period is less apparent for e2vO₃ sensors. Since the training data was collected immediately after, the test data period, and since the test data period was relatively short (approximately one month) sensor drift could be negligible across the combined training/testing time frame.”

(a) Case Study	Model Training and Test Deployment Timelines												
	Year	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Dawson Summer 2014	2014												
SJ Basin Spring 2015	2015												
SJ Basin Summer 2015	2015												
BAO Summer 2015	2015												
BAO Summer 2016	2016 2017												
GRET Fall 2016	2016 2017												
GRET Spring 2017	2017												

(b) Case Study	Training Location	Test Location	O ₃ # U-Pods	O ₃ U-Pod Names	CO ₂ # U-Pods	CO ₂ U-Pod Names
Dawson Summer 2014	CAMP	Dawson	1	BE	NA	NA
SJ Basin Spring 2015	BAO	SJ Basin	4	BB, BD, BF, BJ	NA	NA
SJ Basin Summer 2015	BAO	SJ Basin	7	BA, BB, BD, BE, BF, BH, BI	2	BB, BD
BAO Summer 2015	BAO	BAO	2	BC, BJ	2	BC, BJ
BAO Summer 2016	GRET	BAO	2	BH, BI	2	BH, BI
GRET Fall 2016	GRET	GRET	2	BH, BI	2	BH, BI
GRET Spring 2017	GRET	GRET	2	BH, BI	2	BF, BI

Figure 2: (a) ANN and LM training and test deployment timelines. The Dawson, BAO, and GRET sampling sites are all located in the DJ Basin. Model training periods for each test deployment are shown in blue, and model test periods are shown in magenta. For the BAO Summer 2016 case study, the period outlined in blue shows data that was used to train O₃ model, but not CO₂ models since CO₂ reference data was not available during winter months. (b) Information about each of the case studies presented in the above timelines, including model training and testing locations, as well as the number and names of U-Pods included in each case study for both O₃ and CO₂ models. The U-Pods with names shown in grey were constructed and deployed starting in May of 2014. The U-Pods with names shown in black were constructed and deployed starting in April of 2015.

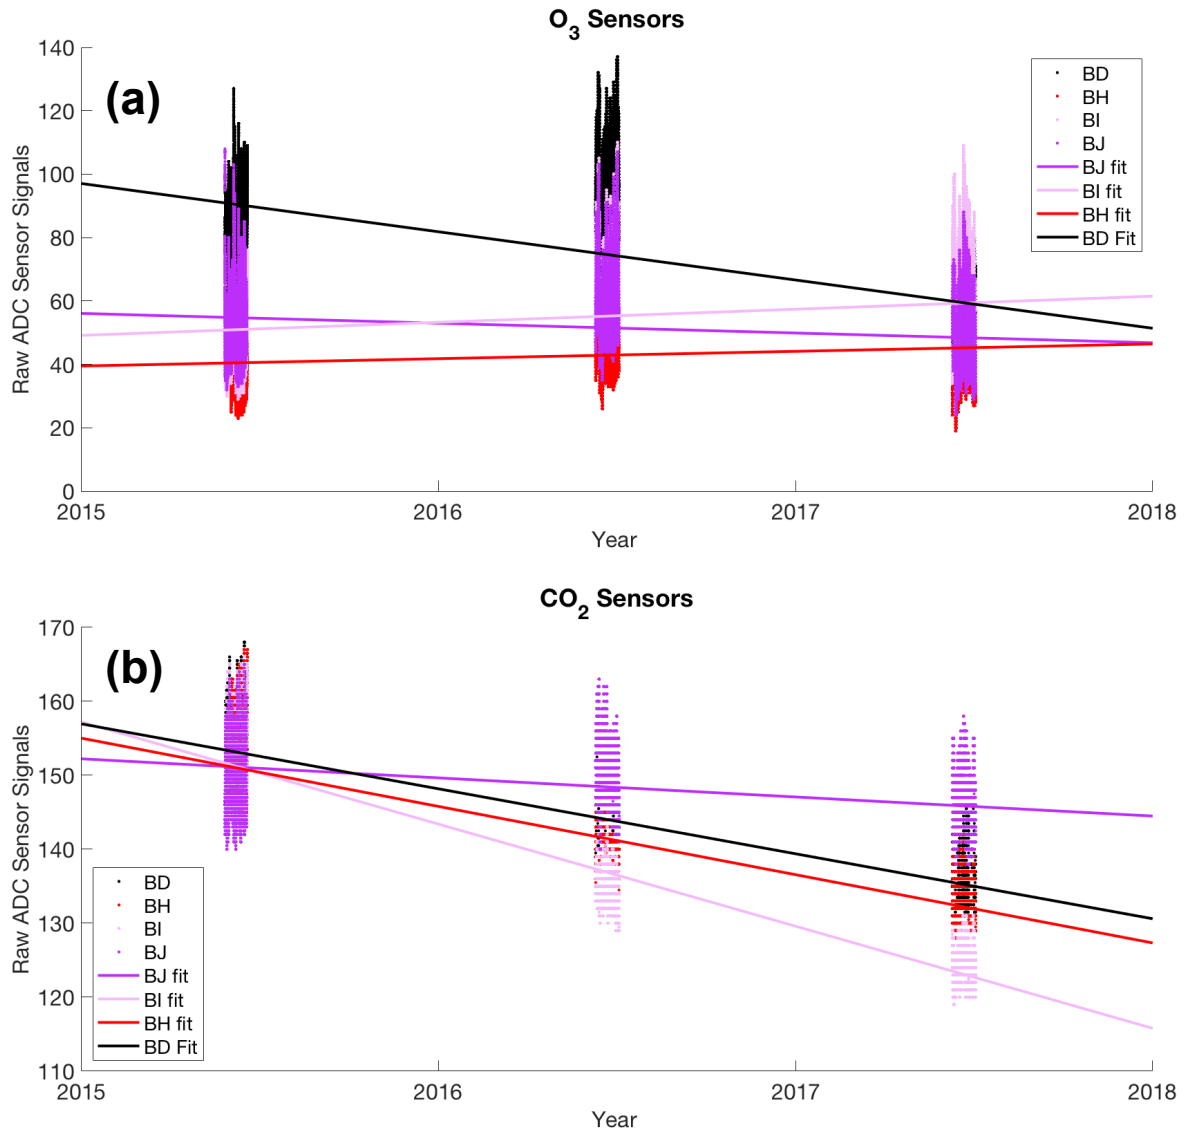


Figure S26 U-Pod sensor drift from 2015 – 2017 for (a) e2v MiCs-2611 O₃ sensors and (b) ELT S300 CO₂ sensors. Data presented are from 23-day periods when U-Pods were co-located together from the summers of 2015, 2016, and 2017. Raw ADC sensor signals were smoothed with rolling hourly medians during these periods, in order to track representative sensor responses across this time period, without the influence of exceptional events. Measurements from summer each year were used to capture sensor response under similar weather conditions.

Comment: Also, when comparing sensor performance spatially, are the U-Pods that are compared the same age?

Response: Thanks to the Reviewer for bringing up this important point.

Edits: We have added the following text accordingly: “Some U-Pods used included in these case studies (indicated in grey font in Fig. 2) were constructed, populated with sensors, and deployed at field sites in the spring of 2014, approximately a year before the rest of the U-Pods were constructed, populated with sensors, and deployed at field sites in the spring of 2015. The relative age of sensor systems included in some case study comparisons could have contributed to some discrepancy in model performance, though systematic differences based on U-Pod age is not apparent.”

Comment: Directly addressing the size of the training and testing windows should be included. It is hard to make generalizable conclusions from the study when there is so much variability in training and testing window size. Is there a reason why some training windows are shorter than others? This should be directly addressed in the manuscript.

Response: Thanks to the Reviewer for the helpful feedback.

Edits: We have added the following explanatory text accordingly: “As available data from each case study allowed, we used approximately one month of training data before and after (pre and post of) a given approximately month-long test period. When training data was not available within several months of a test period, significantly longer training datasets were used in order to attempt capture and effectively represent trends in sensor drift over time, as well as to avoid extrapolation of model parameters (particularly temperature) during the test data period. As a result, model-training durations varied across case studies and sometimes significantly exceeded model-testing durations. Each case study is similar in representing approximately one month-long deployment of sensor systems. This study design serves a primary goal of this work, which is to help support the quantification atmospheric trace gases from low-cost gas sensor data in new locations, relative to model training locations, for periods of approximately one month at a time.”

Comment: I found the discussion of ANN and LM model building to be significantly under-developed, especially considering that this is a measurement techniques journal. This paper relies too heavily on the prior 2017 study, and has too much assumed knowledge that should be summarized in Section 2.4. The resulting LM and coefficients should be provided. As well as some mention of model performance metrics like MAE or r^2 .

Response: Thanks very much to the Reviewer for this helpful feedback. We have developed the discussion of ANN and LM model building significantly, through the addition of the following text in section 2.4 and have added a new subsection 2.5 describing the calibration model evaluation and testing we implemented in this work. We have also added a table summarizing model performance metrics for our previous work, as a case study, among other case studies. Since LM coefficients are unique to individual case studies, and within those groups, unique to gas sensors in individual U-Pods, and since we carried out an analysis of model sensitivity to inputs in our previous work, we have not included LM coefficients in this work.

Edits: “As in [1], direct LMs and ANNs were trained with a number of different sensor input sets to map those inputs to target gas mole fractions measured by reference instruments. Direct LMs implemented were multiple linear regression models given by

$$r = p_1 + p_2 s_1 + p_3 s_2 + \dots + p_n s_{n-1} \quad (1)$$

where r is the target gas mole fraction (measured by a reference instrument) $s_1 - s_{n-1}$ are sensor signals from U-Pods that are included as model predictor variables, and $p_1 - p_n$ are corresponding predictor coefficients. ANNs designed for regression tasks, like those employed in this work, generally consist of artificial neuron nodes that are connected with weights. Weights are initiated with randomly assigned values. An optimization algorithm is then employed to map a given set input values to one or more corresponding target values. An example of a very simple feed forward neural network, and how weights are propagated through it are depicted in Fig. 3. In this work, ANNs were designed by assigning U-Pod sensor signals to artificial neurons in an input layer and assigning target gas mole fractions for an individual gas species, measured by a reference instrument to a single output neuron. Nonlinear, tansig, artificial neurons in one or two hidden layers and a layer of linear neurons were then added between input layer and the network output neuron. Additionally, bias neurons, each assigned a value of 1, were connected to neurons in the hidden layer(s) so that individual connecting weights could be activated or deactivated during the optimization process. The number of neurons in each hidden layer was set equal to the number of inputs included in a given ANN.

For ANN training we employed the Levenberg Marquardt optimization algorithm with Bayesian Regularization [2]. The Levenberg-Marquardt algorithm provides a combination of Gauss-Newton

and Gradient Decent methods, towards incremental minimization of a cost function (the summed squared error between the ANN output and target values as a function of all of the weights in the network). Training begins according to the Gauss-Newton method, in which the Hessian matrix (the second order Taylor series representation of the error surface) is approximated as a function of the Jacobian matrix and its transpose, significantly reducing required training time. Network weights are adjusted accordingly each training step to reduce error. If the cost function is not reduced in a given training step, an algorithm parameter is adjusted so that optimization more closely approximates the gradient decent method (a first order Taylor series representation of the cost function), providing a guarantee of convergence on a cost function minimum. Since local minima may exist across the error surface, it is important to train the same network multiple times (with different randomly assigned starting weights), in order to assess the stability of ANN performance. In this work each ANN was trained 5 times. Fig. 4 shows a diagram of an ANN architecture employed in this work, when there were five inputs.

In the implementation of Bayesian Regularization, a term is added to the sum of squared error cost function as a penalty for increased network complexity in order to guard against over fitting. A two level Bayesian inference framework is employed, operating on the assumptions the noise in the training data is independent, normally distributed, and also that all of the weights in the ANN are small, normally distributed, and unbiased [2]. In preliminary ANN tests we found that over fitting occurred even when Bayesian Regularization was used, so we additionally implemented early stopping, which proved to be effective in the reduction of over fitting. To implement early stopping, a portion of training data is set aside as validation dataset, and during training, an ANN is applied to this validation data after each training step. Training continues so long as the error associated with the validation dataset is reduced. When the error associated with the validation dataset is no longer being reduced, training stops early. For ANNs, training datasets were divided in half on an alternating 24-hr basis, with half used for training and half used as validation data for early stopping. ANNs with two hidden layers were used for CO₂ and ANNs with one hidden layer were used for O₃, in accordance with our earlier findings for each target gas species [1]. Input signals for both LMs and ANNs were normalized so that they ranged in magnitude from -1 to 1 since this practice is recommended for the ANN optimization algorithm used [2].

Calibration Model Evaluation and Testing

LM and ANN performance was evaluated on test datasets. To evaluate the performance of each of the ANN and LM models that were generated using training data then applied to test datasets, we used residuals, the coefficient of determination (r^2), root mean squared error (RMSE), mean bias error (MBE), and centered root mean squared error (CRMSE). The CRMSE is an indicator of the distribution of errors about the mean, or the random component of the error. The MBE, alternatively, is an indicator of the systematic component of the error. The sum of the squares of the CRMSE and the MBE is equal to the square of the total error, the square root of which is defined by the RMSE.

First, we generated and applied the best performing model, as determined in our previous work (presented in Table 4), to data from each new case study. Each new case study was selected to challenge models in different ways in order to evaluate the resiliency of the findings from our previous study when challenged by different circumstances.

Next, we generated, applied, and evaluated the performance of a number of LMs and ANNs with different sets of inputs for each case study in order to see which specific model performed the best for each individual case study. The r^2 , RMSE, and MBE for each of these alternative models when applied to test data are presented in the supplemental materials (SM) in Fig. S2 through Fig. S7, along with representative scatter plots and time series comparing the performance LMs and ANNs for a given set of inputs. In Fig. S2 through Fig. S7, the best performing model inputs for each train/test data pair are shaded in purple. The type of model that performed the best (ANN vs. LM) is indicated in the caption of each figure. We discuss both the performance of the previously determined best fitting model (generated using data from the GRET Spring 2017 case study) when applied and

generated to data from new case studies, and the performance of models that were tuned to perform the best for each individual case study. From these comparisons, we draw insight into circumstances that challenge model performance in terms of relative local emissions characteristics, location, and timing between model training and testing pairs.”

“For each of the case studies, we present the performance three groups of models. The first of these are linear models with only the primary gas sensor signal, along with temperature, and absolute humidity as inputs. The next group of models includes those that were found to perform best in our previous work. The third group of models tested for each case study includes models that were optimized specifically for each case study. Tables 5 and 6 show the mean and standard deviation of model performance metrics for each of the case studies presented.”

Table 5: O₃ model performance metrics.

Case Study	N	R ²	RMSE (ppb)	MBE (ppb)	Standard Deviation R ²	Standard Deviation RMSE	Standard Deviation MBE
O₃ Models							
Best O₃ Model (Casey et al., 2017)							
ANN with inputs: e2vO3 temp absHum e2vVOC e2vCO FigCH4 FigCxHy							
Dawson Summer 2014	1	0.83	6.46	-0.91	0.00	0.00	0.00
SJ Basin Spring 2015	4	0.86	7.74	3.69	0.05	3.82	5.78
SJ Basin Summer 2015	7	0.85	7.03	4.89	0.10	1.10	1.73
BAO Summer 2015	2	0.93	4.26	1.45	0.00	0.31	0.07
BAO Summer 2016	2	0.92	12.21	-11.14	0.00	0.31	0.07
GRET Fall 2016	2	0.96	12.87	12.02	0.01	2.30	2.35
GRET Spring 2017	2	0.98	2.59	1.49	0.00	0.69	1.02
Simple Model (Single Gas Sensor)							
LM with inputs: e2vO3 temp absHum							
Dawson Summer 2014	1	0.95	3.59	-0.46	0.00	0.00	0.00
SJ Basin Spring 2015	4	0.83	17.95	16.09	0.06	6.10	5.83
SJ Basin Summer 2015	7	0.86	6.30	3.53	0.06	1.40	2.06
BAO Summer 2015	2	0.87	5.50	0.94	0.00	0.78	1.56
BAO Summer 2016	2	0.89	5.78	-2.71	0.00	0.78	1.56
GRET Fall 2016	2	0.93	12.73	11.92	0.01	0.62	0.88
GRET Spring 2017	2	0.89	6.00	-3.19	0.00	0.73	1.38
Models Optimized For Case Studies							
Dawson Summer 2014	1	0.95	3.59	-0.46	0.00	0.00	0.00

SJ Basin Spring 2015	4	0.86	7.74	3.69	0.05	3.82	5.78
SJ Basin Summer 2015	7	0.85	7.03	4.89	0.10	1.10	1.73
BAO Summer 2015	2	0.93	4.26	1.45	0.02	0.51	1.54
BAO Summer 2016	2	0.87	6.25	-0.20	0.02	0.51	1.54
GRET Fall 2016	2	0.95	3.99	2.14	0.00	0.28	0.89
GRET Spring 2017	2	0.98	2.59	1.49	0.00	0.69	1.02

Table 6: CO₂ model performance metrics.

Case Study	N	R ²	RMSE (ppm)	MBE (ppm)	Standard Deviation R ²	Standard Deviation RMSE	Standard Deviation MBE
CO₂ Models							
Best CO₂ Model from (Casey et al., 2017)							
ANN with inputs: eltCO2 temp absHum							
SJ Basin Summer 2015	2	0.65	8.42	-0.62	0.00	1.81	1.41
BAO Summer 2015	2	0.75	9.98	-2.60	0.05	13.00	13.89
BAO Summer 2016	2	0.69	54.38	48.37	0.05	13.00	13.89
GRET Fall 2016	2	0.74	42.37	39.58	0.02	2.44	2.57
GRET Spring 2017	2	0.83	6.31	0.59	0.03	0.13	2.61
Simple Model (Single Gas Sensor)							
LM with inputs: eltCO2 temp absHum							
SJ Basin Summer 2015	2	0.71	7.84	0.27	0.01	1.43	0.42
BAO Summer 2015	2	0.69	10.62	-1.26	0.06	1.52	10.67
BAO Summer 2016	2	0.73	11.82	0.73	0.06	1.52	10.67
GRET Fall 2016	2	0.82	8.62	-3.46	0.00	0.69	1.45
GRET Spring 2017	2	0.55	9.88	-0.33	0.03	0.29	1.91
Models Optimized For Case Studies							
SJ Basin Summer 2015	2	0.72	7.45	-0.11	0.04	2.06	0.31
BAO Summer 2015	2	0.80	8.85	-2.29	0.10	6.47	7.08
BAO Summer 2016	2	0.73	11.82	0.73	0.06	1.52	10.67
GRET Fall 2016	2	0.82	8.62	-3.46	0.00	0.69	1.45
GRET Spring 2017	2	0.83	6.31	0.59	0.03	0.13	2.61

Comment: It would be good to include an explicit discussion of % reduction in error by using established models vs. “best fit” models. Can we generalize? What is the quantitative impact of using your prior models vs making a new model every time. My interpretation from this paper is that we need a new model for every U-Pod for every deployment – is there any way around this? I feel there is a significantly missed opportunity to be quantitative here. Section 3.3 could be substantially enhanced using some sort of summary figure/table (other than a target diagram) that gives percent change in bias, random error, r2, mae etc. by switching from pre/post to just post, or by switching location. Given that there are many pairs of sensors looking at impact of pre/post vs. just post or impact of location switching, you could show average % change in model fitting statistics as well as confidence intervals or standard deviations to show the spread across the case studies. This might be a helpful way of streamlining the paper.

Response: Thanks very much to the Reviewer for this helpful comment, which will help us be more quantitative as well as clarify and focus the narrative and results we present.

Edits: Accordingly, we have added a table showing the percent change in R2, RMSE, and MBE, when one set of models is used instead of another, as well as the following text: “Table 7 shows the percent change in model performance metrics when one model-training paradigm is used in place of another, highlighting relative benefits associated with the implementation of different models for O₃ and CO₂.”

Table 7: Relative benefits associated with the implementation of different models for O₃ and CO₂.

Case Study	Mean % Increase in R ²	Mean % Decrease in RMSE	Mean % Decrease in MBE	Mean % Increase in R ²	Mean % Decrease in RMSE	Mean % Decrease in MBE
	CO ₂ Models			O ₃ Models		
Benefit of Models Optimized For Case Studies Over The Best Models from (Casey et al., 2017)						
Dawson Summer 2014				14.51	44.42	50.00
SJ Basin Spring 2015				0.00	0.00	0.00
SJ Basin Summer 2015	10.56	11.52	82.60	0.00	0.00	0.00
BAO Summer 2015	5.84	11.27	11.95	0.00	0.00	0.00
BAO Summer 2016	5.72	78.27	98.49	-5.01	48.82	98.19
GRET Fall 2016	11.17	79.66	108.73	-0.54	68.99	82.22
GRET Spring 2017	0.00	0.00	0.00	0.00	0.00	0.00
Benefit of The Best Models from (Casey et al., 2017) Over Simple Linear Models						
Dawson Summer 2014				-12.67	-79.92	-99.99
SJ Basin Spring 2015				3.20	56.88	77.09
SJ Basin Summer 2015	-8.41	-7.29	331.39	-1.34	-11.53	-38.41
BAO Summer 2015	8.70	6.05	-106.48	6.79	22.48	-53.85
BAO Summer 2016	-5.41	-360.09	-6543.84	2.57	-111.22	-310.71
GRET Fall 2016	-10.05	-391.73	1244.99	2.88	-1.12	-0.86
GRET Spring 2017	51.92	36.13	278.55	10.00	56.90	146.65
Benefit of Models Optimized For Case Studies Over Simple Linear Models						
Dawson Summer 2014				0.00	0.00	0.00
SJ Basin Spring 2015				3.20	56.88	77.09
SJ Basin Summer 2015	1.26	5.06	140.25	-1.34	-11.53	-38.41
BAO Summer 2015	15.04	16.64	-81.80	6.79	22.48	-53.85
BAO Summer 2016	0.00	0.00	0.00	-2.57	-8.10	92.59

GRET Fall 2016	0.00	0.00	0.00	2.33	68.64	82.07
GRET Spring 2017	51.92	36.13	278.55	10.00	56.90	146.65

Comment: This is mentioned in the specific comments, but I would like to see a quantitative assessment of the impact of swapping out RH data if a U-Pod failed. You could accomplish this by taking a U-Pod with valid RH data, replacing it with the Picarro or nearby station RH data, and quantitatively assessing the impact on model performance. That way, you could transition from hypotheticals about the impact of this data swapping to some actual numbers.

Response: Thanks very much to the reviewer for helping us to be less hypothetical about the impact of this data swapping. We have carried out a dummy experiment, testing the effect of this humidity data swapping on data collected during the GRET Spring 2017 case study. A figure, showing the relative performance of models when the humidity data was taken from the U-Pods directly and replaced with measurements from the Picarro CRDS, has been added to the Supplemental Materials. This figure, and associated implication have been cited in the main text.

Edits: “In our previous work, we showed that O₃ models were very sensitive to the humidity signal input (Casey et al., 2017). In this case study, it seems that replacing actual humidity signals with closely approximated humidity signals, negatively influenced model performance. In order to investigate this observation further, we tested the influence of replacing humidity data in the same manner, using mixing ratios from the same co-located Picarro, on test data from the GRET Spring 2017 case study. A comparison of model performance under normal and this ‘borrowed RH’ circumstance are presented in Fig. S27 in the SM. O₃ model performance was negatively impacted when ‘borrowed’ RH values based on Picarro data replaced U-Pod RH sensor signals. From these findings, it seems likely that the inclusion of multiple metal oxide type sensors as inputs in the model, which all respond strongly to humidity fluctuations, helped the ANN to effectively represent the influence of humidity in the system, more so than including a ‘borrowed RH’ signal from another instrument. We tested models with multiple gas sensor signals and no humidity signal as inputs for a number of other case studies as well (as seen in Fig. S2, Fig. S4, and Fig. S5), when good humidity data from U-Pod enclosures was available, but they did not turn out to be the best performing model in any of these other tests.”

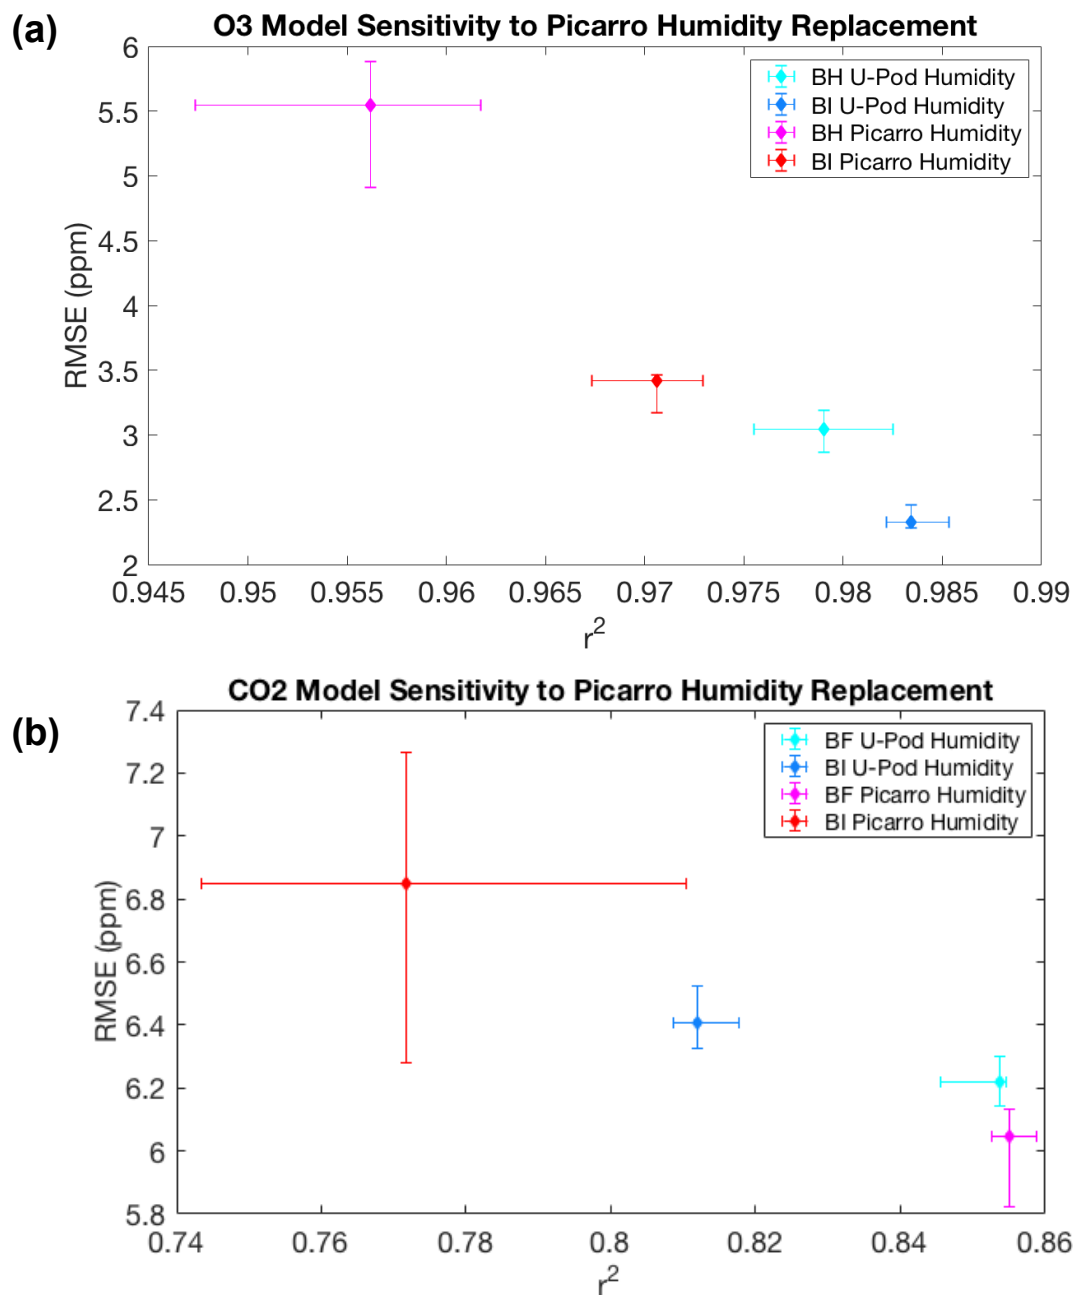


Figure S27 A comparison of model performance when humidity inputs are taken from sensor measurements collected within a given U-Pod sensor system enclosure, vs the performance of models when humidity inputs are replaced using data from a Picarro CRDS for (a) O₃ (b) CO₂

SPECIFIC COMMENTS

Comment: P1 - L13-14: Seems like an oxymoron to say “Generally” if the circumstances for best model performance are case study specific. Recommendation to remove the word “generally”.

Response: Thank you for the helpful feedback.

Edits: We have removed the word ‘Generally’.

Comment: P3 - L19-24: Please discuss why ozone is elevated near O&G production.

Response: Thank you for helping us clarify why ozone can be elevated near oil and gas production activities.

Edits: We have added the following text to augment this discussion: “NO_x and VOC emissions, including those from oil and gas production activities, react in the atmosphere in the presence of sunlight to form tropospheric O₃.”

“Emissions of industry related air pollutants, including O₃ precursors, NO_x and VOCs are expected to occur on spatially distributed scales, across multiple individual components on individual well pads, transmission lines, transportation routes, and gathering stations that are each distributed throughout production basins (Litovitz et al. 2013; Mitchell et al. 2015; Allen et al. 2013). Spatially distributed networks of low-cost sensors have the potential to better inform spatial variability of air quality than existing Regulatory air quality monitoring stations which feasibly cover such spatially resolved measurements continuously, and may not be representative of air quality across smaller spatial scales (Bart et al., 2014; Jiao et al., 2016; Moltchanov et al., 2015).”

Comment: P3 – L27: Can you quantify “small spatial scales” in this context? Is well pad combustion and diesel traffic really contributing so much that it is universally increasing ozone? Most of the construction traffic would occur during active drilling and less so during production when well pad sites are very quiet. I think some further thinking or elaboration on this train of thought it warranted.

Response: Thank you for the helpful comment.

Edits: We have added the following detail, regarding spatial scales that ozone may be influenced near oil and gas emissions sources: “ a modeling study concluded that oil and gas production activities could significantly impact ozone near emissions sources, beginning 2 and 8 km downwind of compressor engine and flaring activities, respectively [3].”

We have also added the following text to address how emissions may change across the lifetime of a given oil and gas production well: “While emissions from truck traffic (and in some cases drilling rig generators), at a given well pad are highest during the drilling, stimulation, and completion phases, industry truck traffic often persists as produced water and condensate tanks are collected from storage tanks on a well pad throughout the life of the well, as do emissions from flaring and compressor engines.”

Comment: P4 – L1-2: What do you mean by “pooling” of compounds - I am not sure I understand this sentence.

Response: Thank you for helping us clarify this statement.

Edits: We have made the following edits accordingly: “While elevated ambient CO₂ levels are not directly harmful to human health, continuous CO₂ measurement can provide information about nearby combustion-related pollution and atmospheric dynamics that lead to the accumulation of potentially harmful compounds associated with the oil and gas production industry during periods of atmospheric stability.”

Comment: P4 – L5-8: I think some short discussion of the operating principles of the sensors would be helpful here.

Response: Thank you for the helpful feedback.

Edits: We have added a discussion of the operating principles of the sensors to section 1.1 accordingly:

“While low-cost sensors have been emerging on the market with sufficient sensitivity to resolve variations in ambient mole fractions of target gases of interest, they are also sensitive to temperature and humidity variations that occur in the ambient environment. NDIR sensors, like the ELT s300 CO₂ sensor employed in this study, have good selectivity, but, since pressure and temperature are not controlled in the optical cavity of ELT s300 CO₂ sensors, the influence of temperature on sensor signals plays an important role. The influence of humidity is also important to address because changes in water vapor are known to influence NDIR measurements of CO₂ in terms of spectral cross-sensitivity due to absorption band broadening (Licor, 2010).

Both metal oxide and electrochemical type sensors operate on the principle of oxidizing or reducing reactions at sensor surfaces. For electrochemical sensors, like the Alphasense CO-B4 sensor

employed in this study, oxidizing or reducing compounds react at the working electrode, resulting in the transfer of ions across an electrolyte solution from the working electrode to the counter electrode, balanced by the flow of electrons across the circuit connecting the working electrode to the counter electrode. A linear relationship is expected between this current and the target gas mole fraction. Electrochemical sensors can be tuned to respond more or less strongly to specific gases by adjusting the materials properties of the working electrode. A membrane is located between the working electrode and the exterior of the sensor in order to control redox reaction rates. Gases diffusion through the membrane to reach the working electrode and the electron transfer rates have been shown to increase at higher temperatures (Xiong and Compton, 2014), and since chemical reaction rates are also influenced by temperature, electrochemical sensor responses can be influenced by sensor operating temperature. Changes in ambient humidity levels can cause sensors to loose or gain of the electrolyte solution, by mass, also influencing electrochemical sensor response (Xiong and Compton, 2014).

For metal oxide sensors, and to a lesser extent for electrochemical sensors, resolving the response of a sensor attributable to the target gas species can also pose a challenge in the presence of interfering gas species. Metal oxide sensors, like those used in this study, have a resistive heater circuit that warms up the sensor surface, causing O₂ molecules to adsorb to the sensor surface, which leads to increased resistance across the surface of the sensor. In the presence of an oxidizing compound, like O₃, more oxygen molecules are adsorbed to the sensor surface and the resistance across the sensor surface is increased further. In the presence of a reducing compound, like CO, oxygen molecules are removed from the sensor surface, allowing electrons to flow more freely, resulting in decreased resistance across the sensor surface. For metal oxide sensors, the resistance across the sensor surface can then be used to determine the mole fraction of a given oxidizing or reducing compound, often according to a nonlinear relationship. Exposure to humidity has been shown to significantly lower the sensitivity of metal oxide gas sensors making it an important parameter to address in a gas quantification model (Wang et al., 2010). Metal oxide sensor operating temperature has also been shown to strongly influence sensor sensitivity and selectivity to different gas species (Wang et al., 2010). Metal oxide type sensors can be tuned to respond differently from one another to oxidizing and reducing gas species by using different metal oxide materials and doping agents for the sensor surface, but selectivity is difficult to achieve.”

We have also added a section to the introduction, section 1.2, entitled “Low-Cost Air Quality Sensor Quantification:

“Because low-cost gas sensor signals are influenced, sometimes significantly, by interfering gas species and changing weather conditions in the ambient environment, field normalization methods to quantify atmospheric trace gases using low-cost sensors have been found to be more effective than lab calibration (Cross et al., 2017; Piedrahita et al., 2014; Sun et al., 2016). Our previous study and several others have compared the efficacy field calibration models generated using LMs (simple and multiple linear regression) relative to supervised learning methods (including ANNs and random forests), all finding that ANNs (Casey et al., 2017; Spinelle et al., 2015, 2017) and random forests (Zimmerman et al., 2017) outperformed LMs in the ambient field calibration of low-cost sensors. Like earlier laboratory based studies (Brudzewski, 1999; Gulbag and Temurtas, 2006; Huyberechts and Szeco, 1997; Martín et al., 2001; Niebling, 1994; Niebling and Schlachter, 1995; Penza and Cassano, 2003; Reza Nadafi et al., 2010; Srivastava, 2003; Sundgren et al., 1991), ANN-based calibration models, incorporating signals from an array of gas sensors with overlapping sensitivity as inputs, have been able to effectively compensate for the influence of interfering gas species and resolve the target gas mole fraction.

ANNs are known to be able to very effectively represent complex, nonlinear, and collinear relationships among input and output variables in a system (Larasati et al., 2011). ANNs are useful in the field calibration of low-cost sensors because, through pattern recognition of a training dataset, they are able to effectively represent the complex processes and relationships among sensors and the ambient environment that would be very challenging to represent analytically or based on empirical

representation of individual driving relationships. In practice though, the reason multiple gas sensors are able to improve the performance of calibration models may be in part the result of correlation between mole fractions of target gases themselves that hold for one model training location, but might not remain effective at alternative sampling sites or during other time periods.”

Comment: P4 – L25: Not sure what is meant by “toward” here

Response: Thank you for the helpful feedback.

Edits: We have replaced “toward” with “that were used for”

Comment: P5 – L27: Is this “clean air” normalization done dynamically/in real-time in parallel with the actual measurement? Or is the clean air measurement established during some calibration/maintenance? Please clarify.

Response: Thank you for helping us clarify.

Edits: We have added the following text accordingly: “For metal oxide type sensors, voltage signals were converted into resistance, and then normalized by the resistance of the sensor in clean air, R_0 . A single value for R_0 was used for each sensor across the study duration. This R_0 value was taken as the resistance of each sensor at the GRET field deployment site when the target pollutant had approached background levels (at night for the metal oxide O_3 sensors and midday for all other metal oxide sensors), and when the ambient temperature was approximately 20° C and relative humidity of approximately 25%.”

Comment: P5 – L30-32: Is there expected to be spatial variability of RH?

Response: Thank you for helping us to clarify.

Edits: We have added the following text to section 2.3: “The closest U-Pod with good humidity sensors ranged from several feet, when U-Pods were co-located during deployments in the DJ Basin, to approximately fifty miles during deployments in the San Juan Basin.”

In Section 3.1 of we have added this text also: “Since the Ignacio site was located approximately twenty-two and fifty miles away from the Navajo Dam and Sub Station sites respectively, this could have introduced some additional error into the application of a calibration equation, particularly since we showed earlier that O_3 ANNs like the ones we employed here are very sensitive to humidity inputs (Casey et al., 2017). Spatial variability in humidity across tens of miles could be significant as isolated storms (which are on average 15 miles in diameter) propagate throughout the region in the summer.”

Comment: P5 – L30-32: Why not just replace the RH sensors directly?

Response: Good question.

Edits: In answer, we have added the following text: “RH sensors were not replaced during field deployments in order to preserve consistency across different deployment periods, allowing for the possibility of a single comprehensive model to apply to all data from a single U-Pod. After some experimentation in generating a ‘master model’ that could be applied to data from a given U-Pod for all collected field measurements, across several years, we determined that individual models for each deployment would be more effective, and replacing RH sensors that had drifted down would have been appropriate in support of the methods presented here. We have since upgraded to Sensirion AG SHT25 sensors, which appear to be more robust and consistent over the course of long-term field deployments.”

Comment: P7 – Section 3.0 first paragraph – Are there some general conclusions from the SM that you can discuss here? Some discussion of model performance is warranted vs just describing what figures are in the SM.

Response: Thanks very much for the feedback.

Edits: We have moved the paragraph in question to the methods section and have added the following sentence, letting the reader know that these plots are discussed in the results and

discussion section in context with each case study presented: “The best-performing model for each case study are highlighted below in the Results and Discussion section.”

Comment: P8 – L17: What is eltCO2?? Can you better define all the model parameter inputs? This comes up in Figure 9 as well.

Response: Yes, thank you for the feedback.

Edits: Description added here and at the first mention of other model input codes in the manuscript in the text: “eltCO2 (ELT S300 CO2 sensor)”

We have also defined these model input codes in the caption for Figure 9 (now Figure 12) as well as Figure 8 (now Figure 10 and Figure 11): “Model input definitions: eltCO2 (ELT S300 CO2 sensor), e2vCO (e2v MiCs-5525 sensor), e2vVOC (e2v MiCs-5521 sensor), e2vO3 (e2v MiCs-2611 sensor), figCH4 (Figaro TGS 2600 sensor), figCxHy (Figaro TGS 2602 sensor), alphaCO (Alphasense CO-B4 sensor), temp (temperature), absHum (absolute humidity), rh (relative humidity), and time (absolute time).”

Comment: P8 – L30-33: Is this early morning under prediction really true? Bloomfield doesn’t look like it is exhibiting any diurnal variation in residual error at all... I feel like given the small number of U-Pods, it is hard to make this conclusion definitively.

Response: We agree, thank you for pointing the trend in Bloomfield out.

Edits: We have edited the text to say “three of four U-Pods” instead of “all four U-Pods”

Comment: P9 – L4-13: I am confused now – why did you use the model with three inputs (eltCO2, abshum, and temp) if the best performing model had more variables? I feel like the model selection discussion is substantially underdeveloped. There could be many good reasons to not choose a more complex model, but any discussion of this seems to be completely omitted, or the reasoning is too difficult to follow.

Response: Thank you for the feedback and for helping us to clarify. You have helped us see that some important details were missing from the methods section regarding model selection and testing procedures.

Edits: We have added a subsection 2.5 to the end of the methods section entitled “Calibration Model Evaluation and Testing”. In this section, we first define the r2, RMSE, MBE, and CRMSE metrics that are used to evaluate the performance of a given model when it is applied to a test dataset. Next we added a paragraph describing how we first tested models that were found to perform best for each gas species in our previous work, and then evaluated the performance of the best model for each specific case study. We then describe the methodology behind model selection and testing for each case study, in the following text and in the newly added Table 4:

“First, we generated and applied the best performing model, as determined in our previous work (presented in Table 3), to data from each new case study. Each new case study was selected to challenge models in different ways in order to evaluate the resiliency of the findings from our previous study when challenged by different circumstances.

Table 3: Best performing models, as determined for each gas species, in the previous study (Casey et al., 2017)

Gas Species	Model Type	Sensor Signal Model Inputs	
CO ₂	ANN	eltCO2 temp absHum	(ELT S300 CO2 Sensor) (temperature) (absolute humidity)
O ₃	ANN	e2vO3 e2vCO e2vVOC figCH4 figCxHy	(e2v MiCs-2611) (e2v MiCs-5525) (e2v MiCs-5521) (Figaro TGS 2600) (Figaro TGS 2602)

		temp	(temperature)
		absHum	(absolute humidity)

Next we tested LMs for CO₂ and O₃ that contained only the primary target gas sensor for each species, as well as temperature and absolute humidity as inputs. Finally, we generated, applied, and evaluated the performance of a number of LMs and ANNs with different sets of inputs for each case study in order to see which specific model performed the best for each individual case study. The r^2 , RMSE, and MBE for each of these alternative models when applied to test data are presented in the supplemental materials (SM) in Fig. S2 through Fig. S7, along with representative scatter plots and time series comparing the performance LMs and ANNs for a given set of inputs. In Fig. S2 through Fig. S7, the best performing model inputs for each train/test data pair are shaded in purple. The type of model that performed the best (ANN vs. LM) is indicated in the caption of each figure. We discuss both the performance of the previously determined best fitting model (generated using data from the GRET Spring 2017 case study) when applied and generated to data from new case studies, and the performance of models that were tuned to perform the best for each individual case study. From these comparisons, we draw insight into circumstances that challenge model performance in terms of relative local emissions characteristics, location, and timing between model training and testing pairs. Table 4 lists the relative timing and parameter coverage between model training and testing periods for dataset pairs, highlighting instances of incomplete coverage during training that led to model extrapolation during testing.”

Comment: P9 – L28-30: It would be good to do a more comprehensive assessment of the impact of replacing RH sensor signal on model performance. Could you conduct a dummy experiment where you replace RH data you actually logged with that of a nearby or alternate monitor and then quantify the impact on model outcome? Given that it seems that a) RH/abshum is an important variable and b) that you had significant data loss issues, I feel a more quantitative assessment of the impact of these data substitutions is needed.

Response: Thanks very much to the reviewer for helping us to be less hypothetical about the impact of this data swapping. We have carried out a dummy experiment, testing the effect of this humidity data swapping on data collected during the GRET Spring 2017 case study. A figure, showing the relative performance of models when the humidity data was taken from the U-Pods directly and replaced with measurements from the Picarro CRDS, has been added to the Supplemental Materials. This figure, and associated implication have been cited in the main text.

Edits: “The fall 2016 GRET test period coincided with the time period U-Pod absolute humidity was replaced using mixing ratios from a co-located Picarro due to missing humidity sensor data. Interestingly, when this ‘borrowed’ humidity signal was not included as an input, the model performance markedly increased and became competitive with other ‘same location’ test deployment case studies. In our previous work, we showed that O₃ models were very sensitive to the humidity signal input (Casey et al., 2017). In this case study, it seems that replacing actual humidity signals with closely approximated humidity signals, negatively influenced model performance. In order to investigate this observation further, we tested the influence of replacing humidity data in the same manner, using mixing ratios from the same co-located Picarro, on test data from the GRET Spring 2017 case study. A comparison of model performance under normal and this ‘borrowed RH’ circumstance are presented in Fig. S27 in the SM. O₃ model performance was negatively impacted when ‘borrowed’ RH values based on Picarro data replaced U-Pod RH sensor signals. From these findings, it seems likely that the inclusion of multiple metal oxide type sensors as inputs in the model, which all respond strongly to humidity fluctuations, helped the ANN to effectively represent the influence of humidity in the system, more so than including a ‘borrowed RH’ signal from another instrument. We tested models with multiple gas sensor signals and no humidity signal as inputs for a number of other case studies as well (as seen in Fig. S2, Fig. S4, and Fig. S5), when good humidity data from U-Pod enclosures was available, but they did not turn out to be the best performing model in any of these other tests.”

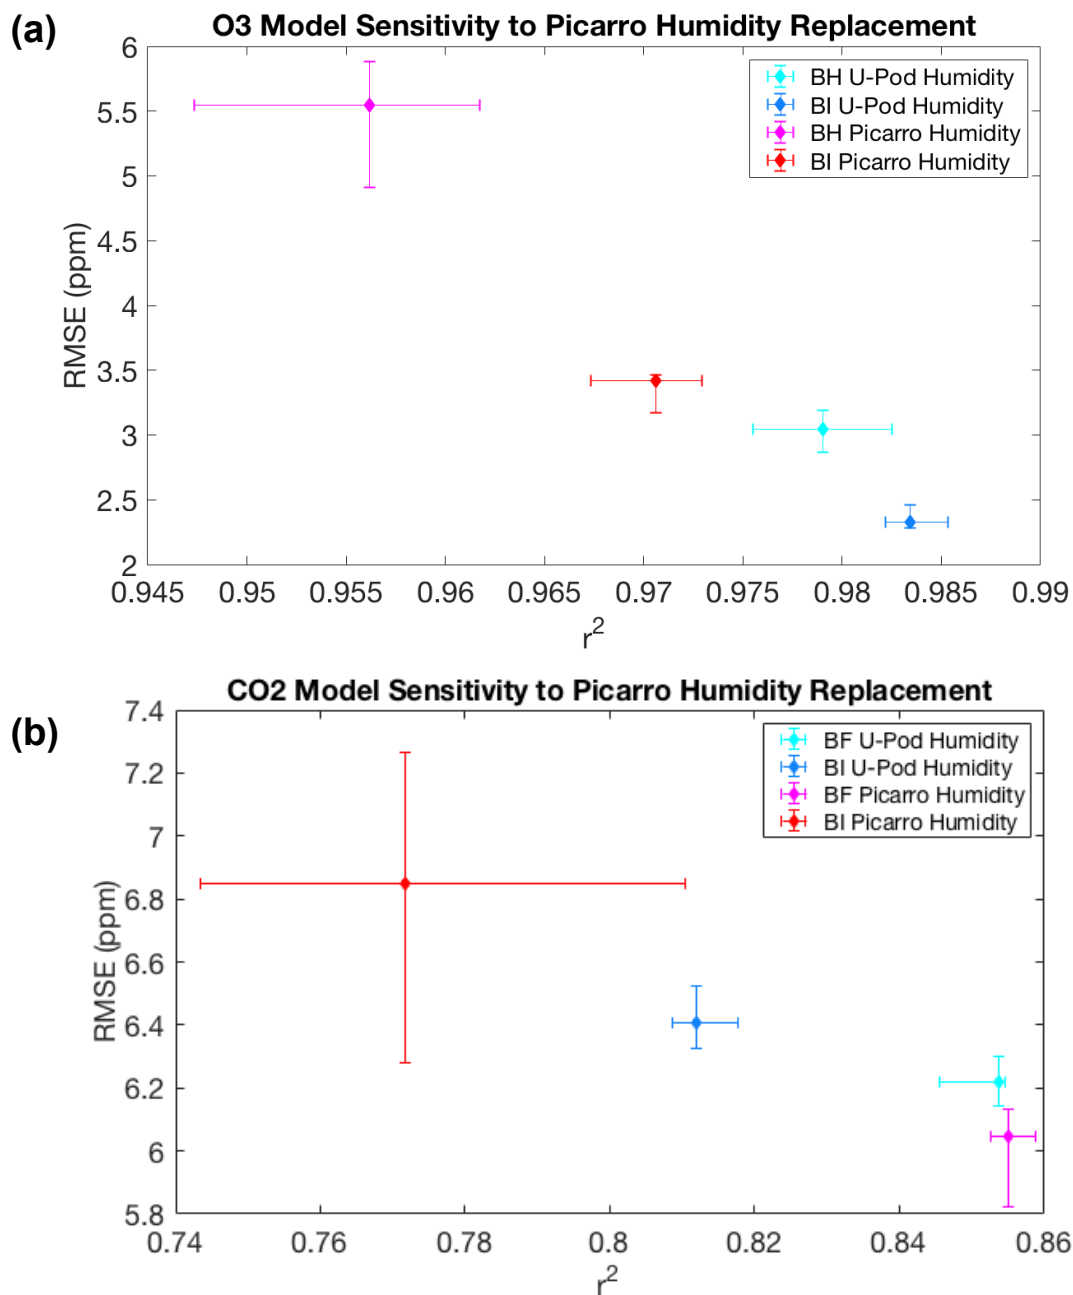


Figure S27 A comparison of model performance when humidity inputs are taken from sensor measurements collected within a given U-Pod sensor system enclosure, vs the performance of models when humidity inputs are replaced using data from a Picarro CRDS for (a) O₃ (b) CO₂

Comment: P11 – L2-13: Can you comment on the quality of the fit at Dawson vs CAMP in addition to the ideal model. The discussion is fairly qualitative. Also, the LM should be much better at extrapolating vs the ANN (which cannot extrapolate I think...not sure) – can you comment on this difference? Does LM perform better because it can extrapolate?

Response: Thank you for this useful comment. We agree it is true that LMs should accommodate extrapolation more effectively than ANNs, and that a quantitative description of model performance is warranted.

Edits: We have added the following text, accordingly: “The fact that LMs performed better than ANNs in this case (with an r^2 of .95 and RMSE of 0.35 ppb for LMs, as opposed to an r^2 of .9 and an RMSE of 5.1 ppb for ANNs) may have to do with the general expectation that LMs be more resilient to extrapolation than ANNs.”

Comment: P11 – Section 3.2.2: If the calibration is immediately after deployment, I am not surprised that there wasn’t much of an effect. Do you anticipate there should be a significant time effect on such short time scales? What is the lifespan of the U-Pods?

Response: Thank you for the useful comment.

Edits: We have added the following text accordingly: “Gas sensor manufactures don’t clearly define sensor lifetimes, but sensors are generally expected to lose sensitivity over time. For example, Alphasense CO-B4 electrochemical sensors are expected to have 50% of their original sensitivity after two years (Alphasense, 2015). The heater resistance in a given metal oxide type sensor is expected to drift over time, influencing sensor measurements (e2v Technologies Ltd., 2007). Masson and colleagues observed a significant drift in a metal oxide sensor heater resistance over the course of a 250 day sampling period in a laboratory setting (Masson et al., 2015). While we did not measure and record metal oxide sensor heater resistance for sensors included in U-Pods, we have investigated eltCO_2 and e2vO_3 sensor signal drift from the summer of 2015 through the summer of 2017. These data are presented in Fig. S26. Systematic downward drift in all eltCO_2 sensor signals is apparent over this time frame. A clear and consistent pattern of systematic drift over this time period is less apparent for e2vO_3 sensors. Since the training data was collected immediately after, the test data period, and since the test data period was relatively short (approximately one month) sensor drift could be negligible across the combined training/testing time frame.”

Comment: P11 – L25: I am confused by the introduction of discussion around figC_xH_y – should we expect this sensor to play an important role?

Response: Thank very much to the reviewer for helping us to clarify.

Edits: We have added the following text: “Again the model for O_3 that was found to perform best in our previous (Casey et al., 2017), an ANN with temp, absHum and all metal oxide sensor signals as inputs, performed the best at sites included in this case study, with one exception. At the Sub Station site the inclusion of the figC_xH_y sensor signal decreased model performance. Additionally, the performance of all models tested at the Sub Station site during the SJ Basin Spring 2015 deployment was significantly worse in terms of MBE than model performance at other sites, both LMs and ANNs with different sets of inputs. Since this sensor signal input augmented model performance at the same sampling location during the summer deployment period, this finding could be attributable to the extrapolation with respect to temperature that occurred during the test period of this case study. As discussed in the introduction, metal oxide sensor sensitivity to different gas species can vary along with sensor surface temperature. Models were trained to use the figC_xH_y sensor signal, across the ambient temperatures encompassed by the training data, to help account for the influence of confounding gas species at the BAO site. We think it is possible that the different temperatures in combination with the unique mix of gas species present at the Sub Station site, which the figC_xH_y sensors are highly sensitive to, caused the ANN to perform worse.”

Comment: P12 – L5: How long is “so long”? This is related to my comment on

Response: Thank you for helping us to be more specific.

Edits: We have changed “so long” to “several months”.

P11 – Section 3.2.2.

Comment: P12 – L21-24: I am confused – did you switch to humidity measured by Picarro or omit humidity entirely? It is not clear to me what happened here.

Response: Thank you for pointing out this confusion.

Edits: We have more clearly described the humidity replacement process (if any) for each individual case study dataset pair in the methods section. We have additionally added to the text in section 2.3 as follows: “Water mole fractions measured by the Picarro were converted into mass-based mixing

ratios to match the units of the absolute humidity signal in the U-Pod data. We applied an adjustment to this absolute humidity signal so that it matched observations in U-Pods during the following month when good RH sensor data was available, to account for the fact that temperatures were higher in U-Pod enclosures than the ambient environment. We then replaced the relative humidity signal in each U-Pod from August 23rd through October 1st in 2016 with the mixing ratios derived from Picarro measurements. Using the temperature and pressure logged in each U-Pod along with the absolute humidity from the Picarro, relative humidity was calculated for each U-Pod during this period.”

And to section 3.2.4: “Interestingly, when humidity this ‘borrowed humidity signal was not included as an input, the model performance markedly increased and became competitive with other ‘same location’ test deployment case studies.”

Comment: P13 – L13: I don’t really understand what is meant by “relative circumstances” – could you be more explicit about each of the case studies? Perhaps a table that outlines case study, with a one sentence description, and a column describing limitations would be more appropriate (and should be introduced at the beginning of the paper).

Response: Thank you for the feedback.

Edits: We have changed “relative circumstances” to “relative timing and parameter coverage”. We have also adapted Table 4 according to this feedback and described the ‘relative circumstances’ present in each case study much more thoroughly in the methods section.

Comment: P13 – L18: What is meant by “extrapolated significantly?” Can you be specific?

Response: Sure, thank you for the comment.

Edits: We have changed “extrapolated significantly” to “extrapolated more than several months”.

Comment: Table 1: I find Table 1 almost impossible to follow. It is not very clear which sensor measures which pollutant, as the input codes are frequently indecipherable. I am honestly not sure what I am supposed to get out of this table.

Response: Thank you for this useful feedback that will help us improve the quality and clarity of Table 1.

Edits: We have added a row at the top of the table indicating the target gases for each sensor. We have added to the Table caption an explanation of the input codes for the sensors: “Gas sensors included in U-Pods along with the model input codes we assigned each. The input code for each gas sensor is simply an abbreviation for the make of the sensor, followed by the target gas species(s).”

Comment: Table 3: I find it difficult to interpret this table. What do the black diamonds mean? What do you mean by “relative circumstances”??

Response: Thank you for the helpful comment.

Edits: We have updated Table 4 by replacing the first column so it is more clearly an indication of which case studies covered which target gases (O₃ and CO₂, or just O₃). We have also updated the caption of Table 4 to be more descriptive and informative, and less confusing: “Relative timing and parameter coverage between model training and test deployment dataset pairs. Incomplete coverage of time occurred if training only took place before or after the test data period and not before and after (pre and post). Incomplete coverage of location occurred when training took place in one location and testing took place in another. Incomplete coverage of parameters, including the target gas mole fraction, temperature, time, and pressure occurred when the values observed during training did not encompass the values observed during testing.”

Comment: Figure 8: There is way too much going on in this figure, it is almost difficult to look at. Is there a more streamlined way of presenting the findings that is less complicated? I feel there is valuable information in the Figure, but it’s hard to determine what that is, due to information overload. Ditto for

Figure 9, though it isn't as bad. Also you would do well to remind the reader what each variable represents.

Response: Thank you very much for the feedback and helping us to simplify and clarify figures.

Edits: We have split what was previously Figure 8 into two figures (now Figure 10 and 11) in order to simplify the graphics and highlight the content of each. We have also added definitions for the sensor inputs in the Figure captions for what were Figures 8 and 9 (now Figures 10, 11, and 12).

TECHNICAL CORRECTIONS

Comment: P1 – L18: “in time than to...” vs. in time that to

Response: Thanks very much for catching this typo.

Edits: We have changed ‘that’ to ‘than’ accordingly.

Comment: P2 – L18: change informal language “hold up” to something more scientific

Response: Thank you for the helpful feedback.

Edits: We have replaced ‘hold up’ with ‘remain effective’.

Comment: P2 – L20: Delete word “Specifically”

Response: Thank you for the helpful edit.

Edits: We have deleted the word “Specifically”

Comment: P7 – L8: Rephrase “...showed successfully reduced over fitting”

Response: Thank you for the helpful comment.

Edits: We have replaced “...showed successfully reduced over fitting” with proved to be effective in the reduction of over fitting”

Comment: Figure 1: Enhance figure caption to explicitly state that blue is training, pink is testing

Response: Thank you, we will do this.

Edits: “Model training periods for each test deployment are shown in blue, and model test periods are shown in magenta.”

Dear Reviewer 2,

We would like to offer our sincere thanks for spending your time in the review of our work and helping us to significantly improve the quality and clarity of the manuscript with your very detailed comments and suggestions. Each of your comments is listed below in black text, followed by our response and edits in blue text.

Review 2 Comments:

Overview: Casey and Hannigan explore the spatial and temporal transferability of field calibration models (specifically linear models (LMs) and artificial neural networks (ANNs)) for two sensors, O₃ (e2vO₃) and CO₂ (eltCO₂), reported by the integrated U-POD sensor package. By 'spatial/temporal transferability' they mean a determination as to whether a calibration model trained from sensor colocation (with reference instrumentation measuring target species) at one location works effectively when that same sensor system is then deployed at a different location. As the authors point out, changing the micro-environment (and local air pollution source contributions to that unique environment) may pose additional complications/challenges when trying to reconcile quantitative measurements with low-cost sensors. The authors make some attempt to separately describe temporal and spatial extension to better understand whether time-alone undermines the accuracy of the calibration models or change of location.

While the topic of sensor calibration and extension of calibration models across a diverse set of deployment scenarios is of fundamental importance to the field of low-cost AQ sensing, the paper, as written, largely fails to pull together a coherent narrative from which active participants in the low-cost AQ measurement space could easily glean useful, actionable information. To be clear, the topic of sensor quantification is inherently complex, and the authors undertake an ambitious analysis spanning 3 years of data from 10 U-POD systems deployed across 4 micro-environments. There are important lessons to be learned from their efforts, but at present these lessons are not brought to the fore of the paper and as a result are easily lost to the reader.

Response: Owing to Reviewer comments and through careful reconstruction, we think the updated version of the paper does a much better job of pulling together a coherent narrative that will be useful for others in the field of low-cost AQ sensing, in terms of useful, actionable information. We hope the Reviewer and the editor will find that important elements of the paper and take-away lessons are now brought to the fore of the paper so that readers can more easily note and make use of our findings.

Edits: We have clarified and added significant detail to the methods section. We have added more context, explanation, and discussion of specific findings in the results and discussions section, and we have explicitly highlighted a number of take away points and recommendations connected to specific findings in the conclusions, as well as highlighted these points in the abstract.

Comment: Throughout the manuscript the authors refer back to their published work (Casey et al., 2017). In the vast majority of instances in which this reference is provided, there is little to no contextual detail explicitly drawing the lines of connectivity between the current work and the previous work. Seeking out the exact evidence that exists in the earlier work and relating its relevance to the current work is left entirely up to the reader. Overall, this referencing needs to be done in a manner that is not vague and does not require that the reader be intimately familiar with the previous work.

Response: Thanks very much to the Reviewer for this helpful comment highlighting confusion between the contributions and citations of our previous work and the unique contributions of the current work.

Edits: In each instance that we have cited our previous work, we have added text to provide contextual detail explicitly drawing the lines of connectivity between the current work and the previous work.

Comment: The paper would also be strengthened if the unique and novel insights that result from the current work were more clearly differentiated from the Casey et al., 2017 effort.

Response: Thanks very much to the Reviewer for this very helpful comment and suggestion.

Edits: We have added text to the conclusions section explicitly summarizing unique and novel insights that result from the current work so that current findings are clearly differentiated from our current work.

Comment: There are seemingly contradictory statements throughout the text. These tend to originate from the authors' desire to provide a clear-cut answer as to whether or not a given model 'worked' in a given case study under a given environmental sampling condition. The fact of the matter is, low-cost AQ sensor quantification is extremely convoluted and often times the validity of data can be somewhat ambiguous. Faced with this level of complexity, the current manuscript fails to provide a succinct and systematic evaluation/reporting approach, and as such main (and important) take-home lessons from their work are lost.

Response: Thanks very much to the Reviewer for this helpful feedback. Many of the specific comments below have helped us to clarify what previously appeared to be contradictory statements throughout the text.

Edits: Throughout the manuscript, we have made edits to clarify what could have been perceived as contradictory statements. We have added significant detail to better match the level of complexity in the findings we present and have attempted to more systematically and succinctly report these findings and associated take away messages for readers.

Specific comments:

- Abstract.** We assessed the performance of ambient ozone (O₃) and carbon dioxide (CO₂) sensor field calibration techniques when they were generated using data from one location and then applied to data collected at a new location. We also explored the sensitivity of these methods to the timing of field calibrations relative to deployments they are applied to.
- 10 Employing data from a number of field deployments in Colorado and New Mexico that spanned several years, we tested and compared the performance of field-calibrated sensors using both linear models (LMs) and artificial neural networks (ANNs) for regression. Sampling sites covered urban, rural/peri-urban, and oil and gas production influenced environments. Generally, we found that the best performing model inputs and model type depended on circumstances associated with individual case studies. In agreement with findings from our previous study that was focused on data from a single location
- 15 (Casey et al., 2017), ANNs remained more effective than LMs for a number of these case studies but there were some exceptions. In almost all cases the best CO₂ models were ANNs that only included the NDIR CO₂ sensor along with temperature and humidity. The performance of O₃ models tended to be more sensitive to deployment location than to extrapolation in time while the performance of CO₂ models tended to be more sensitive to extrapolation in time than to deployment location. The performance of O₃ ANN models benefited from the inclusion of several secondary metal oxide
- 20 type sensors as inputs in many cases.

Comment: L9. Avoid ending sentence with 'to'

Response: Thank you for catching this mistake.

Edits: We changed the sentence structure accordingly. "We also explored the sensitivity of these methods in response to the timing of field calibrations relative to deployments periods."

Comment: L13: this is one of the core conclusions: the resilience of a given calibration model depends on the circumstances of the deployment for that same sensor system. As such, the paper would be strengthened if the authors focused the narrative on succinctly describing such dependences and circumstances relating these factors back to the sensitivity, selectivity, and stability of each sensor system and sensor type.

Response: Thanks very much to the Reviewer for this helpful comment.

Edits: Throughout the paper, we have added text to help focus the narrative in the context of relative circumstances present in individual case studies, and how model performance in each case study relates to sensor sensitivity, selectivity, and stability.

Comment: This language is far too vague, especially for an abstract. What circumstances?

Response: Thank you for noting how we can make the abstract more informative and less vague.

Edits: We have added descriptions of specific circumstances: “We found that the best performing model inputs and model type depended on circumstances associated with individual case studies, such as differing characteristics of local dominant emissions sources, relative timing of model training and application, and the extent of extrapolation outside of parameter space encompassed by model training.”

Comment: L15: ‘a number’ - again, this is too vague. Define exactly how many of the case studies were characterized as having superior AAN models and how many were just as well served with an LM model

Response: Thanks to the Reviewer for the helpful comment.

Edits: We have added the following detail to the abstract: “Among models that were tailored to cases studies on an individual basis, O₃ ANNs performed better than O₃ LMs in 6 out of 7 case studies, while CO₂ ANNs performed better than CO₂ LMs in 3 out of 5 case studies.”

Comment: L16: This line suggests that people should model CO₂ with ANNs not LMs. The more detailed discussion in the body of the paper contradicts this assertion.

Response: Thanks to the Reviewer for the helpful comment.

Edits: After further consideration, we determined that this statement was an oversimplification and so have removed it from the abstract.

Comment: L19: subscript O₃

Response: Thanks to the Reviewer for catching this.

Edits: O₃ subscripted: O₃

5 1.1 Low-Cost Sensors For Air Quality Measurements

The use of low-cost metal oxide, electrochemical and non-dispersive infrared sensors to characterize air quality is becoming increasingly common across the globe (Clements et al., 2017; Kumar et al., 2015). Field normalization methods to quantify atmospheric trace gases using low-cost sensors have been found to be more effective than lab calibration when sensors are deployed in the ambient environment, and subject to changing temperature and humidity (Cross et al., 2017; Piedrahita et al., 2014; Sun et al., 2016). Our previous study and several others have compared the efficacy of LMs (simple and multiple linear regression) relative to supervised learning methods, all finding that ANNs (Casey et al., 2017; Spinelle et al., 2015, 2017) and random forests (Zimmerman et al., 2017) outperformed LMs in the ambient field calibration of low-cost sensors. These effective supervised learning techniques often incorporate multiple gas sensor signals as inputs in order to quantify each target gas, in addition to environmental variable sensor signals, with the goal of compensating for the effects of interfering gas species and environmental factors. In practice though, the reason multiple gas sensors are able to improve the performance of supervised learning type regression, and linear models for that matter, may be in part the result of correlation, or correlation for some time periods, between mole fractions of target gases themselves that hold for one model training location, but might not hold up at alternative sampling sites or during other time periods.

Comment: L11: What is the difference between supervised learning methods and ANNs? This warrants a more detailed description / definition.

Response: Thanks to the Reviewer for the helpful comment.

Edits: We have added the following text to help clarify that ANNs are an example of a supervised learning method, as are random forests: “supervised learning methods (including ANNs and random forests)”

Comment: L15: This sentence (bracketed in red) - is very important, but also very wordy and hard to follow.

Response: Thanks to the Reviewer for the helpful comment.

Edits: We have simplified and clarified the sentence accordingly: “In practice though, the reason multiple gas sensors are able to improve the performance of calibration models may be in part the result of correlation between mole fractions of target gases themselves that hold for one model training location, but might not remain effective at alternative sampling sites or during other time periods.”

Comment: Related to this assertion, it is not clear how the authors disentangle the temporal and spatial domain from one another, particularly the temporal domain. Time-decay patterns in the data are going to be present whether or not the sensor system has been moved to a different location. How would one ascribe difference in that case to a spatial domain and not temporal domain?

Response: In this work, we attempt to disentangle the temporal and spatial domain by including some case studies where models were only extended outside of their training spatial domain or only extended outside of their temporal domain, but not both. We attempt to represent time-decay patterns effectively in models by using pre/post training data for some case studies.

Edits: We have clarified this strategy by more clearly defining what is meant by ‘extrapolation in time’, and by more clearly identifying which case studies were subject to extrapolation in time. We have added the following text to section 2.3 accordingly: “A model was extrapolated in time when ever training data does not take place both before and after a given test deployment period. In several case studies we present, model training only took place after the test deployment period, comprising a ‘post only’ calibration. In Colorado, and more broadly in the western United States, ambient temperatures change significantly across the seasons throughout the year, so if a model is extrapolated in time, extrapolation in temperature often results as well.”

Comment: ‘hold up’ this language is too casual and used throughout the text. Consider re-wording.

Response: Thank you for the feedback.

Edits: We have replaced ‘hold up’ with ‘remain effective’.

Section 1.2

Comment: L28: ‘A number of enclosures..’ define the number.

Response: Thank you for the feedback and helping us to clarify how many U-Pods were included in each case study and in our previous work.

Edits: In section 1.3 (previously section 1.2) we have added the following text: “The study tested and compared calibration models using data from two U-Pod sensor systems”. We have also updated the text in the methods section and Fig. 2 (previously Fig.1) to explicitly list the number of U-Pods included in each case study.

Comment: If Casey et al., 2017 demonstrated the ANN results for CO₂ and O₃ in the Spring of 2017 in Greeley, CO; is that same data being presented as a portion of this paper (as Figure 1 suggests).

Response: Yes, it is the same data. Thank you for pointing out that this needs clarity.

Edits: We have clarified this, adding sub subsections within section 2.2 describing each case study individually, including our previous study in section 2.2.7: “We include findings from our previous work as a case study in order to provide context. Models for CO₂ and O₃ were tested using data from two U-Pods collected over the course of approximately one month at the GRET site in the spring of 2017. Data from two U-Pods during approximately month-long periods pre and post of the test period were used to train O₃ and CO₂ models. This case study provides another example of model performance when training took place both pre and post of the test period, and testing took place in the same location as training.”

Comment: The concluding sentences of this section nicely frame the motivation/need for the current work, consider bringing this to the fore of the paper / abstract, etc.

Response: Thanks to the Reviewer for the helpful comment.

Edits: We have augmented the abstract with this piece of motivation by adding the following text: “This was motivated by a previous study (Casey et al., 2017) which highlighted the importance of

determining the extent to which field calibration regression models could be aided by relationships among atmospheric trace gases at a given training location, which may not hold if a model is applied to data collected in a new location. We also explored the sensitivity of these methods in response to the timing of field calibrations relative to deployments periods.”

We have also augmented the for of the paper by added the following text to the end of the first paragraph in the introduction: “ANNs, as powerful pattern recognition tools, were found to perform better than both inverted and direct LMs in our previous study, but concerns arose when findings suggested that the performance of ANNs was being augmented by the relationships among gas mole fractions in the atmosphere at a given location. Low-cost gas sensor systems have the potential to inform spatial and temporal variability of pollution, when calibration equations for each sensor system are generated in one location based on co-located measurements with reference instruments, then moving the sensor systems into a spatially distributed network. Since the relationships among gas mole fractions at different sampling sites across a spatially distributed network, calibration models may not hold at new sampling sites. In this work, we test calibration model performance when extended to new locations.”

Section 1.3

Comment: Final sentence: It’s unclear why, if all of the U-POD sensor systems were equipped to measure CO and CH₄ alongside CO₂ and O₃, analogous training/test matrix pairs are unavailable for these other species.

Response: Thanks to the Reviewer for the helpful comment.

Edits: We have clarified that while the sensors for CO and CH₄ were included in the U-Pods during all the presented case studies, reference measurements for these species were not available: “In previous work (Casey et al., 2017) we have additionally addressed the quantification of CO and CH₄ using arrays of low-cost sensors together with field normalization methods, but these species are not included in the present analysis because reference data for model training and testing deployment pairs, diverging in location and timing and analogous to those we present for O₃ and CO₂, were not available CO and CH₄.”

Section 1.4

Comment: L20: ‘Very high levels of ozone’ – specify the actual concentration or concentration range

Response: Thanks to the Reviewer for the helpful comment.

Edits: We have replaced ‘Very high levels of ozone’ with ‘Mole fractions of ozone in as high as 140 ppb and 117 ppb during winter months have also been observed and attributed directly to oil and gas production emissions in the Upper Green River Basin of Wyoming and Utah’s Uinta Basin, respectively”

Comment: L23: ‘a modeling study’ – is there really only one modeling study that shows this?

Response: Thanks for this comment. This is the only modeling study we know of that was focused on the effects of oil and gas production emissions on ozone, with potentially high spatial near emissions sources.

Comment: Final sentence: ‘pooling’ avoid using words with common association different from the intended meaning. Consider re-wording. ‘accumulating’?

Response: Thanks to the Reviewer for the helpful comment.

Edits: We replaced ‘pooling’ with accumulation’

Section 2.1

Comment: L5: “with a number of low-cost gas sensors” – specify the actual number of sensors integrated in each U-POD. The authors identify that 10 U-POD systems were used in the previous and current work, but the vast majority of case studies (outlined in Figure 1.) utilize just 2 U-PODs at each

location. The authors need to more clearly describe in the text how the U-PODs were distributed throughout the work and whether all 10 U-PODs used in the current work had the same characteristic O₃ and CO₂ response when measuring the same air.

Response: Thanks to the Reviewer for the helpful comment.

Edits: We have replaced added subsections to section 2.2 describing each case study, including how many U-Pods were included in each:

2.2.1 Dawson Summer 2014

The first distributed measurement campaign took place during the summer of 2014 when five U-Pods were sited at locations around Boulder County, with four distributed along the eastern boundary of the county, adjacent to Weld County where dense oil and gas production activities were underway. A background site, further from oil and gas production activities was also included to the west, near a busy traffic intersection on the north end of the City of Boulder. Co-locations with reference measurements that were used for field calibration of sensors took place at the Continuous Ambient Monitoring Program (CAMP) Colorado Department of Health and Environment (CDPHE) air quality monitoring site in downtown Denver. One of the distributed sampling sites, Dawson School, was also equipped with an O₃ reference instrument (a Thermo Electron 49) ,operated by Detlev Helmig's research group from the Institute for Arctic and Alpine Research (INSTAAR). In this work, a case study is developed using data from one U-Pod located at the CAMP site in downtown Denver for O₃ model training, and data from one U-Pod, located at the Dawson School for O₃ model testing. This case study is used to test model performance when extrapolated in terms of O₃ mole fractions and applied in a new location, transferred from an urban to a peri-urban environment.

2.2.2 SJ Basin Spring 2015

In the spring of 2015 we augmented our original fleet of five U-Pods (BA, BB, BD, BE, and BF) with five more (BC, BG, BH, BI, and BJ) and deployed these sensor systems in the SJ Basin while a targeted field campaign was underway to understand more about a CH₄ 'hot spot' that was discovered from satellite based remote sensing measurements (Frankenberg et al., 2016; Kort et al., 2014). The primary goal of this sensor deployment was to inform spatial and temporal patterns in atmospheric trace gases like CH₄, O₃, CO, and CO₂ across the SJ Basin. Most U-Pods were located at existing air quality monitoring sites operated by the New Mexico Air Quality Bureau (NM AQB), the Southern Ute Indian Tribe Air Quality Program (SUIT AQP), and the Navajo Environmental Protection Agency (NEPA), which supported validation of sensor measurements for O₃ After this deployment period, all U-Pods were moved to the BAO site in the DJ Basin for approximately one month, and were co-located with reference instruments there that were operated by National Oceanic and Atmospheric Administration (NOAA) researchers. A case study is developed with data from the BAO site to train O₃ models for four U-Pods, and data from SJ Basin sites to test O₃ models for four U-Pods. This case study is used to test model performance when extrapolated in temperature and time, and extended to a new location, extended from one oil and gas production basin to another across Colorado

2.2.3 SJ Basin Summer 2015

In the summer of 2015, after an approximately month-long co-location with reference instruments at the BAO site, seven U-Pods were deployed again at existing regulatory monitoring sites for approximately one month, after which they were moved back to the BAO site for another month of co-location with reference instruments there. We equipped two of the regulatory monitoring sites in the SJ Basin with LI-COR LI-840A CO₂ analysers to provide reference measurements for CO₂. A case study is developed with data from the BAO site, pre and post of the SJ Basin summer 2015 deployment to train models, and data from SJ Basin sites during the summer deployment period, to test models. Data from seven U-Pods were used to train and test O₃ models and data from two U-Pods were used to train and test CO₂ models. This case study is used to test model performance when training took place both pre and post of the test period, and when extended to a new location, from one oil and gas production basin to another across Colorado

2.2.4 BAO Summer 2015

During the SJ Basin Summer 2015 deployment period, two U-Pods remained at the BAO site. A case study is developed using data from two U-Pods the BAO site, pre and post of the summer 2015 deployment to train models for O₃ and CO₂, and data from two U-Pods the BAO site during the summer deployment period to test models for O₃ and CO₂. This case study is used to test model

performance when training took place both pre and post of the test period, and when the model was tested on data that was collected in the same location as model training.

2.2.5 BAO Summer 2016

U-Pods were deployed at the BAO site again in 2016 for several months during the summer. In August of 2016 the U-Pods were moved to the Greeley Tower (GRET) CDPHE air quality monitoring site in Greeley, Colorado, a location which, like the BAO site, is also strongly influenced by DJ Basin oil and gas production activities; the U-Pods remained there for a year. For the GRET co-location period, CDPHE shared reference measurements for O₃. Additionally, Jeffrey Collett and Katherine Benedict of Colorado State University (CSU) shared CO₂ reference measurements from an instrument they operated at the site before October 1st in 2016 and after March 7th in 2017, when the instrument was located at the GRET site. A case study is developed using data from two U-Pods during the yearlong deployment at the GRET site to train models for O₃, and data from two U-Pods during the BAO summer 2016 deployment to test models for O₃. Because reference data for CO₂ was not available at the GRET site during winter months, data from two U-Pods during eight months at the GRET site was used to train models for CO₂, and data from two U-Pods during the BAO summer 2016 deployment was used to test models for CO₂. A significantly longer training duration is implemented in this case study because the training period took place more than several months after the model testing period. We reasoned that a longer training duration would be better able to represent patterns in sensor drift over time, as well as encompass the temperature range of test dataset period. This case study is used to test model performance when extrapolated significantly (more than several months) in time and extended to a new location, from one location in DJ oil and gas production basin to another.

2.2.6 GRET Fall 2016

In order to test model performance, under similar circumstances in terms of relative model training and testing durations and timing, to the BAO Summer 2016 case study, but with no extension of models to a new location, we developed another case study. This time, models for O₃ and CO₂ were trained using data from two U-Pods at GRET over the course of eight months and models for O₃ and CO₂ were tested using data from two U-Pods at GRET over the course of approximately a month in the fall of 2016. This case study is used to test model performance when extrapolated significantly (more than several months) in time and applied in the same location as training took place.

2.2.7 GRET Spring 2017

We include findings from our previous work as a case study in order to provide context. Models for CO₂ and O₃ were tested using data from two U-Pods collected over the course of approximately one month at the GRET site in the spring of 2017. Data from two U-Pods during approximately month-long periods pre and post of the test period were used to train O₃ and CO₂ models. This case study provides another example of model performance when training took place both pre and post of the test period, and testing took place in the same location as training.”

2.2 Deployment Locations and Timelines

10 These ten U-Pods were deployed at a number of sampling sites in and around the DJ and SJ Basins over the course of several years, from 2014 - 2017. Deployments generally consisted of co-location with reference measurements prior to and following a period of spatially distributed measurements. During some the distributed measurement periods, a subset of U-Pods remained co-located with reference instruments where the field calibrations took place. During other distributed measurement periods, U-Pods were deployed in new locations that were equipped with reference measurements. We
15 opportunistically employ data from a number of these sensor deployments, treating them as case studies in order to characterize the performance of field calibration models when they are extended to new locations. For each case study, data was divided into training and test periods.

Table 2 lists the O₃ and CO₂ reference instruments that were co-located with U-Pods at each sampling site, along with
20 instrument operators, calibration procedures, and reference data time resolution. The first distributed measurement campaign took place during the summer of 2014 when five U-Pods were sited at locations around Boulder County, with four distributed along the eastern boundary of the county, adjacent to Weld County where dense oil and gas production activities were underway. A background site, further from oil and gas production activities was also included to the west, near a busy traffic intersection on the north end of the City of Boulder. Co-locations with reference measurements toward field
25 calibration of sensors took place at the Continuous Ambient Monitoring Program (CAMP) Colorado Department of Health and Environment (CDPHE) air quality monitoring site in downtown Denver. One of the distributed sampling sites, Dawson School, was also equipped with an optical O₃ instrument, operated by Detlev Helmig's research group from the Institute for Artic and Alpine Research (INSTAAR). This study was funded by Boulder County with the combined aims of gaining a better understanding of how oil and gas emissions affect air quality in Boulder County and learning more about how low-
30 cost air quality measurement methods can help inform air quality in this context.

Comment: The sensor system age (time since manufacture date) and environmental-hysteresis (lifetime environmental exposure of a given UPOD system) is not mentioned anywhere in the text. Do these factors not matter when analyzing the temporal extension of a given calibration model? When considering the fundamental measurement principles of these particular gas sensors, does degradation occur due to gradual (or rapid) deposition of material onto active catalytic sites within the sensors? If so, then the age of a given sensor and what's it's been exposed to over its lifetime, ought to factor in.. or at least deserve a mention.

Response: We agree with the importance and relevance of sensor challenges highlighted in this comment.

Edits: We have added the following text in section 2.2 accordingly: "Making quantitative measurements of atmospheric trace gases with low-cost sensors is challenged by unique variations in individual sensor responses associated with variations in the manufacturing process, sensor age, and sensor exposure history. For these reasons, we generated unique calibration models using data from sensors in each individual U-Pod sensor system. The closest available data prior and or subsequent to a test data period was used for model training to avoid complications associated with significant sensor drift and degradation in sensor sensitivity to target gas species over time if possible."

We have additionally added the following text has been added to section 3.2.2: "Gas sensor manufactures don't clearly define sensor lifetimes, but sensors are generally expected to loose sensitivity over time. For example, Alphasense CO-B4 electrochemical sensors are expected to have 50% of their original sensitivity after two years (Alphasense, 2015). The heater resistance in a give metal oxide type sensor is expected to drift over time, influencing sensor measurements (e2v Technologies Ltd., 2007). Masson and colleagues observed a significant drift in a metal oxide sensor heater resistance over the course of a 250 day sampling period in a laboratory setting (Masson et al., 2015). While we did not measure and record metal oxide sensor heater resistance for sensors included in U-Pods, we have investigated eltCO2 and e2vO3 sensor signal drift from the summer of

2015 through the summer of 2017. These data are presented in Fig. S26. Systematic downward drift in all eltCO_2 sensor signals is apparent over this time frame. A clear and consistent pattern of systematic drift over this time period is less apparent for e2vO_3 sensors. Since the training data was collected immediately after, the test data period, and since the test data period was relatively short (approximately one month) sensor drift could be negligible across the combined training/testing time frame. “

Comment: The explanation of the training vs test sampling periods is confusing as written. Given the nature of the experiment, doesn't each UPOD system have to be co-located with reference instrumentation for the full duration of the period of study? It sounds as though the authors aimed to bookend the distributed network measurements ('testing period' with a period of colocation at a reference site in the general vicinity of the deployment ('training period') – but in order evaluate their models, they would have to retain a co-located reference measurement of O_3 and CO_2 at all times in all locations.

Response: Yes, we present data, opportunistically, from test periods when sensor systems were co-located with O_3 and CO_2 reference instruments.

Edits: We have added text to section 2.2 to try to clarify these details: “Five to ten U-Pods were deployed at sampling sites in and around the DJ and SJ Basins over the course of several years, from 2014 - 2017. Deployments generally consisted of co-location with reference measurements prior to and following approximately one-month periods of spatially distributed measurements. During some of the distributed measurement periods, a subset of U-Pods remained co-located with reference instruments where the field calibrations took place. As well, during some distributed measurement periods, some U-Pods were deployed in new locations that were equipped with reference measurements. In between periods of distributed sensor system deployments, sensor systems were co-located with reference instruments for as long as possible, as logistics, and coordination with other regulatory agencies and researchers would allow. In this way, we hoped to maximize our ability to encompass full ranges of temperature, humidity, and trace gases that occur across seasons, in order to minimize extrapolation with respect to these parameters when models were applied to measurements from distributed deployment periods. The locations where all or a subset of U-Pods were co-located with reference instruments are indicated in Fig. 1. In this exploratory study, we opportunistically employ data from these sensor deployments, treating them as case studies in order to characterize the performance of field calibration models when they are extended to new locations. For each case study, described below, data was divided into training and test periods. Timelines for these dataset pairs detailed in Fig. 2. Some U-Pods used included in these case studies (indicated in grey font in Fig. 2) were constructed, populated with sensors, and deployed at field sites in the spring of 2014, approximately a year before the rest of the U-Pods were constructed, populated with sensors, and deployed at field sites in the spring of 2015. The relative age of sensor systems included in some case study comparisons could have contributed to some discrepancy in model performance, though systematic differences based on U-Pod age is not apparent.

As available data from each case study allowed, we used approximately one month of training data before and after (pre and post of) a given approximately month-long test period. When training data was not available within several months of a test period, significantly longer training datasets were used in order to attempt capture and effectively represent trends in sensor drift over time, as well as to avoid extrapolation of model parameters (particularly temperature) during the test data period. As a result, model-training durations varied across case studies and sometimes significantly exceeded model-testing durations. Each case study is similar in representing approximately one month-long deployment of sensor systems. This study design serves a primary goal of this work, which is to help support the quantification atmospheric trace gases from low-cost gas sensor data in new locations, relative to model training locations, for periods of approximately one month at a time.

Making quantitative measurements of atmospheric trace gases with low-cost sensors is challenged by unique variations in individual sensor responses associated with variations in the manufacturing process, sensor age, and sensor exposure history. For these reasons, we generated unique calibration models using data from sensors in each individual U-Pod sensor system. The closest

available data prior and or subsequent to a test data period was used for model training to avoid complications associated with significant sensor drift and degradation in sensor sensitivity to target gas species over time if possible. Table 2 lists the O₃ and CO₂ reference instruments that were co-located with U-Pods at each sampling site, along with instrument operators, calibration procedures, and reference data time resolution. The selected case studies, described in sections 2.2.1 through 2.2.7 below are aimed at supporting methods to quantify atmospheric trace gases during the distributed deployments we carried out from 2014 through 2017 as well as future distributed sensor network measurements. Fig. 1 shows sampling site locations in context with urban areas and oil and gas production wells. Fig. 2 shows the timeline of each of these deployments, highlighting the training and testing periods defined for both O₃ and CO₂.”

Comment: Looking at the deployment timelines displayed in Figure 1, it is also evident from the Figure (but not from the text) that the vast majority (~75% or greater) of the total deployment time was used to train the nodes not test the resultant calibration models (~25% of the total time). These train-to-test ratios appear to undermine the general applicability of the models to longer duration, distributed sensor measurements in which no co-located reference measurements are available. The authors should make an effort to bridge the gap between how they were able to execute their experiments and how distributed low-cost AQ sensor systems will ultimately be deployed.

Response: Thank you for helping us clarify why varying and sometime long durations of training data were used for each case study, and how we hope this study design can help support future sensor measurement efforts.

Edits: We have added the following text to section 2.2 toward this end:

“In between periods of distributed sensor system deployments, sensor systems were co-located with reference instruments for as long as possible, as logistics, and coordination with other regulatory agencies and researchers would allow. In this way, we hoped to maximize our ability to encompass full ranges of ambient temperature, humidity, and trace gases that occur across seasons, in order to minimize extrapolation with respect to these parameters when models were applied to measurements from distributed deployment periods.”

“In an effort to fully encompass the parameter space present and during each individual test deployment case study, as well as sensor drift over time, model-training durations varied across case studies and sometimes significantly exceeded model-testing durations. Each case study is similar in representing approximately one month-long deployment of sensor systems. This study design serves a primary goal of this work, which is to help support the quantification atmospheric trace gases from low-cost gas sensor data in new locations, relative to model training locations, for periods of approximately one month at a time.”

From section 2.2.5: “A significantly longer training duration is implemented in this case study because the training period took place more than several months after the model testing period. We reasoned that a longer training duration would be better able to represent patterns in sensor drift over time, as well as encompass the temperature range of test dataset period. Significantly less training time is needed when training occurs directly pre and/or post of a given model application period.

Highlighted in passage above:

Comment: L10: define the number of sampling sites. Eliminate vague language in the text. L15: same comment.

Response: The number of sampling sites during each case study varied, so to help clarify we directly reference the map showing each of the sampling sites included in the study.

Edits: We have renamed this ‘Figure 1’ and renamed the timeline ‘Figure 2’ accordingly.

Comment: L21: 5 UPOD systems are purportedly used in the Boulder / CAMP 2014 work. Figure 1 lists 1 UPOD system as being active during that test. Reconcile this.

Response: Thank you for helping us clarify that while 5 U-Pods were deployed during the Boulder County study, only one of the U-Pods was deployed at a location that had co-located reference measurements for O₃.

Edits: The number of U-Pods used in the Dawson Summer 2014 case study and others has been clearly updated in sections 2.2.1 – 2.2.7.

Comment: L27: Identify the actual ref. O₃ measurement in the text here

Response: Thank you, we agree indicating the specific instrument used would be useful to the reader.

Edits: “Thermo Electron 49”

Comment: Last sentence: is this relevant to the current paper/study? Not clear what ‘study’ the authors are referring to in this sentence.

Response: We agree this sentence lacked specific relevance to the current study.

Edits: We have removed this sentence.

seems to be connected to the sampling sites or the circumstances discussed previously as opposed to the quality of individual sensors in each of those U-Pods.

5 All SJ Basin U-Pod O₃ measurements systematically over estimate lower levels of O₃ each night, a trend apparent in the scatter plots in Fig. 5 and in the residuals by time of day plot in Fig. 6. Upon examination of the scatter plots in Fig. 5, U-Pods at some sampling sites had positive bias for higher O₃ measurements as well (Shiprock, Ignacio, Sub Station, and Bloomfield), while for others, bias at the higher end of O₃ distributions did not appear to be present (Navajo Dam, Fort Lewis, and Bondad). The residuals by time of day plot in Fig. 6 shows that the two U-Pods at BAO did not have significant trends in their residuals according to the time of day, but that U-Pods deployed at SJ Basin sites consistently over estimated
10 nighttime O₃. The residuals are also plotted with respect to temperature in Fig. 6, where all U-Pods, even those at BAO to a lesser extent, appear to over predict O₃ at lower temperatures, which generally occurred at night. The times of day that generally correspond to the highest O₃ levels generally had the lowest residuals, with some exceptions at the Fort Lewis and Navajo Dam sites.

15 Fig. 6 includes a plot of the residuals across the duration of the deployment period, showing no significant sensor drift in measurements for any of the U-Pods. This plot also shows that the highest residuals observed generally occurred over short periods in time, particularly for the Fort Lewis (blue) and Sub Station (magenta) sites. In order to further explore the performance of field calibration models for O₃ at SJ Basin sites relative to BAO, the combined parameter space of temperature with O₃ reference mole fractions and temperature with absolute humidity are presented in Fig. 7. The combined
20 temperature and reference O₃ parameter space appears to be similar for all of the U-Pods, both at BAO and the SJ Basin sites. However, there appears to be some outlying combined temperature and humidity parameter space at the Sub Station site and at the Navajo Dam site. Brief excursions of high humidity may be connected to some of the large short-term residuals observed at these two sampling sites.

We deployed a similar distributed network of sensor systems throughout the DJ Basin in the summer of 2015 and the winter/spring of 2016 to explore spatial and temporal patterns in air quality in another region influenced by oil and gas production activities. For these 2015 and 2016 deployments, we co-located U-Pods with reference instruments operated by National Oceanic and Atmospheric Administration (NOAA) researchers at the Boulder Atmospheric Observatory (BAO) Tower toward field calibration, though none of the distributed sampling sites were equipped with reference instruments to support validation. In August of 2016 the ten U-Pods were moved to the Greeley Tower (GRET) CDPHE air quality monitoring site in Greeley, Colorado, a location which is also strongly influenced by DJ Basin oil and gas production activities; the U-Pods remained there for a year. For the GRET co-location period, CDPHE shared reference measurements for O₃. Additionally, Jeffrey Collett and Katherine Benedict of Colorado State University (CSU) shared CO₂ reference measurements from an instrument they operated at the site before October 1st in 2016 and after March 7th in 2017, when the instrument was located at the GRET site.

20

The work presented here is aimed at supporting methods to quantify atmospheric trace gases during the distributed deployments described above as well as future distributed sensor network measurements. Fig. 1 shows the timeline of each of these deployments, highlighting the training and testing periods defined for both O₃ and CO₂. Fig. 2 shows sampling site locations in context with urban areas and oil and gas production wells.

25 2.3 Reference and Sensor Data Preparation

Each of the U-Pod sensor signals was logged to an onboard micro SD card. For metal oxide type sensors, voltage signals were converted into resistance, and then normalized by the resistance of the sensor in clean air. Relative humidity, temperature, and pressure measured in each U-Pod were used to calculate absolute humidity. Over the course of several field deployments, relative humidity sensors in a few of the U-Pods drifted down, causing the lower humidity levels to be cut off or 'bottomed out'. For measurements collected in the spring and summer of 2015 and the spring of 2017, we replaced the relative humidity (RH) signal of U-Pods with malfunctioning humidity sensors with signals from nearby U-Pods with good humidity sensors and complete data coverage as noted in Table S1.

Comment: L9: The authors claim that the SJ Basin network was similarly executed for the DJ Basin. DJ Basin is absent from Figure 1., replaced presumably by BAO. It is unclear how many UPODs were deployed to the DJ Basin. It's very confusing trying to track in time and location the distribution of the 10 UPODs. If I try and decipher the information in Figure 1, either 2 or 4 UPOD units were deployed to the DJ Basin, which on the face of it, does not constitute a similar network deployment of 10x UPODs deployed to the SJ Basin (although, it seems that only 4 and/or 7 UPOD units were deployed to the SJ Basin.

Response: Thanks to the Reviewer for the helpful comment.

Edits: We have added to the caption of Figure 2 (previously Figure 1) to help clarify that all sampling sites outside of the SJ Basin group were in the DJ Basin. "The Dawson, BAO, and GRET sampling sites are all located in the DJ Basin."

Comment: L13: The authors identify the BAO site as the relevant co-location site for the DJ Basin-deployed UPODS, but then point out that there were NO co-located reference instrumentation accessible for any of the distributed sampling sites. What does this mean for evaluating / testing their models in the distributed network application?

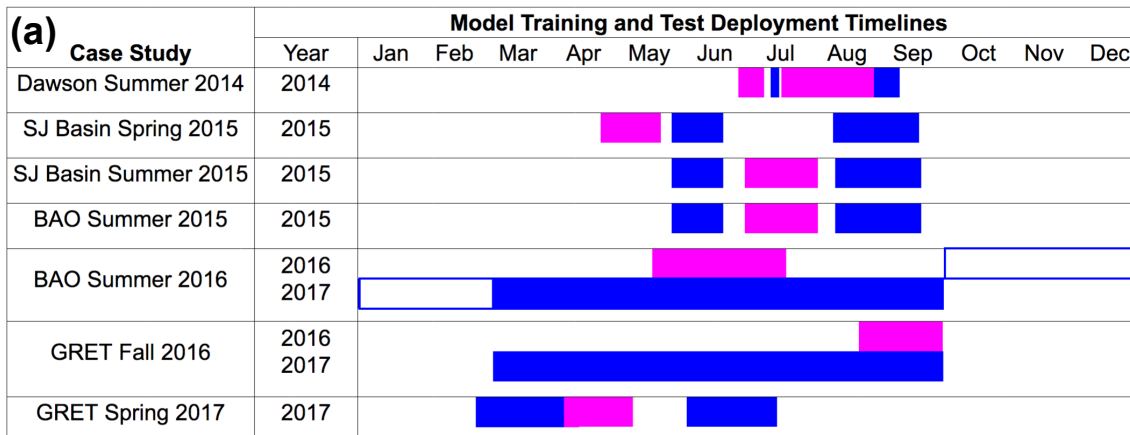
Response: Thanks to the Reviewer for the helpful comment.

Edits: We have made significant edits in section 2.2, more clearly defining which data is included in this work and in each case study. All discussion about the distributed deployment sites that did not have reference measurements has been removed from the text, since these deployments helped to motivate, but are not directly relevant to the present work.

Comment: L14-16: The authors state the GRET site housed all 10x UPOD systems for a year, but Figure 1 indicates that only 2-6? UPOD systems were used at this location and only for shorter periods of time. Again, the text is extremely hard to follow and the information in Figure 1 does not make it any clearer.

Response: Thanks very much to the Reviewer for pointing out the confusing nature of how the information is presented.

Edits: We have updated Figure 2 (previously Figure 1) to help clarify which U-Pods were included in each case study.



(b) Case Study	Training Location	Test Location	O ₃ # U-Pods	O ₃ U-Pod Names	CO ₂ # U-Pods	CO ₂ U-Pod Names
Dawson Summer 2014	CAMP	Dawson	1	BE	NA	NA
SJ Basin Spring 2015	BAO	SJ Basin	4	BB, BD, BF, BJ	NA	NA
SJ Basin Summer 2015	BAO	SJ Basin	7	BA, BB, BD, BE, BF, BH, BI	2	BB, BD
BAO Summer 2015	BAO	BAO	2	BC, BJ	2	BC, BJ
BAO Summer 2016	GRET	BAO	2	BH, BI	2	BH, BI
GRET Fall 2016	GRET	GRET	2	BH, BI	2	BH, BI
GRET Spring 2017	GRET	GRET	2	BH, BI	2	BF, BI

Figure 2: (a) ANN and LM training and test deployment timelines. The Dawson, BAO, and GRET sampling sites are all located in the DJ Basin. Model training periods for each test deployment are shown in blue, and model test periods are shown in magenta. For the BAO Summer 2016 case study, the period outlined in blue shows data that was used to train O₃ model, but not CO₂ models since CO₂ reference data was not available during winter months. (b) Information about each of the case studies presented in the above timelines, including model training and testing locations, as well as the number and names of U-Pods included in each case study for both O₃ and CO₂ models. The U-Pods with names shown in grey were constructed and deployed starting in May of 2014. The U-Pods with names shown in black were constructed and deployed starting in April of 2015.

Comment: L26: The only metal oxide sensor that's relevant to the current work is the e2vO3 sensor. The operational fundamentals of this sensor should be described: the raw signal processing, circuitry considerations, and known theoretical operational conditions that undermine the sensitivity, selectivity, and/or stability of the e2vO3 metal oxide sensor.

Response: Thanks to the Reviewer for the helpful comment.

Edits: Since models in this work included signals from multiple gas sensors, we have added a discussion of the operating principles of metal oxide, electrochemical, NDIR the sensors accordingly,

as well as discuss these sensor properties in context with model development in section 1.1, and 1.2. Additionally, we discuss these sensor considerations in context with unique challenges associated with measurements in oil and gas production basins in section 1.5:

“While low-cost sensors have been emerging on the market with sufficient sensitivity to resolve variations in ambient mole fractions of target gases of interest, they are also sensitive to temperature and humidity variations that occur in the ambient environment. NDIR sensors, like the ELT s300 CO₂ sensor employed in this study, have good selectivity, but, since pressure and temperature are not controlled in the optical cavity of ELT s300 CO₂ sensors, the influence of temperature on sensor signals plays an important role. The influence of humidity is also important to address because changes in water vapor are known to influence NDIR measurements of CO₂ in terms of spectral cross-sensitivity due to absorption band broadening (Licor, 2010).

Both metal oxide and electrochemical type sensors operate on the principle of oxidizing or reducing reactions at sensor surfaces. For electrochemical sensors, like the Alphasense CO-B4 sensor employed in this study, oxidizing or reducing compounds react at the working electrode, resulting in the transfer of ions across an electrolyte solution from the working electrode to the counter electrode, balanced by the flow of electrons across the circuit connecting the working electrode to the counter electrode. A linear relationship is expected between this current and the target gas mole fraction. Electrochemical sensors can be tuned to respond more or less strongly to specific gases by adjusting the materials properties of the working electrode. A membrane is located between the working electrode and the exterior of the sensor in order to control redox reaction rates. Gases diffuse through the membrane to reach the working electrode and the electron transfer rates have been shown to increase at higher temperatures (Xiong and Compton, 2014), and since chemical reaction rates are also influenced by temperature, electrochemical sensor responses can be influenced by sensor operating temperature. Changes in ambient humidity levels can cause sensors to lose or gain the electrolyte solution, by mass, also influencing electrochemical sensor response (Xiong and Compton, 2014).

For metal oxide sensors, and to a lesser extent for electrochemical sensors, resolving the response of a sensor attributable to the target gas species can also pose a challenge in the presence of interfering gas species. Metal oxide sensors, like those used in this study, have a resistive heater circuit that warms up the sensor surface, causing O₂ molecules to adsorb to the sensor surface, which leads to increased resistance across the surface of the sensor. In the presence of an oxidizing compound, like O₃, more oxygen molecules are adsorbed to the sensor surface and the resistance across the sensor surface is increased further. In the presence of a reducing compound, like CO, oxygen molecules are removed from the sensor surface, allowing electrons to flow more freely, resulting in decreased resistance across the sensor surface. For metal oxide sensors, the resistance across the sensor surface can then be used to determine the mole fraction of a given oxidizing or reducing compound, often according to a nonlinear relationship. Exposure to humidity has been shown to significantly lower the sensitivity of metal oxide gas sensors making it an important parameter to address in a gas quantification model (Wang et al., 2010). Metal oxide sensor operating temperature has also been shown to strongly influence sensor sensitivity and selectivity to different gas species (Wang et al., 2010). Metal oxide type sensors can be tuned to respond differently from one another to oxidizing and reducing gas species by using different metal oxide materials and doping agents for the sensor surface, but selectivity is difficult to achieve.

1.2 Low-Cost Air Quality Sensor Quantification

Because low-cost gas sensor signals are influenced, sometimes significantly, by interfering gas species and changing weather conditions in the ambient environment, field normalization methods to quantify atmospheric trace gases using low-cost sensors have been found to be more effective than lab calibration (Cross et al., 2017; Piedrahita et al., 2014; Sun et al., 2016). Our previous study and several others have compared the efficacy field calibration models generated using LMs (simple and multiple linear regression) relative to supervised learning methods (including ANNs and random

forests), all finding that ANNs (Casey et al., 2017; Spinelle et al., 2015, 2017) and random forests (Zimmerman et al., 2017) outperformed LMs in the ambient field calibration of low-cost sensors. Like earlier laboratory based studies (Brudzewski, 1999; Gulbag and Temurtas, 2006; Huyberechts and Szeco, 1997; Martín et al., 2001; Niebling, 1994; Niebling and Schlachter, 1995; Penza and Cassano, 2003; Reza Nadafi et al., 2010; Srivastava, 2003; Sundgren et al., 1991), ANN-based calibration models, incorporating signals from an array of gas sensors with overlapping sensitivity as inputs, have been able to effectively compensate for the influence of interfering gas species and resolve the target gas mole fraction.

ANNs are known to be able to very effectively represent complex, nonlinear, and collinear relationships among input and output variables in a system (Larasati et al., 2011). ANNs are useful in the field calibration of low-cost sensors because, through pattern recognition of a training dataset, they are able to effectively represent the complex processes and relationships among sensors and the ambient environment that would be very challenging to represent analytically or based on empirical representation of individual driving relationships. In practice though, the reason multiple gas sensors are able to improve the performance of calibration models may be in part the result of correlation between mole fractions of target gases themselves that hold for one model training location, but might not remain effective at alternative sampling sites or during other time periods.”

“In this work, we present and compare models designed to address the unique challenges that come with using low-cost sensors, in the quantification of atmospheric trace gases of interest in oil and gas production basins, where ambient hydrocarbon mole fractions are potentially elevated, exerting uniquely confounding influence on low-cost gas sensors. We investigate how well models can be transferred from one microenvironment to another, with different dominant local emissions source characteristics, and different relative abundance of oxidizing and reducing compounds. Microenvironments explored in this work include an oil and gas basin where both natural gas and heavier hydrocarbons are produced (the DJ Basin), and an oil and gas production basin where prominently natural gas is produced (the SJ Basin), with much smaller proportional emissions of heavier hydrocarbons, and in turn, lower atmospheric concentrations of alkanes. Within and bordering the DJ Basin, additional microenvironments include an urban location, with significant mobile sources emissions (NO_x, CO, and VOCs), and a peri-urban site with fewer mobile emissions and closer proximity to oil and gas production activities. We explore how robust model performance is when a model is trained in one microenvironment and transferred to another; challenged by different relative abundance of oxidizing and reducing gas species. Additionally we test how well models can represent and address sensor stability over time and the potential for drift. “

Comment: L29: ‘in a few’ Quantify the number of UPODs with faulty RH sensors

Response: Thanks to the Reviewer for the helpful comment.

Edits: We have replaced ‘in a few’ with ‘in four’.

Comment: L31: ‘nearby’: Define the exact position relative to the faulty UPOD

Response: Thanks to the Reviewer for the helpful comment.

Edits: We have updated this passage to include specific information about the relative positions of U-Pods when faulty humidity signals were replaced: “The closest U-Pod with good humidity sensors ranged from several feet, when U-Pods were co-located during deployments in the DJ Basin, to approximately fifty miles during deployments in the San Juan Basin.”

When the U-Pods were initially deployed at the GRET site, on August 23rd of 2016, the RH sensors in all ten U-Pods malfunctioned, logging an error code of -99 instead of the relative humidity. This malfunction seemed to coincide with the implementation of radio communication from each U-Pod to a central node in an effort to reduce trips to the field site to download data and to identify issues with data acquisition promptly. RH signals in the U-Pods began logging correctly again in November when we stopped remote communication. We replaced RH values for the U-Pods during this time period by utilizing data from the Picarro Cavity Ring-Down Spectrometer that was co-located at GRET with the U-Pods. Water mole fractions measured by the Picarro were converted into mass-based mixing ratios to match the units of the absolute humidity signal in the U-Pod data. We then replaced the absolute humidity signal in each U-Pod from August 23rd through October 1st in 2016 with the mixing ratios derived from Picarro measurements. Using the temperature and pressure logged in each U-Pod along with the absolute humidity from the Picarro, relative humidity was calculated for each U-Pod during this period.

To perform regressions toward field calibration of sensors, the reference and U-Pod data needed to be aligned. When reference measurements with minute time resolution were available for both training and corresponding testing periods, minute median data from the U-Pods were used. Medians were used as opposed to averages in order to reduce the potential influence of sensor noise as well as to remove short duration spikes in the reference and sensor data that resulted from air masses that may not have been well mixed across the reference instrument inlets and the U-Pod enclosures. When reference data were instead available with only 5-minute or 60-minute time resolution, U-Pod medians were calculated for the same time step. Medians were also calculated for reference measurements with finer time resolution to match the time resolution of corresponding training/testing data. The first 15 minutes of data after any period that the U-Pods had not recorded data for the previous 5 minutes was removed in order to filter transient behavior associated with sensor warm-up.

Comment: Did the implementation of radio communication for the UPODs have any impact on any of the other measurements in the system, beyond RH?

Response: We have added the following text to help address this question for the Reviewer and other readers.

Edits: “No other impacts to sensor systems were observed in connection with radio communications.”

Comment: At the beginning of the paragraph, the authors state that the radio communication was active until November, but the substitute RH values from the Picarro were only applied up to October 1 (later part of the paragraph). This is confusing.

Response: Thank you for catching this conflict.

Edits: We have corrected it by changing “November” to “October” in the first instance.

Comment: Generally speaking, faulty or absent RH measurements on-board the UPOD (or any low-cost AQ sensor system that suffers from environmental interference) is a potentially widespread issue across the emerging field. I think the authors missed an opportunity to discuss their work-around in more detail and comment on the importance of maintaining stable RH measurements within any given low-cost AQ sensor system.

Response: Thanks to the reviewer for this helpful comment. We added to the text to help clarify the work around that we implemented for faulty RH sensor data. Over the course of multiple field deployments of U-Pod sensor systems, including those described in this work, RHT03 sensors signals were found to drift down over time, and “bottom out” in some cases. Following this observation, we have since upgraded to Sensirion AG SHT25 sensors which appear to be more robust and consistent over the course of long-term field deployments. Hopefully this information will be as helpful to readers as the more through discussion of the work around we have added.

Edits: We have added the following text accordingly:

“Over the course of multiple field deployments, relative humidity sensors in four of the U-Pods drifted down, causing the lower humidity levels to be cut off or ‘bottomed out’. RH sensors were not replaced during field deployments in order to preserve consistency across different deployment

periods, allowing for the possibility of a single comprehensive model to apply to all data from a single U-Pod. After some experimentation in generating a 'master model' that could be applied to data from a given U-Pod for all collected field measurements, across several years, we determined that individual models for each deployment would be more effective, and replacing RH sensors that had drifted down would have been appropriate in support of the methods presented here. We have since upgraded to Sensirion AG SHT25 sensors, which appear to be more robust and consistent over the course of long-term field deployments."

Comment: The completely unusable radio communication RH values and the drifting RH values mentioned in section 2.3 beg the question – do the authors think this is a failure on the RHT component itself or the circuitry of the UPODs. Again, if the evidence suggests the former, that is useful empirical data for others in the field.

Response: Thanks to the Reviewer for bringing up this important question. We have not yet determined whether the failure of the RHT sensor signals during periods of active radio communications were connected to the sensors themselves or to the circuitry of the UPODs. This will be important to determine for the sensor community and before we try to implement radio communications again. As indicated in the previous comment, the drift of RHT03 sensors over time appeared to be an issue associated with the sensor model itself.

Comment: Where is RH measured specifically within each UPOD. Is the measurement internal to the box or positioned in a manner to provide a true ambient RH measurement? What are the implications of using alternative RH data sources that are not on-board the same UPOD?

Response: RH is measured within each U-Pod enclosure, in the microenvironment where the gas sensors are located. Using an alternative source for RH data that are not onboard and individual U-Pod has the potential to increase uncertainty of quantified gas mole fractions.

Edits: We have added the following text accordingly: "Temperature and RH sensor measurements are usually collected from within each U-Pod sensor system, in order to gain representative information about the environment the gas sensors are being operated in. Using an alternative source for RH data that are not onboard and individual U-Pod has the potential to increase uncertainty of quantified gas mole fractions. We used replacement RH data from the closest available U-Pod instead of ambient measurements in order to more closely match operating temperature within a U-Pod enclosure."

Comment: If median values were used for the co-located reference instruments, but the data from those instruments was 1-min averages, how did the authors obtain reference measurement medians at 1- min (the vast majority of temporal resolution used in the current work).

Response: Thank you for pointing out that this passage was confusing. We have changed the text to help clarify.

Edits: "In order to test models using the same time resolution they were trained with, the time resolution of reference and sensor measurements for corresponding training/testing datasets were matched, if necessary, by taking medians of the dataset with higher time resolution to match the data with the longer time resolution."

Comment: L19: What % of the total data used in training/testing each UPOD was removed due to this 5-min null data condition?

Response: We agree this information is useful for readers.

Edits: Accordingly, we have added the following text: "During a given deployment, the data removed to avoid sensor warm-up transients constituted less than 1%."

Section 2.4

Comment: L32 'using methods described previously', given the importance of the LMs and ANNs in the current work, each model should be described in more detail in the manuscript.

Response: Thanks to the Reviewer for the feedback. We are happy to provide more detail.

Edits: Two useful figures from our previous paper, showing ANN architecture, have been added and cited (now Figures 2 and 3) to help clarify. The following text has also been added:

“ As in (Casey et al., 2017), direct LMs and ANNs were trained with a number of different sensor input sets to map those inputs to target gas mole fractions measured by reference instruments. Direct LMs implemented were multiple linear regression models given by

$$r = p_1 + p_2s_1 + p_3s_2 + \dots + p_ns_{n-1} \quad (1)$$

where r is the target gas mole fraction (measured by a reference instrument) $s_1 - s_{n-1}$ are sensor signals from U-Pods that are included as model predictor variables, and $p_1 - p_n$ are corresponding predictor coefficients.

ANNs designed for regression tasks, like those employed in this work, generally consist of artificial neuron nodes that are connected with weights. Weights are initiated with randomly assigned values. An optimization algorithm is then employed to map a given set input values to corresponding target values. An example of a very simple feed forward neural network, and how weights are propagated through it are depicted in Fig. 3.

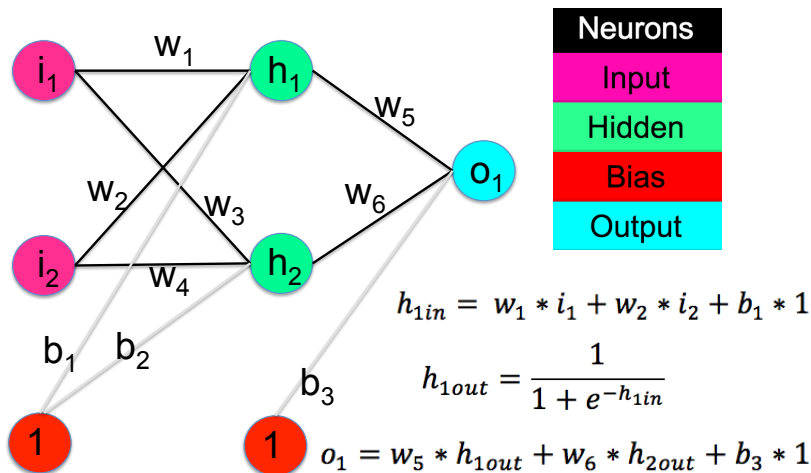


Figure 3. Example of a simple feed forward neural network, showing how inputs are propagated through the network during each of the training iterations (Casey et al., 2017)

In this work, ANNs were designed by assigning U-Pod sensor signals to artificial neurons in an input layer and assigning target gas mole fractions for an individual gas species, measured by a reference instrument to a single output neuron. Nonlinear, tansig, artificial neurons in one hidden layer for O_3 or two hidden layers for CO_2 (accordance with our earlier findings for each target gas species (Casey et al., 2017)) were then added between input layer and the network output neuron. Additionally, bias neurons, each assigned a value of 1, were connected to neurons in the hidden layer(s) so that individual connecting weights could be activated or deactivated during the optimization process. The number of neurons in each hidden layer was set equal to the number of inputs included in a given ANN. Fig. 4 shows a diagram of an ANN architecture employed in this work, when there were five inputs.




Figure 4. Diagram of an example ANN with the same color-coded components as are presented in Figure SM3 in section 2.2 of the SM. This ANN has 5 inputs, 1 hidden layer with 5 tansig hidden neurons, and 1 linear output layer leading to 1 output. The network is fully connected with weights and biases (Casey et al., 2017).

For ANN training we employed the Levenberg Marquardt optimization algorithm with Bayesian Regularization (Hagan et al., 1997). The Levenberg-Marquardt algorithm combines the Gauss-Newton and Gradient Decent methods, towards incremental minimization of a cost function (the summed squared error between the ANN output and target values as a function of all of the weights in the network). Training begins according to the Gauss-Newton method, in which the Hessian matrix (the second order Taylor series representation of the error surface) is approximated as a function of the Jacobian matrix and its transpose, significantly reducing required training time. Network weights are adjusted accordingly each training step to reduce error. If the cost function is not reduced in a given training step, an algorithm parameter is adjusted so that optimization more closely approximates the gradient decent method (a first order Taylor series representation of the cost function), providing a guarantee of convergence on a cost function minimum. Since local minima may exist across the error surface, it is important to train the same network multiple times (with different randomly assigned starting weights), in order to access the stability of ANN performance. In this work each ANN was trained 5 times.”

Comment: P7L6 – need reference for Bayesian Regularization

Response: We agree this would be useful for readers.

Edits: Test added: “In the implementation of Bayesian Regularization, a term is added to the sum of squared error cost function as a penalty for increased network complexity in order to guard against over fitting. A two level Bayesian inference framework is employed, operating on the assumptions the noise in the training data is independent, normally distributed, and also that all of the weights in the ANN are small, normally distributed, and unbiased (Hagan et al., 1997).”

Comment: The concepts of early stopping, hidden neurons, and hidden layers need to be described

Response: Thanks for this useful comment. Hidden neurons and hidden layers have been depicted in diagrams and described in more detail, embedded in the new text describing ANNs in general cited two comments above.

Edits: We have added some text to describe the concept of early stopping: “In preliminary ANN tests we found that over fitting occurred even when Bayesian Regularization was used, so we additionally implemented early stopping, which proved to be effective in the reduction of over fitting. To implement early stopping, a portion of training data is set aside as validation dataset, and during training, an ANN is applied to this validation data after each training step. Training continues so long as the error associated with the validation dataset is reduced. When the error associated with the validation dataset is no longer being reduced, training stops early. For ANNs, training datasets were divided in half on an alternating 24-hr basis, with half used for training and half used as validation data for early stopping.”

3 Results and Discussion

To evaluate the performance of each of the ANN and LM models that were generated using training data then applied to test datasets, we used residuals, the coefficient of determination (r^2), root mean squared error (RMSE), mean bias error (MBE), and centered root mean squared error (CRMSE). The CRMSE is an indicator of the distribution of errors about the mean, or the random component of the error. The MBE, alternatively, is an indicator of the systematic component of the error. The sum of the squares of the CRMSE and the MBE is equal to the square of the total error, the square root of which is defined by the RMSE. As in our previous work (Casey et al., 2017), we compared performance of LMs and ANNs with a number of different sets of inputs for each train/test data pair. The r^2 , RMSE, and MBE for each of these alternative models when applied to test data are presented in the supplemental materials (SM) in Fig. S2 through Fig. S7, along with representative scatter plots and time series comparing the performance LMs and ANNs for a given set of inputs. In Fig. S2 through Fig. S7, the best performing model inputs for each train/test data pair are shaded in purple. The type of model that performed the best (ANN vs. LM) is indicated in the caption of each figure. Presented below is an analysis and comparison of the best performing model for each species as determined in our previous work, as well as performance metrics for the best performing model associated with each new training/testing dataset pairs described in section 2.2.

Comment: Highlighted sentence is confusing as written. How can there be multiple 'best' performing models?

Response: Thanks to the Reviewer for the helpful comment.

Edits: We have added section 2.5 entitled "Calibration Model Evaluation and Testing in order to help clarify:

"To evaluate the performance of each of the ANN and LM models that were generated using training data then applied to test datasets, we used residuals, the coefficient of determination (r^2), root mean squared error (RMSE), mean bias error (MBE), and centered root mean squared error (CRMSE). The CRMSE is an indicator of the distribution of errors about the mean, or the random component of the error. The MBE, alternatively, is an indicator of the systematic component of the error. The sum of the squares of the CRMSE and the MBE is equal to the square of the total error, the square root of which is defined by the RMSE.

First, we generated and applied the best performing model, as determined in our previous work (presented in Table 3), to data from each new case study. Each new case study was selected to challenge models in different ways in order to evaluate the resiliency of the findings from our previous study when challenged by different circumstances. Next we tested LMs for CO₂ and O₃ that contained only the primary target gas sensor for each species, as well as temperature and absolute humidity as inputs. Finally, we generated, applied, and evaluated the performance of a number of LMs and ANNs with different sets of inputs for each case study in order to see which specific model performed the best for each individual case study. The r^2 , RMSE, and MBE for each of these alternative models when applied to test data are presented in the supplemental materials (SM) in Fig. S2 through Fig. S7, along with representative scatter plots and time series comparing the performance LMs and ANNs for a given set of inputs. In Fig. S2 through Fig. S7, the best performing model inputs for each train/test data pair are shaded in purple. The type of model that performed the best (ANN vs. LM) is indicated in the caption of each figure. We discuss both the performance of the previously determined best fitting model (generated using data from the GRET Spring 2017 case study) when applied and generated to data from new case studies, and the performance of models that were tuned to perform the best for each individual case study. From these comparisons, we draw insight into circumstances that challenge model performance in terms of relative local emissions characteristics, location, and timing between model training and testing pairs. Table 4 lists the relative timing and parameter coverage between model training and testing periods for dataset pairs, highlighting instances of incomplete coverage during training that led to model extrapolation during testing."

Comment: Does section 2.2 really succinctly describe each training/testing dataset pair? This is the first place in the text of the manuscript where the limited extent of co-location upon distributed field deployment is described and how the 10 UPODs are reconciled against such limitations.

Response: Thanks so much for pointing out that this is needed. Instead of describing why measurements were planned and carried out, we change the focus in section 2.2 to describe measurements and how they are used in this work.

Edits: We have added subsections to section 2.2, in which we describe each case study (training/testing dataset pair) in the context of the work presented here.

Section 3.1

Comment: For the purposes of the current study, if there is no co-location with reference, is it still a relevant data point? Can the authors effectively 'test' their model under these circumstances?

Response: Thanks to the reviewer for pointing out this confusion. We only have the ability to evaluate models when we have co-located reference instruments, and we only include data in this work that had co-located reference instruments.

Edits: We have added details in section 2.2 about how many U-Pods are included in each case study presented and which reference instruments were co-located with each.

Comment: This section P8L3 is also the first mention of reducing/oxidizing interfering gas species – this potential deserves a more detailed explanation in the context of the specific micro-environment source contributions

Response: Thanks for this important comment.

Edits: We have added a discussion of the operating principles of the sensors to section 2.1 accordingly, detailed in response to an earlier comment above.

Comment: The overall discussion of factors impacting differences between the two Basin deployments is fairly scattered. It would be more beneficial to the reader if the authors could draw more specific lines of connectivity between environmental or pollution source contributions and the robustness (or lack of robustness) in the model.

Response: Thanks very much to the Reviewer for this helpful comment.

Edits: We have improved the manuscript by describing differences between the gas basins in more detail, and in the context of sensor sensitivity and selectivity. Here is text from one place in the manuscript where we have made these improvements: “In this work, we present and compare models designed to address the unique challenges that come with using low-cost sensors, in the quantification of atmospheric trace gases of interest in oil and gas production basins, where ambient hydrocarbon mole fractions are potentially elevated, exerting uniquely confounding influence on low-cost gas sensors. We investigate how well models can be transferred from one microenvironment to another, with different dominant local emissions source characteristics, and different relative abundance of oxidizing and reducing compounds. Microenvironments explored in this work include an oil and gas basin where both natural gas and heavier hydrocarbons are produced (the DJ Basin), and an oil and gas production basin where prominently natural gas is produced (the SJ Basin), with much smaller proportional emissions of heavier hydrocarbons, and in turn, lower atmospheric concentrations of alkanes. Within and bordering the DJ Basin, additional microenvironments include an urban location, with significant mobile sources emissions (NO_x, CO, and VOCs), and a peri-urban site with fewer mobile emissions and closer proximity to oil and gas production activities. We explore how robust model performance is when a model is trained in one microenvironment and transferred to another; challenged by different relative abundance of oxidizing and reducing gas species. Additionally we test how well models can represent and address sensor stability over time and the potential for drift.”

17
The U-Pod CO₂ data presented in Fig. 3 and Fig. 4 were quantified using ANNs that were trained using data from the BAO Tower with the following inputs from each U-Pod: eltCO₂, temp, and absHum. This set of model inputs were found to be the best ANN inputs that we highlighted in our previous study, using data from the GRET site in the spring of 2017 (Casey et al., 2017). Fig. 3 shows scatter plots of U-Pod CO₂ vs. reference CO₂ during the test data period for sensors located at BAO as well as sensors that were located at distributed sampling sites throughout the SJ Basin. The scatter plots show that while there was generally a smaller dynamic range of CO₂ at the SJ Basin sites relative to BAO, model performance did not appear to be impacted or degraded by spatial extension to these locations in the SJ Basin. The line of best fit for Fort Lewis site (periwinkle) is even closer to the 1:1 than the lines of best fit for two U-Pods located at BAO (black and grey). Overlaid histograms of residuals in the bottom right corner of Fig. 3 show that CO₂ residuals from each of the SJ Basin U-Pods are generally centered and evenly distributed about zero with similar spread.

20
U-Pod CO₂ average residuals from the same data presented in Fig. 3, quantified using ANNs with eltCO₂, temp, and absHum signals as inputs, are plotted according to time of day and date in Fig. 4. While the use of ANNs in place of LMs was shown to reduce U-Pod CO₂ residuals significantly with respect to temperature, some daily periodicity in the residuals for all four U-Pods is apparent in the upper plot in Fig. 4 that shows residuals by date. The lower plot in Fig. 4, showing residuals by time of day, demonstrates that CO₂ from all four U-Pods was generally under predicted during early hours of the morning and generally over predicted during afternoon and evening hours. Interestingly, this trend in residuals by time of day is more pronounced for the two U-Pods that remained at BAO. The majority of U-Pods stopped logging data, unfortunately, at one point or another during these deployments. The periods of missing data are reflected in the plots of

8

Comment: eltCO₂, temp, absHum should be human readable, this is the first time these parameters appear in the text. I understand that they were listed in the table describing UPOD guts, but they should be spelled out here.

Response: Thank you for the feedback.

Edits: Descriptions added here and at the first mention of other model input codes in the manuscript in the text: “eltCO₂ (ELT S300 CO₂ sensor) , temp (temperature) , and absHum (absolute humidity)”

Comment: L18: it is unclear to what extent the current work and the previous work are duplicated here? Does the previous work form the basis for determining the optimal set of input parameters to train the ANN model and those same set of input parameters were found to be optimal again in this second application or are the actual applications overlapping and therefore the result is redundant? This is an example where I find the self-referential context to Casey et al., 2017 confusing (and lacking specific differentiating information).

Response: Thank you for the useful comment. We applied the model that we found to perform best in our previous work to new data. The application circumstances did not overlap and are not redundant.

Edits: We have added the following text to help clarify: “We began by testing the best-performing CO₂ model, as determined in our previous work (Casey et al., 2017), on this data, collected under a different set of circumstances, during the summer of 2015.”

Comment: The under-prediction / over-prediction behavior of all four UPODs warrants more discussion. What environmental conditions are pushing the model beyond its limits? What is the fundamental (under-the-hood) reason for the interference in the first place (based on sensor fundamentals)?

Response: Thank you for this interesting comment. After analysis and careful consideration, we have added the following text:

Edits: “Upon examination of overlaid histograms showing distributions of parameters during model testing and training periods, in Fig. S12, and model time series and residuals plots in Fig. S3, there is no indication of model extrapolation at the BAO site, and no significant trends of concern with respect to residuals. Bias introduced to mole fraction estimates are likely attributable to differences in hydrocarbon mixtures in the SJ Basin relative to the DJ Basin.”

Comment: Why did the majority of UPODs stop logging data during the deployment? Did the system overheat? What fraction of the total possible sample time was missed?

Response: Thank you for the comment and helping us to improve the details of the study.

Edits: The following text has been added accordingly: “Periods of missed data during the month-long deployment included approximately 1 day at the Shiprock site, 2 days at the Bloomfield site, 4 days at the Sub Station site, 9 days at the Fort Lewis site, and 17 days at the Navajo Dam site. We carried out frequent sampling site visits (on a weekly or biweekly basis as logistics and travel to remote locations in some cases allowed) in order to identify and fix problems as they arose during field deployments. Operational issues were predominantly attributable to power supply problems associated with BNC bulkhead fittings and corrupted micro SD cards.”

Section 3.1 continued..

5 Differing from our previous findings, for this group of training and testing data pairs from the summer of 2015 at the BAO and SJ Basin sites, the inclusion of the e2vVOC and alphaCO signals noticeably improved the RMSE in the quantification of CO₂. While the inclusion of these two secondary sensor signals didn't result in the best performance in our previous study, using data from the GRET site (Casey et al., 2017), we found that their inclusion did not degrade performance relative to the models that included just eltCO₂, temp, and absHum signals as inputs. Generally, using rh vs. absHum signals as ANN inputs did not seem to make a big difference in model performance, though linear models were sometimes found perform
10 better when the absHum signal is used instead of the rh signal. From Fig. S2, it is apparent that inputs including e2vCO₂, temp, rh, e2vVOC, and alphaCO sensor signals as model inputs resulted in the lowest RMSE for U-Pods at BAO as well as at the two SJ Basin sites. Plots analogous to those presented in Fig. 3 and Fig. 4, but with this best performing set of inputs for the present data set pairs are presented in the SM, in Fig. S24 and Fig. S25 respectively.

15 O₃ was quantified for all the U-Pods deployed at BAO and SJ Basin sampling sites using an ANN with the following inputs: e2vO₃, temp, absHum, e2vCO, e2vVOC, figCH₄, and figCxHy. ANNs with this configuration were found to perform best in the quantification of O₃ in our previous study (Casey et al., 2017). These same inputs and model configuration were also found to be the best performing among others tested for SJ Basin 2015 dataset pairs as noted in Fig. S2. Interestingly though, LMs with this same set of inputs performed competitively well for a number of the U-Pods in the SJ Basin. For
20 three of seven U-Pods in the SJ Basin, LMs even outperformed ANNs in terms of RMSE. When the BAO trained U-Pods field calibrations for O₃ were extended to sites in the SJ Basin, we found that U-Pods at some of the sites performed better than others across all models that were tested, as seen in Fig. S2.

Scatter plots and trends in residuals are presented in Fig. 5 and Fig. 6 for O₃. These plots show the performance of U-Pods at
25 BAO relative to those at SJ Basin sites in the quantification of O₃ during the test data period. U-Pod O₃ measurements at Fort Lewis, Navajo Lake, and the Sub Station did not agree with reference measurements as well as U-Pod O₃ measurements from the other four SJ Basin sites. U-Pods at Navajo Lake and Sub Station had bad humidity sensor data, as noted in section 6.2.3 and Table S1, so humidity from the U-Pod located at the Ignacio site was used in place of their humidity signals. Since the Ignacio site was located relatively far away from the Navajo Dam and Sub Station sites, this could have introduced some
30 additional error into the application of a calibration equation, particularly since we showed earlier that O₃ ANNs like the ones we employed here are very sensitive to humidity inputs (Casey et al., 2017). The Fort Lewis site had a different reference instrument than those used at other sites, which may have contributed to observed discrepancies. Fig. S1 shows that differences among U-Pod O₃ performance during the test deployment period were larger than those observed during the training phase among the same U-Pods; therefore, the incongruous field calibration performance phenomena we observed

Highlighted above:

Comment: L6-9: Discussion is confusing and language is too casual: “did not make a big difference” – too vague. Quantify based on the statistical analysis of the model test data. When considering the

benefit of including extra sensor inputs in the training matrix for their models, again the Authors are drawing comparisons to their earlier work (Casey et al, 2017) but it's not really clear how this improves/informs the current work – besides stating that the inclusion of the parameters didn't make the data product worse.

Response: Thank you for this useful comment. With this work, we are testing methods that we developed in our previous work under new circumstances that have the potential to challenge and degrade model performance. The finding we are highlighting in this instance is that in the current work, two additional sensor signals result in improved performance of a model under different circumstances, relative to our previous work. Since the addition of these two signals do not reduce the performance of models in our previous work, the addition of these two sensor signals in models for the quantification of CO₂ may be warranted more broadly.

Edits: We have changed 'did not make a big difference' to 'did not have a measurable affect'. Additionally we have added the following text: "so including these sensor signals may be appropriate as a general rule, in areas that are strongly influenced by oil and gas production activities."

Comment: L10 e2vCO₂ does not exist as a sensor metric in the UPODs.

Response: Thanks so much to the Reviewer for catching this mistake. It should be e2vCO.

Edits: We have changed 'e2vCO₂' to 'e2vCO'.

Comment: L15: 'all the UPODs' how many is this again?

Response: Thank for the clarifying comment.

Edits: The following edits have been made: "O₃ was quantified for the 2 U-Pods deployed at BAO and 7 of the U-Pods deployed at SJ Basin sampling sites"

Comment: L19: 'For a number of UPODs': state the number.

Response: Thanks to the reviewer for helping us clarify.

Edits: The text has been edited accordingly: "Interestingly though, LMs with this same set of inputs performed competitively well for 3 of the 7 U-Pods in the SJ Basin in terms of RMSE and r²"

Comment: L21: 'for some of the sites.': which sites?

Response: Thanks to the Reviewer for helping us clarify.

Edits: The following edits have been made to the text: "When the BAO trained U-Pods field calibrations for O₃ were extended to sites in the SJ Basin, we found that U-Pods at the Bloomfield, Bondad, Shiprock and Ignacio sites performed better than others across all models that were tested, as seen in Fig. S2."

Comment: L15-22: this paragraph seems to say that the ANN training matrix determined to be optimal in Casey et al., 2017 was also found to be optimal in the current work, with inclusion of all peripheral sensors to the input training matrix for O₃. But they also state that the LMs data products were just as good (or better) when compared to the ANN models. This result seems important, but not really discussed further. The results are left vague. Conclusions as to why this might be the case are absent.

Response: Thank you very much for helping us to make our conclusions more detailed and less vague.

Edits: We have added the following text accordingly: "The observation that LMs performed competitively well at a subset of SJ Basin sites is likely connected to the relative abundance of hydrocarbons and other potentially interfering oxidizing and reducing gas species at individual sampling sites, diverging from conditions present during model training at the BAO site. ANNs can better represent the influence of these interfering species than LMs during training, but appear to have lost their ability to do so for this subset of microenvironments in the SJ Basin."

Comment: L27: 'had bad RH data' – as noted in a section that doesn't exist. What is bad RH data?

L29: 'relatively far away' – how far? Again. These details matter.

Response: Thank you for helping us clarify the text and catching the error regarding the section referenced.

Edits: We have made the following edits to the text: “As noted earlier, U-Pods at the Navajo Dam and Sub Station sites had faulty relative humidity sensor data, so humidity from the U-Pod located at the Ignacio site was used in place of their humidity signals. Since the Ignacio site was located approximately twenty-two and fifty miles away from the Navajo Dam and Sub Station sites respectively, this could have introduced some additional error into the application of a calibration equation, particularly since we showed earlier that O₃ ANNs like the ones we employed here are very sensitive to humidity inputs (Casey et al., 2017). Spatial variability in humidity across tens of miles could be significant as isolated storms (which are on average 15 miles in diameter) propagate throughout the region in the summer.”

Comment: L30-31: Apparently one of the major results from Casey et al., 2017 is an extreme sensitivity to RH when using ANN's to quantify O₃. Given the failure of the RH sensor throughout much of the work presented in the current work, it seems critically important that this RH sensitivity be discussed in much greater detail in the current work, not simply stated in an off-handed matter with a reference to the prior work.

Response: Thanks very much to the Reviewer for this helpful comment. We have added significant detail throughout the text describing humidity influences on sensors in the context of model development and testing.

Edits: Here is an example of some text we have added accordingly in section 3.2: “In our previous work, we showed that O₃ models were very sensitive to the humidity signal input (Casey et al., 2017). In this case study, it seems that replacing actual humidity signals with closely approximated humidity signals, negatively influenced model performance. In order to investigate this observation further, we tested the influence of replacing humidity data in the same manner, using mixing ratios from the same co-located Picarro, on test data from the GRET Spring 2017 case study. A comparison of model performance under normal and this ‘borrowed RH’ circumstance are presented in Fig. S27 in the SM. O₃ model performance was negatively impacted when ‘borrowed’ RH values based on Picarro data replaced U-Pod RH sensor signals. From these findings, it seems likely that the inclusion of multiple metal oxide type sensors as inputs in the model, which all respond strongly to humidity fluctuations, helped the ANN to effectively represent the influence of humidity in the system, more so than including a ‘borrowed RH’ signal from another instrument. We tested models with multiple gas sensor signals and no humidity signal as inputs for a number of other case studies as well (as seen in Fig. S2, Fig. S4, and Fig. S5), when good humidity data from U-Pod enclosures was available, but they did not turn out to be the best performing model in any of these other tests.”

Comment: L32: ‘had a different reference instrument’ what was the instrument and why do the authors think that this particular reference instrument was in error, subsequently disrupting the validity of their calibration model?

Response: Thank you for the useful comment. We only want to acknowledge that discrepancies among different reference instruments that are operated according to different protocols and by different agencies are possible.

Edits: The following text has been added to help clarify: “At the Fort Lewis site, a 2b Technologies model 202 O₃ analyser was employed as a reference instrument, differing from the Thermo Scientific 49i, Thermo Scientific 49is, and Teledyne API T400 instruments utilized for reference measurements, elsewhere in the SJ Basin, and the Thermo Scientific 49c that was operated at the BAO site and used for model training. Of all the reference instruments, only the 2b Technologies model 202 O₃ at the Fort Lewis site was operated in a room that was not temperature controlled. Some bias may have been introduced to the Fort Lewis O₃ reference measurements as the temperature in the room it was housed in varied. Different instruments, operators, calibration and data quality checking procedures could have contributed to observed discrepancies. It is also possible that the microenvironment at each of these three sites contributed lower model performance.”

Comment: L34 – carried into highlighted passage below: The authors indicate that the sampling sites or the circumstances discussed previously are the reason for the poor model performance, not the sensors comprising the UPODs. First, WHAT circumstances specifically, and what specifically about the sampling sites? This level of non-explanation is unacceptable.

Response: Thanks to the Reviewer for helping us clarify.

Edits: The following edits have been made to the text: “therefore, the incongruous field calibration performance phenomena we observed seems to be connected to unique characteristics associated with individual sampling sites; possibly the relative abundance of oxidizing and reducing molecules in the local atmosphere, which could interfere with sensor responses to their target gas species, as opposed to the quality of individual sensors in each of those U-Pods.”

seems to be connected to the sampling sites or the circumstances discussed previously as opposed to the quality of individual sensors in each of those U-Pods.

- 5 All SJ Basin U-Pod O₃ measurements systematically over estimate lower levels of O₃ each night, a trend apparent in the scatter plots in Fig. 5 and in the residuals by time of day plot in Fig. 6. Upon examination of the scatter plots in Fig. 5, U-Pods at some sampling sites had positive bias for higher O₃ measurements as well (Shiprock, Ignacio, Sub Station, and Bloomfield), while for others, bias at the higher end of O₃ distributions did not appear to be present (Navajo Dam, Fort Lewis, and Bondad). The residuals by time of day plot in Fig. 6 shows that the two U-Pods at BAO did not have significant trends in their residuals according to the time of day, but that U-Pods deployed at SJ Basin sites consistently over estimated
- 10 nighttime O₃. The residuals are also plotted with respect to temperature in Fig. 6, where all U-Pods, even those at BAO to a lesser extent, appear to over predict O₃ at lower temperatures, which generally occurred at night. The times of day that generally correspond to the highest O₃ levels generally had the lowest residuals, with some exceptions at the Fort Lewis and Navajo Dam sites.
- 15 Fig. 6 includes a plot of the residuals across the duration of the deployment period, showing no significant sensor drift in measurements for any of the U-Pods. This plot also shows that the highest residuals observed generally occurred over short periods in time, particularly for the Fort Lewis (blue) and Sub Station (magenta) sites. In order to further explore the performance of field calibration models for O₃ at SJ Basin sites relative to BAO, the combined parameter space of temperature with O₃ reference mole fractions and temperature with absolute humidity are presented in Fig. 7. The combined
- 20 temperature and reference O₃ parameter space appears to be similar for all of the U-Pods, both at BAO and the SJ Basin sites. However, there appears to be some outlying combined temperature and humidity parameter space at the Sub Station site and at the Navajo Dam site. Brief excursions of high humidity may be connected to some of the large short-term residuals observed at these two sampling sites.

Comment: L22: brief excursions of high humidity – how brief? How high?

Response: Thank you for helping us improve clarity.

Edits: The following details have been added: “Brief excursions, lasting approximately 2 – 4 hours, of high humidity (up to 0.025 kg/kg, relative to the upper bound of absolute humidity observed at other sampling sites of 0.013 kg/kg) may be connected to some of the large short-term residuals observed at these two sampling sites.”

Comment: Can the authors comment on the role that humidity transients play in fundamental sensor response? The description of the high and low bias resulting from the models at different locations and different times of day is difficult to follow. What are the common response characteristics and failings of the model that manifest across the case studies featured here? What are the lessons learned and how can these lessons better inform ANN model development moving forward?

Response: Thanks very much to the Reviewer for this helpful comment.

Edits: We have added the following details about fundamental sensor response in section 2.1, including the role that humidity transients play.

Section 3.2.1

Comment: Extrapolation of the ANN and LM models is problematic. Why? If the full-span of O₃ (or CO₂) concentration encountered in the field deployment is not covered in the training set for the model, is the model incapable of reasonably extrapolating?

Response: Yes, thank you for the helpful comment. The Dawson Summer 2014 case study suggested that, when a model is transferred to a new location, with different dominant local emission sources, both ANNs and LMs fail to extrapolate effectively with respect to high O₃ mole fractions.

Edits: We have added the following text has been added accordingly: "Across applications, ANNs have been found to be unreliable when extrapolated, due to the nonlinear nature and complexity of the relationships they represent. Though they are generally expected to be more robust to extrapolation than ANNs, increased uncertainty in measurements can also be introduced to LMs when parameters are extrapolated. In order to have high confidence in measurements of uncommonly high mole fractions of a target gas, the model -raining period has to encompass the full possible range. Combining both field calibration and lab calibration data together in a training dataset could accomplish this type of coverage. If extrapolation is expected to occur with respect to the target gas mole fraction, as in this case study, the use of an inverted LM may yield better results than LMs or ANNs. We describe inverted LMs and their potential advantages in our previous work (Casey et al., 2017)."

Section 3.2.2

Comment: L15: post-test deployment co-locations: It's unclear what is meant by 'post-test', please clarify.

Response: We are happy to clarify this concept.

Edits: The following text has been added to the end of section 2.3: "A model was extrapolated in time when ever training data does not take place both before and after a given test deployment period. In several case studies we present, model training only took place after the test deployment period, comprising a 'post only' calibration. In Colorado, and more broadly in the western United States, ambient temperatures change significantly across the seasons throughout the year, so if a model is extrapolated in time, extrapolation in temperature often results as well."

Comment: L16: state the # of UPODs

Response: Thank you for helping us to add clarity.

Edits: The following edits have been carried out: "We present data from four U-Pods that were co-located with reference instruments in the SJ Basin in the spring of 2015, at the Navajo Dam, Sub Station, and Bloomfield sites. Two U-Pods at the Bloomfield site provide a set of duplicate measures."

Comment: The concept of extrapolation in time is confusing. Please clarify what is meant by this? Generating a model at time X and then applying that same model to time X-Y?

Response: Thank you for helping us to clarify what is meant by extrapolation in time.

Edits: The edits have been made in the manuscript in section 2.3: "A model is extrapolated in time when ever training data does not take place both before and after a given test deployment period. In this case study, model training only took place after the test deployment period, comprising a 'post only' calibration. In Colorado, and more broadly in the western United States, ambient temperatures change significantly across the seasons throughout the year, so if a model is extrapolated in time, extrapolation in temperature often results as well."

Comment: The authors identify coal-fired power plants as an important near-field ('close-by') pollutant source that could contribute a specific (unique) pollutant signature that could render the utility of the Figaro sensor useless. Did the CO₂ response of the UPODs or reference instruments or CO response of the sensor measurements indicate a near-field power plant plume across the deployment area?

Response: Thanks to the Reviewer for this useful comment. We did observe evidence of a near-field power plant plume in the raw CO₂ and CO sensor signals as well as the NO and NO₂ reference measurements (the site was not equipped with a CO reference instrument).

Edits: We have added the following text accordingly: "Several-hour long enhancements or spikes are apparent in the raw eltCO₂ and alphaCO sensor signals in the U-Pod deployed at the Sub Station site, indicating the presence of a near-by combustion-related emissions source. Another indication of indicate a near-field power plant plume across the deployment area is apparent, in the form of several-hour long enhancements reference measurements of NO and NO₂ at the site."

Section 3.2.3

Comment: How specifically was 'time' included as a raw input vector in the training matrix? Absolute time? Time since start of deployment? Time since calibration? Time since sensor manufacture?

Response: Thank you for the helpful comment.

Edits: We have added the following text to the end of section 2.3 to help clarify, since the time model input is discussed there first: "When time was included in a model as an input, the absolute time was used. Specifically, we used the datenum value from the MATLAB environment, which is defined by the number of days that have elapsed since the start of January 1st, in the year 0000."

Comment: L11-12: "...LMs outperformed ANNs with notable instability associated with the performance of ANNs when time was included as an input." In the previous sentence the authors stated that time was useful predictor of CO₂.. but the last sentence appears to contradict this assertion. The fact that LMs outperformed ANNs for CO₂ also contradicts general assertions made in the abstract.

Response: Thanks very much to the Reviewer for making these important points.

Edits: We have added the following text to section 3.2.3 to help clarify: "In the case of CO₂, LMs outperformed ANNs, which could be largely attributable to notable instability associated with the performance of ANNs when time was included as an input."

We have also added the following text to the abstract to help clarify: "For CO₂ models, exceptions included, case studies in which training data used took place more than several months subsequent to the test data period. For O₃ models, exceptions included studies in which the characteristics of dominant local emissions sources (oil and gas vs. urban) were significantly different at model training and testing locations."

Comment: The authors should comment on the notion that time-sensitive response patterns in sensors indicates that some level of time-decay. Is this the case with the CO₂ sensor and that's why time as a input parameter in the model makes such a big difference? Is there some fundamental reason why the ANNs would be poorly suited to model time-decay patterns in the sensors?

Response: Thanks very much to the reviewer for this suggestion as well as interesting and relevant questions.

Edits: We have added the following text to address each: "For CO₂, we expected the inclusion of time as an input to be a useful to model performance across this time frame, owing to observed trends of decreased CO₂ sensor sensitivity in time. To keep the power requirements for the U-Pod sensor systems low, and to keep systems quiet, fans were used to exchange air in the enclosures as opposed to pumps. As a result, the air entering the enclosures was not filtered, and sensors were exposed to some dust over time. This dust exposure is likely largely responsible for observed decreases in CO₂ sensors sensitivity over time, shown in Fig. S26. Decreases in infrared lamp intensity over time may also play a role. In the case of CO₂ sensors, the implementation of pumps to draw new, filtered air into sensor enclosures could likely significantly reduce lose rates in the sensitivity of an individual sensor over periods of continuous deployment in ambient environment. While we are not sure why ANN performance tended not to benefit from the addition of a time input, while LM performance did, it is likely attributable to the extrapolation of the time input, since only data that was collected significantly subsequent to the test data period was used for training. ANNs are expected to be able

to better represent time decay trends if data from measurements both prior and subsequent to the test period are used in training, so that there is no extrapolation with respect to the time input.”

Section 3.2.4

Comment: L23-24 – final sentence in this section is very important. Where the faulty RH (and necessity of substituting RH from alternate sources) degraded the models, if enough RH variability was captured with the suite of peripheral metal oxides sensors, the RH-interference could be effectively modeled without explicit RH inputs. It would seem important to emphasize this point a bit more prominently and discuss further – especially in the context of overcoming some of the RH-measurement shortfalls elsewhere in the manuscript through similar means.

Response: Thanks very much for this helpful comment. We found this to be an interesting result also.

Edits: We have added the following text accordingly: “We tested models with multiple gas sensor signals and no humidity signal as inputs for a number of other case studies as well (as seen in Figures S2, S4, and S5), when good humidity data from U-Pod enclosures was available, but they did not turn out to be the best performing model in any of these other tests.”

4. Conclusions

Comment: Supervised learning techniques – generally, the manuscript lacks a description of what is meant by this -

Response: Thanks very much to the Reviewer for pointing this out.

Edits: We have added the following text to the introduction and the conclusions to help clarify that ANNs are an example of a supervised learning method, as are random forests: “We investigated how well a supervised learning technique (ANNs) hold up when sensors are moved to a new location, different from where calibration model training took place.”

Comment: L19-20 the concepts of temporal and spatial extension are still a bit confusing here. Earlier statements to clarify exactly what is meant by each condition would be helpful.

Response: Thanks to the Reviewer for pointing out this confusion.

Edits: We have added the following text, early in the manuscript, at the end of section 1.4: “In the present work, we test model performance under conditions of spatial extension, wherein a model is trained using data from one location then applied to a test dataset using data from a new location. In testing spatial extension of a model we investigate how well the field calibration of low-cost sensors can inform target gas mole fractions when sensors are deployed in a new location and a new microenvironment of oxidizing and reducing compounds. We also test model performance under conditions of temporal extension, wherein a model is trained using data that was collected only prior or subsequent to the model application period. In testing temporal extension of models, we investigate how model performance is influenced by sensor drift over time.”

Comment: L24: how does one move something in terms of its temporal coverage?

Response: Thanks to the Reviewer for pointing out that this statement is confusing and unclear.

Edits: We have updated the text accordingly: “While ANNs and other supervised learning techniques have been shown to consistently outperform linear models in previous studies when training and testing took place in the same location, we find that this trend does not always hold when field calibration models are applied in a new location, with significantly different local emissions source signatures for O₃ models, or when model training data takes more than several months subsequent to the model application period for CO₂ models.”

Comment: L1-3P16: LMs appear to be more robust when applied to a changing deployment condition – but then the authors hedge and say that they “... were not able to fully represent some of the complex nonlinear response behavior exhibited by the arrays of sensors.” So a linear model can’t model nonlinear behavior? The statement needs to be more specific.

Response: Thanks to the Reviewer for pointing out the vague nature of the statement. After some consideration, we realize that a more important point to make at the end of this paragraph has less to do with nonlinear response behavior, and more to do with extrapolation of observed ozone mole fraction.

Edits: We have updated the text accordingly: “While these LMs seemed to be more stable under circumstances of significant extrapolation in terms of local air chemistry and timing, we found that they did not extrapolate well in terms of the O₃ mole fraction, resulting in underproduction of O₃ values during the test period that exceeded those encompassed in the training data.”

Comment: L7: “..data is almost a band running vertically in a range of CRMSEs.” Data running in ‘a band’ doesn’t aid in the interpretation of the data. Re-phrase to address the statistical product that results from the bias that was encountered.

Response: We agree re-phrasing this statement in terms of statistical attributes will help clarify.

Edits: The text has been updated accordingly: “As seen in Fig. 12, plot markers from all case studies have very similar CRMSE values, but plot markers from case studies in which models were tested in new locations have larger MBE values than models that were tested in the same location as they were trained.”

Comment: Final paragraph: how ‘generalizable’ are the models developed here? It would seem that despite having done an exhaustive amount of work, each individual UPOD system still required its own ANN or LM based on co-located data and raw sensor data from that individual sensor system. While the input matrix of raw sensor signals may be more generalizable, the models themselves appear to be very much node-specific, at least in so far as what has been shown in the paper.

Response: Thanks to the reviewer for highlighting this important point. We have added text to help address it.

Edits: Text added: “In order to account for unique variations in sensor responses, in each individual sensor system, due to variations in manufacturing along with elapses time and specific exposure subsequent to manufacturing, we present models that are generated for each sensor system on an individual basis. Future studies exploring the potential for universal calibration models would be very useful to the field.”

Comment: It is unclear how the extension of the model frameworks discussed in the current paper can be used in the context of low-cost electrochemical sensors

Response: Thanks to the Reviewer for this very important comment. We have added five key take away points from this work and associated recommendations that we hope can be used by others in the field of low-cost gas sensors.

Edits: “The following findings from this work, and associated recommendations, are made to help inform the logistics of future studies that employ field calibration methods of low-cost gas sensors.

1. **Finding:** For O₃ models, LMs perform better than ANNs when the chemical composition of local emissions sources is significantly different in the model-training location relative to the model-application location. We found that when models were trained in an urban area with significant mobile sources, then tested in a peri-urban area, more strongly influenced by oil and gas emissions, the differences in local sources of pollution were significantly different enough that LMs outperformed ANNs. Alternatively, when models were trained in one oil and gas production region and tested in another the different composition of local emissions (lighter vs. heavier hydrocarbons) was not significant enough for LM performance to surpass the performance of ANNs, though some positive bias was evident in predicted O₃ mole fractions.

Explanation: ANNs are very effective at compensating for the influence of interfering gas species through pattern recognition of a training dataset. However, if different patterns, in terms of the relative abundance of various oxidizing and reducing compounds in the air, are present in the testing location relative to the training location, ANNs may not be able to compensate for the influence of interfering gas species as effectively. The relative

abundance of interfering oxidizing and reducing compounds are not included as model parameters, but ANN performance is challenged by these circumstances.

Recommendation: When measuring O₃ or other gas species with a metal oxide type sensor, if the nature of dominant emissions sources at the model training location is significantly different than the nature of dominant emissions sources in the model application location, use an LM instead of an ANN. For the best performance, try to train models in locations with similar emissions sources to a desired sampling location. If the nature of dominant emissions sources at the model training and application locations are similar, signals from an array of multiple unique metal oxide sensors will likely augment model performance.

2. **Finding:** For CO₂ models, LMs perform better than ANNs when model training occurs significantly (more than several months) prior to or subsequent to the model application period.

Explanation: CO₂ sensors drift over time in terms of sensitivity and baseline response. When models are extrapolated in time (when training takes place more than several months prior or subsequent to the model application period), ANN performance can be compromised to a greater extent than LM performance because ANNs are able to represent relationships during training very effectively, and with significant more complexity and nonlinear relationships among time and other model inputs than LMs. The more complex the model, the less likely it can be extrapolate effectively. LMs, with no interaction terms like we employ in this work, are not able to fit data and potentially complex patterns inherent in sensor drift over time during training as closely as an ANN, but the simple linear relationships they represent between the time input and the target gas mole fraction over the course of training are more likely to hold prior or subsequent to the training period.

Recommendation: When measuring CO₂ with a NDIR sensor, if model-training data is only available more than several months prior or subsequent to the model application period, use a LM instead of an ANN. For the best model performance, use training data that is collected directly pre or post of the model application period, and preferably data from both pre and post of the model application period. Training models using data from both pre and post of a given model application period helps models to encompass sensor drift over time as well as increases the likelihood of covering the full range of environmental parameter space that occurs during the model application period so that extrapolation of these parameters is avoided.

3. **Finding:** Extrapolation of an O₃ or CO₂ model in time, and especially significant extrapolation in time, can change both the type of model that is most effective, as well as the specific model input signals that are most effective.

Explanation: Low-cost sensors change over time, both in terms of their baseline response and in terms of their sensitivity to target and interfering gas species. Different sensor types drift due to different physical phenomenon so further a generalization across sensor types is difficult.

Recommendation: Use training data collected directly pre and post of the model application period in order to implement a 'best performing model' for each gas species that can be applied using data from different model training and application pairs.

4. **Finding:** ANNs yield less bias and more accurate gas mole fraction quantification than LMs, even when transferred to a new location under the following circumstances: when extrapolation of training parameters is avoided during the model application period, when training takes place for several weeks to a month prior and subsequent to the model application period, and when the dominant local emissions sources are similar in the model training and application locations.

Explanation: Our previous study and multiple other ambient and laboratory based experiments have shown, arrays of low-cost sensors in combination with ANN regression models can support useful quantification of gases in mixtures and in the ambient environment because ANNs can more effectively represent complex nonlinear relationships

among environmental variables and signals in a sensor system like a U-Pod than LMs. With this work, we have explored limitations associated with these methods when challenged in different ways, as we present with a number of case studies.

Recommendation: If minimizing error and bias in measurements of gas mole fractions using low-cost sensors systems is a primary goal, design sensor system training and field deployment experiments so that extrapolation of model training parameters is avoided during the model application period, so that training takes place for several weeks to a month directly prior and directly subsequent to the model application period, and so that the dominant local emissions sources are similar in the model training and application locations. When these conditions are satisfied, ANNs can be robustly implemented, with better performance than LMs.

It is also imperative that sensor users keep in mind the primary importance of minimizing extrapolation of temperature, humidity and sensor signal from model training to application.”

Testing the performance of field calibration techniques for low-cost gas sensors in new deployment locations: across a county line and across Colorado

Joanna Gordon Casey¹, Michael P. Hannigan¹

5 ¹Department of Mechanical Engineering, University of Colorado at Boulder, Boulder, 80309, United States of America

Correspondence to: Joanna Gordon Casey (joanna.casey@colorado.edu)

10 **Abstract.** We assessed the performance of ambient ozone (O₃) and carbon dioxide (CO₂) sensor field calibration techniques when they were generated using data from one location and then applied to data collected at a new location. This was motivated by a previous study (Casey et al., 2017) which highlighted the importance of determining the extent to which field calibration regression models could be aided by relationships among atmospheric trace gases at a given training location, which may not hold if a model is applied to data collected in a new location. We also explored the sensitivity of these methods in response to the timing of field calibrations relative to deployments periods. Employing data
15 from a number of field deployments in Colorado and New Mexico that spanned several years, we tested and compared the performance of field-calibrated sensors using both linear models (LMs) and artificial neural networks (ANNs) for regression. Sampling sites covered urban, rural/peri-urban, and oil and gas production influenced environments. We found that the best performing model inputs and model type depended on circumstances associated with individual case studies, such as differing characteristics of local dominant emissions sources, relative timing of model training and application, and the extent of extrapolation outside of parameter space encompassed by model training. In agreement with findings from our previous study that was focused on data from a single location (Casey et al., 2017), ANNs remained more effective than LMs for a number of these case studies but there were some exceptions. For CO₂ models, exceptions included, case studies in which training data collection took place more
20 than several months subsequent to the test data period. For O₃ models, exceptions included studies in which the characteristics of dominant local emissions sources (oil and gas vs. urban) were significantly

Joanna Gordon 8/15/18 11:07 AM

Deleted:

Joanna Gordon 8/15/18 11:07 AM

Deleted: We also explored the sensitivity of these methods to the timing of field calibrations relative to deployments they are applied to.

Joanna Gordon 9/30/18 12:48 PM

Deleted: WeGenerally,

Joanna Gordon 8/27/18 11:13 AM

Deleted: w

Joanna Gordon 9/30/18 12:48 PM

Deleted: collection

different at model training and testing locations. Among models that were tailored to case studies on an individual basis, O₃ ANNs performed better than O₃ LMs in 6 out of 7 case studies, while CO₂ ANNs performed better than CO₂ LMs in 3 out of 5 case studies. The performance of O₃ models tended to be more sensitive to deployment location than to extrapolation in time while the performance of CO₂ models tended to be more sensitive to extrapolation in time than to deployment location. The performance of O₃ ANN models benefited from the inclusion of several secondary metal oxide type sensors as inputs in 5 of 7 case studies.

1 Introduction

In places like the Denver Julesburg (DJ) and San Juan (SJ) Basins, along Colorado's Front Range and in the Four Corners Region, oil and gas production activities have been increasing with the advent of horizontal drilling that can be effectively used in conjunction with hydraulic fracturing to produce hydrocarbons from unconventional geologic formations. Public health concerns have arisen about the increasing number of people living alongside these industrial activities and emissions (Adgate et al., 2014; McKenzie et al., 2014; McKenzie et al., 2012, 2017). We previously developed methods to quantify ozone (O₃), carbon dioxide (CO₂), methane (CH₄), and carbon monoxide (CO) using low-cost gas sensors in an area where the ambient mole fractions of these species are influenced by oil and gas production activities (Casey et al., 2017). Such low-cost sensor measurements could enable greater understanding of air quality in oil and gas production basins, informing the spatial and temporal scales that people live and work in a way that current technologies used by regulatory agencies cannot feasibly accomplish. In our previous work, we tested and compared the performance of direct and inverted linear models (LMs) as well as artificial neural networks (ANNs) as regression tools in the field calibration of low-cost sensor arrays to quantify these target gas species using month-long test datasets, training each model with approximately one month of data prior to and one month of data subsequent to this test period. ANNs are powerful pattern recognition tools. They were found to perform better than both inverted and direct LMs in our previous study, but concerns arose when findings suggested that the performance of ANNs was being augmented by the relationships among gas mole fractions in the atmosphere at a given location. Low-cost gas sensor systems have the potential to inform spatial and

Joanna Gordon 9/2/18 11:39 AM

Deleted:

Joanna Gordon 9/4/18 11:23 AM

Deleted: In almost all cases the best CO₂ models were ANNs that only included the NDIR CO₂ sensor along with temperature and humidity.

Joanna Gordon 9/30/18 12:48 PM

Deleted: thanthat

Joanna Gordon 9/4/18 11:23 AM

Deleted: many

Joanna Gordon 9/4/18 11:24 AM

Deleted: cases.

Joanna Gordon 9/2/18 11:43 AM

Deleted: We

temporal variability of pollution. Calibration equations for each sensor system can be generated in one location based on co-located measurements with reference instruments, and then the sensor systems can be moved into a spatially distributed network. Since the relationships among gas mole fractions will differ at different sampling sites across a spatially distributed network, calibration models may not hold at new sampling sites. In this work, we test calibration model performance when extended to new locations.

Joanna Gordon 9/30/18 12:47 PM

Deleted: will differ

1.1 Low-Cost Sensors for Air Quality Measurements

The use of low-cost metal oxide, electrochemical and non-dispersive infrared sensors to characterize air quality is becoming increasingly common across the globe (Clements et al., 2017; Kumar et al., 2015).

While low-cost sensors have been emerging on the market with sufficient sensitivity to resolve variations in ambient mole fractions of target gases of interest, they are also sensitive to temperature and humidity variations that occur in the ambient environment. NDIR sensors, like the ELT s300 CO₂ sensor employed in this study, have good selectivity, but, since pressure and temperature are not controlled in the optical cavity of ELT s300 CO₂ sensors, the influence of temperature on sensor signals plays an important role. The influence of humidity is also important to address because changes in water vapor are known to influence NDIR measurements of CO₂ in terms of spectral cross-sensitivity due to absorption band broadening (Licor, 2010).

Joanna Gordon 9/30/18 12:47 PM

Deleted: for

Both metal oxide and electrochemical type sensors operate on the principle of oxidizing or reducing reactions at sensor surfaces. For electrochemical sensors, like the Alphasense CO-B4 sensor employed in this study, oxidizing or reducing compounds react at the working electrode, resulting in the transfer of ions across an electrolyte solution from the working electrode to the counter electrode, balanced by the flow of electrons across the circuit connecting the working electrode to the counter electrode. A linear relationship is expected between this current and the target gas mole fraction. Electrochemical sensors can be tuned to respond more or less strongly to specific gases by adjusting the materials properties of the working electrode. A membrane is located between the working electrode and the exterior of the sensor in order to control redox reaction rates. The rates at which gases diffuse through

Joanna Gordon 9/30/18 12:49 PM

Deleted: diffuse

the membrane to reach the working electrode and the electron transfer rates have been shown to increase at higher temperatures (Xiong and Compton, 2014), and since chemical reaction rates are also influenced by temperature, electrochemical sensor responses can be influenced by sensor operating temperature. Changes in ambient humidity levels can cause sensors to lose or gain of the electrolyte solution, by mass, also influencing electrochemical sensor response (Xiong and Compton, 2014).

For metal oxide sensors, and to a lesser extent for electrochemical sensors, resolving the response of a sensor attributable to the target gas species can also pose a challenge in the presence of interfering gas species. Metal oxide sensors, like those used in this study, have a resistive heater circuit that warms up the sensor surface, causing O_2 molecules to adsorb to the sensor surface, which leads to increased resistance across the surface of the sensor. In the presence of an oxidizing compound, like O_3 , more oxygen molecules are adsorbed to the sensor surface and the resistance across the sensor surface is increased further. In the presence of a reducing compound, like CO, oxygen molecules are removed from the sensor surface, allowing electrons to flow more freely, resulting in decreased resistance across the sensor surface. For metal oxide sensors, the resistance across the sensor surface can then be used to determine the mole fraction of a given oxidizing or reducing compound, often according to a nonlinear relationship. Exposure to humidity has been shown to significantly lower the sensitivity of metal oxide gas sensors making it an important parameter to address in a gas quantification model (Wang et al., 2010). Metal oxide sensor operating temperature has also been shown to strongly influence sensor sensitivity and selectivity to different gas species (Wang et al., 2010). Metal oxide type sensors can be tuned to respond differently from one another to oxidizing and reducing gas species by using different metal oxide materials and doping agents for the sensor surface, but selectivity is difficult to achieve.

1.2 Low-Cost Air Quality Sensor Quantification

Because low-cost gas sensor signals are influenced, sometimes significantly, by interfering gas species and changing weather conditions in the ambient environment, field normalization methods to quantify atmospheric trace gases using low-cost sensors have been found to be more effective than lab calibration (Cross et al., 2017; Piedrahita et al., 2014; Sun et al., 2016). Our previous study and several

Joanna Gordon 9/30/18 12:49 PM

Deleted: is

Joanna Gordon 9/29/18 10:13 AM

Formatted: Heading 2

5 others have compared the performance of field calibration models generated using LMs (simple and multiple linear regression) relative to supervised learning methods (including ANNs and random forests), all finding that ANNs (Casey et al., 2017; Spinelle et al., 2015, 2017) and random forests (Zimmerman et al., 2017) outperformed LMs in the ambient field calibration of low-cost sensors. Like
10 earlier laboratory based studies (Brudzewski, 1999; Gulbag and Temurtas, 2006; Huyberegts and Szeco, 1997; Martín et al., 2001; Niebling, 1994; Niebling and Schlachter, 1995; Penza and Cassano, 2003; Reza Nadafi et al., 2010; Srivastava, 2003; Sundgren et al., 1991), ANN-based calibration models, incorporating signals from an array of gas sensors with overlapping sensitivity as inputs, have been able to effectively compensate for the influence of interfering gas species and resolve the target gas mole fraction.

15 ANNs are known to be able to very effectively represent complex, nonlinear, and collinear relationships among input and output variables in a system (Larasati et al., 2011). ANNs are useful in the field calibration of low-cost sensors because, through pattern recognition of a training dataset, they are able to effectively represent the complex processes and relationships among sensors and the ambient environment that would be very challenging to represent analytically or based on empirical representation of individual driving relationships. In practice though, the reason multiple gas sensors are able to improve the performance of calibration models may be in part the result of correlation between mole fractions of target gases themselves that hold for one model training location, but might
20 not remain effective at alternative sampling sites or during other time periods.

1.3 Summary of Previous Study

25 Our previous study was carried out using sensor measurements collected over the course of several months in the spring of 2017, in Greeley, Colorado, which lies within the Denver Julesburg oil and gas production basin. Others had recently demonstrated the utility of machine learning methods in the quantification of atmospheric trace gases using arrays of low-cost sensors in urban (De Vito et al., 2008, 2009; Zimmerman et al., 2017) and rural (Spinelle et al., 2015, 2017) areas. Our previous study tested the relative performance of machine learning methods and LMs in the quantification of CH₄, O₃, CO₂,

Joanna Gordon 9/30/18 12:49 PM
Deleted: performance of

Joanna Gordon 9/2/18 1:55 PM
Deleted: Field normalization methods to quantify atmospheric trace gases using low-cost sensors have been found to be more effective than lab calibration when sensors are deployed in the ambient environment, and subject to changing temperature and humidity (Cross et al., 2017; Piedrahita et al., 2014; Sun et al., 2016). Our previous study and several others have compared the efficacy of LMs (simple and multiple linear regression) relative to supervised learning methods, all finding that ANNs (Casey et al., 2017; Spinelle et al., 2015, 2017) and random forests (Zimmerman et al., 2017) outperformed LMs in the ambient field calibration of low-cost sensors. These effective supervised learning techniques often incorporate multiple gas sensor signals as inputs in order to quantify each target gas, in addition to environmental variable sensor signals, with the goal of compensating for the effects of interfering gas species and environmental factors. In practice though, the reason multiple gas sensors are able to improve the performance of supervised learning type regression, and linear models for that matter, may be in part the result of correlation, or correlation for some time periods, between mole fractions of target gases themselves that hold for one model training location, but might not hold up at alternative sampling sites or during other time periods. -

Joanna Gordon 9/30/18 12:50 PM
Deleted: OurSpecifically, our

Joanna Gordon 9/30/18 12:50 PM
Deleted: utility

Joanna Gordon 9/30/18 12:50 PM
Deleted: performance

and CO in an area strongly influenced by oil and gas production activities, where enhanced levels of hydrocarbons and other industry related pollutants could potentially confound measurements. Our previous study was also the first [to compare](#) machine learning regression techniques with LMs toward the quantification of CH₄ using arrays of low-cost sensors in any setting. [The study tested and compared calibration models using data from two U-Pod sensor systems](#) containing arrays of low-cost gas sensors; [these systems](#) were co-located with optical gas analysers at a Colorado Department of Public Health and Environment monitoring site. ANNs and LMs were trained using a variety of sensor signal input sets from a month of co-located data collected prior to and following a month long test period. The performance of each model was then evaluated relative to reference instrument measurements during the test period. For quantification of all four trace gases that we tested in this oil and gas-influenced setting, we found that ANNs performed better than LMs. The better performance of ANNs over LMs was likely largely attributable to the ability of ANNs to more effectively represent complex and nonlinear relationships among sensor responses, environmental variables, and trace gas mole fractions than LMs. However, the performance of these powerful regression methods could be aided by relationships among atmospheric trace gases specific to the training location, which would not necessarily hold at different sampling sites.

1.4 Spatially Distributed Networks of Sensors and Spatial Extension of Calibration Models

Distributed spatial networks of low-cost sensor [systems have the potential to inform air quality with high spatial and temporal resolution](#). [As such, studies have begun to deploy spatial networks of low-cost sensor systems](#). [These studies rely](#) on the spatial transferability of quantification techniques. [In the present work, we test model performance under conditions of spatial transferability, wherein a model is trained using data from one location then applied to a test dataset using data from a new location](#). [In testing spatial extension of a model, we investigate how well the field calibration of low-cost sensors can inform target gas mole fractions when sensors are deployed in a new location and a new microenvironment of oxidizing and reducing compounds](#). We also test model performance under [conditions of temporal extension, wherein a model is trained using data that was collected only prior or subsequent to the model application period](#). [In testing temporal extension of models, we investigate](#)

- Joanna Gordon 9/30/18 12:50 PM
Deleted: to compare in time
- Joanna Gordon 9/29/18 10:14 AM
Deleted: The A number of
- Joanna Gordon 9/2/18 2:31 PM
Deleted: enclosures
- Joanna Gordon 9/30/18 12:50 PM
Deleted: ; these systems

- Joanna Gordon 9/2/18 2:39 PM
Deleted:

- Joanna Gordon 9/2/18 2:55 PM
Deleted: measurements
- Joanna Gordon 9/29/18 10:25 AM
Deleted:
- Joanna Gordon 9/30/18 12:51 PM
Deleted: As such, studies have begun to deploy spatial networks of low-cost sensor systems.
- Joanna Gordon 9/30/18 12:51 PM
Deleted: These
- Joanna Gordon 9/2/18 2:55 PM
Deleted: that
- Joanna Gordon 9/30/18 12:51 PM
Deleted: rely contingent
- Joanna Gordon 9/30/18 12:51 PM
Deleted: transferability

how model performance is influenced by sensor drift over time. We opportunistically utilize measurements collected with low-cost sensors in Denver, Boulder County, and the DJ and SJ oil and gas production basins in recent years. This effort focuses on the analysis for O₃ and CO₂ using both LMs and ANNs, including a comparison of models with a number of different input sets. In previous work (Casey et al., 2017), we have additionally addressed the quantification of CO and CH₄ using arrays of low-cost sensors together with field normalization methods, but these species are not included in the present analysis because analogous reference data to those we present for O₃ and CO₂ were not available CO and CH₄.

10 1.5 Oil and Gas Production and Air Quality

Oil and gas production related emissions, namely nitrogen oxides (NO_x) and volatile organic compounds (VOCs), have been shown to influence tropospheric ozone (O₃), which is particularly relevant in regions that are in non-attainment of the United States Environmental Protection Agency (USEPA) National Ambient Air Quality Standards (NAAQS) for ozone, like the Colorado Front Range where the DJ Basin is situated. NO_x and VOC emissions, including those from oil and gas production activities, react in the atmosphere in the presence of sunlight to form tropospheric O₃. A number of studies have demonstrated that oil and gas related emissions contribute to increased O₃ in the DJ Basin (Cheadle et al., 2017; Gilman et al., 2013; McDuffie et al., 2016). Mole fractions of ozone as high as 140 ppb and 117 ppb during winter months have also been observed and attributed directly to oil and gas production emissions in the Upper Green River Basin of Wyoming and Utah's Uinta Basin, respectively (Ahmadov et al., 2015; Edwards et al., 2013, 2014; Field et al., 2015; Oltmans et al., 2016; Schnell et al., 2009). Additionally, a modeling study concluded that oil and gas production activities could significantly impact ozone near emissions sources, beginning 2 and 8 km downwind of compressor engine and flaring activities, respectively (Olague, 2012).

Emissions of industry related air pollutants, including O₃ precursors, NO_x and VOCs, are expected to occur on spatially distributed scales, across components on well pads, transmission lines, transportation

Joanna Gordon 9/30/18 12:51 PM

Deleted: influenced

Joanna Gordon 9/4/18 1:09 PM

Deleted: In this work we test how well the field calibration of low-cost sensors can inform target gas mole fractions when sensors are deployed in a new location.

Joanna Gordon 9/30/18 12:51 PM

Deleted: measurements

Joanna Gordon 9/2/18 2:59 PM

Deleted: M

Joanna Gordon 9/2/18 2:58 PM

Deleted: Basins

Joanna Gordon 9/30/18 12:52 PM

Deleted: This effort focuses support on the

Joanna Gordon 9/30/18 12:52 PM

Deleted: using

Joanna Gordon 9/30/18 12:52 PM

Deleted: .

Joanna Gordon 9/29/18 10:29 AM

Deleted: model training and testing deployment pairs, diverged/diverging in location and timing, as such and

Joanna Gordon 9/30/18 12:52 PM

Deleted: reference data

Joanna Gordon 9/4/18 1:11 PM

Formatted: English (UK)

Joanna Gordon 8/27/18 11:16 AM

Deleted:

Joanna Gordon 8/15/18 1:53 PM

Deleted: Very high levels of

Joanna Gordon 8/27/18 12:17 PM

Deleted:

5 routes, and gathering stations that are each distributed throughout production basins (Litovitz et al. 2013; Mitchell et al. 2015; Allen et al. 2013). Spatially distributed networks of low-cost sensors have the potential to better inform spatial variability of air quality than existing regulatory air quality monitoring stations which cannot feasibly cover such spatially resolved measurements continuously, and may not be representative of air quality across smaller spatial scales (Bart et al., 2014; Jiao et al., 2016; Moltchanov et al., 2015). Abeleira and Farmer show that ozone production throughout much of the Front Range, outside of downtown Denver, is likely to be NO_x limited implying that local NO_x sources are likely influencing ozone on small spatial scales (Abeleira and Farmer, 2017). Oil and gas industry related NO_x sources, like well pad combustion and diesel truck traffic, flaring, and compressor engines could lead to pockets of elevated O₃ throughout the DJ Basin. While emissions from truck traffic (and in some cases drilling rig generators), at a given well pad are expected to be highest during the drilling, stimulation, and completion phases, industry truck traffic often persists as the contents of produced water and condensate tanks are frequently collected from well pads throughout production phase, as are emissions from flaring and compressor engines. Low-cost O₃ sensors could augment the few and far apart regulatory sites that currently monitor O₃ levels in places like the DJ Basin, which has better coverage than many other production basins in the United States. While elevated ambient CO₂ levels are not directly harmful to human health, continuous CO₂ measurement can provide information about nearby combustion-related pollution and atmospheric dynamics that lead to the accumulation of potentially harmful compounds associated with the oil and gas production industry during periods of atmospheric stability.

25 In this work, we present and compare models designed to address the unique challenges that come with using low-cost sensors, in the quantification of atmospheric trace gases of interest in oil and gas production basins, where ambient hydrocarbon mole fractions are potentially elevated, exerting uniquely confounding influence on low-cost gas sensors. Calibration models that were found to perform best in our previous study are applied to data collected in different locations. For the first time, we investigate how well models can be transferred from one microenvironment to another, with different dominant local emissions source characteristics, and different relative abundance of oxidizing and

Joanna Gordon 9/30/18 12:52 PM
Deleted: are

Joanna Gordon 8/15/18 2:00 PM
Deleted: pooling

reducing compounds. Microenvironments explored in this work include a basin where both natural gas and heavier hydrocarbons are produced (the DJ Basin), and a basin where prominently natural gas is produced (the SJ Basin), with much smaller proportional emissions of heavier hydrocarbons, and likely lower atmospheric concentrations of alkanes, alkenes, and aromatics. Within and bordering the DJ Basin, additional microenvironments include an urban location, with significant mobile sources emissions (NO_x, CO, and VOCs), and a peri-urban site with fewer mobile emissions and closer proximity to oil and gas production activities. We explore how robust model performance is when a model is trained in one microenvironment and transferred to another; challenged by different relative abundance of oxidizing and reducing gas species. Additionally, we test how well models can represent and address sensor stability over time and the potential for drift.

2 Methods

2.1 Sensors and U-pods

All U-Pod sensor systems (mobilesensingtechnology.com) employed in the case studies, described below, were populated with seven low-cost gas sensors, as in our previous study (Casey et al., 2017). The gas sensors are listed in Table 1 along with their target gas and the model input codes we assigned to each. A RHT03 sensor was used in each U-Pod to measure temperature (temp) and relative humidity (rh). A Bosch BMP085 sensor was used to measure pressure in each U-Pod.

2.2 Case Studies

Five to ten U-Pods were deployed at sampling sites in and around the DJ and SJ Basins from 2014 - 2017. Deployments generally consisted of co-location with reference measurements prior to and following approximately one-month periods of spatially distributed measurements. During some of the distributed measurement periods, a subset of U-Pods remained co-located with reference instruments where the field calibrations took place. As well, during some distributed measurement periods, some U-Pods were deployed in new locations that were equipped with reference measurements. In between periods of distributed sensor system deployments, sensor systems were co-located with reference instruments for as long as possible, as logistics, and coordination with other regulatory agencies and

Joanna Gordon 9/30/18 12:53 PM

Deleted: a

Joanna Gordon 9/30/18 12:53 PM

Deleted: a

Joanna Gordon 9/30/18 12:53 PM

Deleted: tuern

Joanna Gordon 9/2/18 12:32 PM

Deleted: Ten

Joanna Gordon 9/30/18 12:53 PM

Deleted: systemss

Joanna Gordon 8/15/18 2:31 PM

Deleted: a number of

Joanna Gordon 8/29/18 1:49 PM

Deleted: -

Joanna Gordon 8/24/18 4:15 PM

Deleted: Deployment Locations and Timelines

Joanna Gordon 9/2/18 3:08 PM

Deleted: These ten

Joanna Gordon 8/15/18 2:42 PM

Deleted: a number of

Joanna Gordon 9/2/18 3:09 PM

Deleted:

Joanna Gordon 8/16/18 11:29 AM

Deleted: a

Joanna Gordon 8/16/18 11:29 AM

Deleted:

Joanna Gordon 9/2/18 3:10 PM

Deleted:

Joanna Gordon 8/16/18 11:32 AM

Deleted: During other

researchers would allow. In this way, we hoped to maximize our ability to encompass full ranges of temperature, humidity, and trace gases that occur across seasons, in order to minimize extrapolation with respect to these parameters when models were applied to measurements from distributed deployment periods. The locations where all or a subset of U-Pods were co-located with reference instruments are indicated in Fig. 1. In this exploratory study, we opportunistically employ data from these sensor deployments, treating them as case studies in order to characterize the performance of field calibration models when they are extended to new locations. For each case study, described below, data was divided into training and test periods. Timelines for these dataset pairs are detailed in Fig. 2. Some U-Pods employed in these case studies (indicated in grey font in Fig. 2) were constructed, populated with sensors, and deployed at field sites in the spring of 2014, approximately a year before the rest of the U-Pods were constructed, populated with sensors, and deployed at field sites in the spring of 2015. The relative age of sensor systems included in some case study comparisons could have contributed to some discrepancy in model performance, though systematic differences based on U-Pod age is not apparent.

As available data from each case study allowed, we used approximately one month of training data before and after a given test period. When training data was not available within several months of a test period, significantly longer training datasets were used in order to attempt capture and effectively represent trends in sensor drift over time, as well as to avoid extrapolation of model parameters (particularly temperature) during the test data period. As a result, model-training durations varied across case studies and sometimes significantly exceeded model-testing durations. Each case study is similar in representing approximately one month-long deployment of sensor systems. This study design serves a primary goal of this work, supporting the quantification of atmospheric trace gases from low-cost gas sensor data in new locations, relative to model training locations, for periods of approximately one month at a time.

Making quantitative measurements of atmospheric trace gases with low-cost sensors is challenged by unique variations in individual sensor responses associated with variations in the manufacturing

process, sensor age, and sensor exposure history. For these reasons, we generated unique calibration models using data from sensors in each individual U-Pod sensor system. The closest available data prior and or subsequent to a test data period was used for model training to avoid complications associated with significant sensor drift and degradation in sensor sensitivity to target gas species over time. Table 2 lists the O₃ and CO₂ reference instruments that were co-located with U-Pods at each sampling site, along with instrument operators, calibration procedures, and reference data time resolution. The selected case studies, described in sections 2.2.1 through 2.2.7 below, were aimed to support methods to quantify atmospheric trace gases during the distributed deployments we carried out from 2014 through 2017 as well as future distributed sensor network measurements. Fig. 1 shows sampling site locations in context with urban areas and oil and gas production wells. Fig. 2 shows the timeline of each of these deployments, highlighting the training and testing periods defined for both O₃ and CO₂.

2.2.1 Dawson Summer 2014

The first distributed measurement campaign took place during the summer of 2014 when five U-Pods were sited at locations around Boulder County, with four distributed along the eastern boundary of the county, adjacent to Weld County where dense oil and gas production activities were underway. A background site, further from oil and gas production activities was also included to the west, near a busy traffic intersection on the north end of the City of Boulder. Co-locations with reference measurements that were used for field calibration of the sensors took place at the Continuous Ambient Monitoring Program (CAMP) Colorado Department of Health and Environment (CDPHE) air quality monitoring site in downtown Denver. One of the distributed sampling sites, Dawson School, was also equipped with a Thermo Electron 49 O₃ reference instrument, operated by Detlev Helmig's research group from the Institute for Artic and Alpine Research (INSTAAR). In this work, a case study is developed using data from one U-Pod located at the CAMP site in downtown Denver for O₃ model training, and data from one U-Pod, located at the Dawson School for O₃ model testing. This case study is used to test model performance when extrapolated in terms of O₃ mole fractions and applied in a new location, transferred from an urban to a peri-urban environment.

Joanna Gordon 9/30/18 12:56 PM
Deleted: .
Joanna Gordon 8/24/18 4:34 PM
Deleted: .
Joanna Gordon 9/30/18 12:56 PM
Deleted: ,were
Joanna Gordon 9/30/18 12:57 PM
Deleted: to

Joanna Gordon 8/27/18 12:41 PM
Deleted: toward
Joanna Gordon 9/30/18 12:57 PM
Deleted: the
Joanna Gordon 9/30/18 12:57 PM
Deleted:
Joanna Gordon 9/29/18 11:05 AM
Deleted: aan optical O₃ instrument,

Joanna Gordon 8/16/18 1:11 PM
Deleted: This study was funded by Boulder County with the combined aims of gaining a better understanding of how oil and gas emissions affect air quality in Boulder County and learning more about how low-cost air quality measurement methods can help inform air quality in this context. .

2.2.2 SJ Basin Spring 2015

In the spring of 2015, we augmented our original fleet of five U-Pods (BA, BB, BD, BE, and BF) with five more (BC, BG, BH, BI, and BJ) and deployed these sensor systems in the SJ Basin while a targeted field campaign was underway to understand more about a CH₄ ‘hot spot’ that was discovered from satellite based remote sensing measurements (Frankenberg et al., 2016; Kort et al., 2014). The primary goal of this sensor deployment was to inform spatial and temporal patterns in atmospheric trace gases like CH₄, O₃, CO, and CO₂ across the SJ Basin. Most U-Pods were located at existing air quality monitoring sites operated by the New Mexico Air Quality Bureau (NM AQB), the Southern Ute Indian Tribe Air Quality Program (SUIT AQP), and the Navajo Environmental Protection Agency (NEPA), which supported validation of sensor measurements for O₃. After this deployment period, all U-Pods were moved to the BAO site in the DJ Basin for approximately one month, and were co-located with reference instruments there that were operated by National Oceanic and Atmospheric Administration (NOAA) researchers. A case study is developed with data from the BAO site to train O₃ models for four U-Pods, and data from SJ Basin sites to test O₃ models for four U-Pods. This case study is used to test model performance when extrapolated in temperature and time, and extended to a new location, extended from one oil and gas production basin to another across Colorado.

Joanna Gordon 8/29/18 2:51 PM

Deleted: The majority of

Joanna Gordon 8/29/18 2:46 PM

Deleted: and CO₂.

Joanna Gordon 8/29/18 2:51 PM

Deleted: -

2.2.3 SJ Basin Summer 2015

In the summer of 2015, after an approximately month-long co-location with reference instruments at the BAO site, seven U-Pods were deployed again at existing regulatory monitoring sites for approximately one month, after which they were moved back to the BAO site for another month of co-location with reference instruments there. We equipped two of the regulatory monitoring sites in the SJ Basin with LI-COR LI-840A CO₂ analysers to provide reference measurements for CO₂. A case study is developed with data from the BAO site, pre and post of the SJ Basin summer 2015 deployment to train models, and data from SJ Basin sites during the summer deployment period, to test models. Data from seven U-Pods were used to train and test O₃ models and data from two U-Pods were used to train and test CO₂ models. This case study is used to test model performance when training took place both pre

and post of the test period, and when extended to a new location, from one oil and gas production basin to another across Colorado.

2.2.4 BAO Summer 2015

During the SJ Basin Summer 2015 deployment period, two U-Pods remained at the BAO site. A case study is developed using data from those two U-Pods that remained at the BAO site. This case study is used to test model performance when training took place both pre and post of the test period, and when the model was tested on data that was collected in the same location as model training.

2.2.5 BAO Summer 2016

U-Pods were deployed at the BAO site again in 2016 for several months during the summer. In August of 2016 the U-Pods were moved to the Greeley Tower (GRET) CDPHE air quality monitoring site in Greeley, Colorado, a location which, like the BAO site, is also strongly influenced by DJ Basin oil and gas production activities. The U-Pods remained there for one year. For the GRET co-location period, CDPHE shared reference measurements for O₃. Additionally, Jeffrey Collett and Katherine Benedict of Colorado State University (CSU) shared CO₂ reference measurements from an instrument they operated at the site before October 1st in 2016 and after March 7th in 2017, when the instrument was located at the GRET site. A case study is developed using data from two U-Pods during the yearlong deployment at the GRET site to train models for O₃, and data from two U-Pods during the BAO summer 2016 deployment to test models for O₃. Because reference data for CO₂ was not available at the GRET site during winter months, the data from two U-Pods collected during eight months at the GRET site was used to train models for CO₂, and data from two U-Pods during the BAO summer 2016 deployment was used to test models for CO₂. A significantly longer training duration is implemented in this case study because the training period took place more than several months after the model testing period. We reasoned that a longer training duration would be better able to represent patterns in sensor drift over time, as well as encompass the temperature range of test dataset period. Significantly less training time is needed when training occurs directly pre and/or post of a given model application period. This case

Joanna Gordon 9/30/18 12:57 PM

Deleted: those

Joanna Gordon 9/30/18 12:57 PM

Deleted: that remained at

Joanna Gordon 9/29/18 10:39 AM

Deleted: BAO Summer 2016 We deployed a similar distributed network of sensor systems throughout the DJ Basin in the summer of 2015 and the winter/spring of 2016 to explore spatial and temporal patterns in air quality in another region influenced by oil and gas production activities. For these 2015 and 2016 deployments, we co-located U-Pods with reference instruments operated by National Oceanic and Atmospheric Administration (NOAA) researchers at the Boulder Atmospheric Observatory (BAO) Tower toward field calibration, though none of the distributed sampling sites were equipped with reference instruments to support validation. In August of 2016 the ten U-Pods were moved to the Greeley Tower (GRET) CDPHE air quality monitoring site in Greeley, Colorado, a location which is also strongly influenced by DJ Basin oil and gas production activities; the U-Pods remained there for a year. For the GRET co-location period, CDPHE shared reference measurements for O₃. Additionally, Jeffrey Collett and Katherine Benedict of Colorado State University (CSU) shared CO₂ reference measurements from an instrument they operated at the site before October 1st in 2016 and after March 7th in 2017, when the instrument was located at the GRET site.

Joanna Gordon 9/29/18 10:40 AM

Deleted: . The

Joanna Gordon 9/30/18 12:57 PM

Deleted: onea

Joanna Gordon 9/30/18 12:57 PM

Deleted: the

Joanna Gordon 9/30/18 12:57 PM

Deleted: collected

study is used to test model performance when extrapolated significantly (more than several months) in time and extended to a new location, from one location in the DJ basin to another.

2.2.6 GRET Fall 2016

In order to test model performance, under similar circumstances in terms of relative model training and testing durations and timing to the BAO Summer 2016 case study, but with no extension of models to a new location, we developed another case study. This time, models for O₃ and CO₂ were trained using data from two U-Pods at GRET over the course of eight months and models for O₃ and CO₂ were tested using data from two U-Pods at GRET over the course of approximately a month in the fall of 2016. This case study is used to test model performance when extrapolated significantly (more than several months) in time and applied in the same location as training took place.

2.2.7 GRET Spring 2017

We include findings from our previous work as a case study in order to provide context. Models for CO₂ and O₃ were tested using data from two U-Pods collected over the course of approximately one month at the GRET site in the spring of 2017. Data from two U-Pods during approximately month-long periods pre and post of the test period were used to train O₃ and CO₂ models. This case study provides another example of model performance when training took place both pre and post of the test period, and testing took place in the same location as training.

2.3 Reference and Sensor Data Preparation

Each of the U-Pod sensor signals was logged to an onboard micro SD card. For metal oxide type sensors, voltage signals were converted into resistance, and then normalized by the resistance of the sensor in clean air, R_0 . A single value for R_0 was used for each sensor across the study duration. This R_0 value was taken as the resistance of each sensor during the GRET Spring 2017 field deployment period, when the target pollutant had approached background levels (at night for the metal oxide O₃ sensors and midday for all other metal oxide sensors), and when the ambient temperature was approximately 20° C and relative humidity of approximately 25%. Relative humidity, temperature, and pressure measured in each U-Pod were used to calculate absolute humidity. Over the course of multiple field deployments,

Joanna Gordon 9/30/18 12:57 PM

Deleted: the

Joanna Gordon 8/24/18 4:39 PM

Deleted: .

Joanna Gordon 8/29/18 3:43 PM

Deleted: The work presented here is aimed at supporting methods to quantify atmospheric trace gases during the distributed deployments described above as well as future distributed sensor network measurements. Fig. 2 shows sampling site locations in context with urban areas and oil and gas production wells. Fig. 1 shows the timeline of each of these deployments, highlighting the training and testing periods defined for both O₃ and CO₂.

Joanna Gordon 8/27/18 12:39 PM

Deleted:

Joanna Gordon 8/16/18 1:49 PM

Deleted: several

relative humidity sensors in four of the U-Pods drifted down, causing the lower humidity levels to be cut off or 'bottomed out'. RH sensors were not replaced during field deployments in order to preserve consistency across different deployment periods, allowing for the possibility of a single comprehensive model to apply to all data from a single U-Pod. After some experimentation in generating a 'master model' that could be applied to data from a given U-Pod for all collected field measurements, across several years, we determined that individual models for each deployment would be more effective, and replacing RH sensors that had drifted down would have been appropriate in support of the methods presented here. We have since upgraded to Sensirion AG SHT25 sensors, which appear to be more robust and consistent over the course of long-term field deployments. For measurements collected in the spring and summer of 2015 and the spring of 2017, we replaced the relative humidity (RH) signal of U-Pods with malfunctioning humidity sensors with signals from the closest U-Pod, with a good humidity sensor and complete data coverage as noted in Table S1. Temperature and RH sensor measurements are usually collected from within each U-Pod sensor system in order to gain representative information about the environment the gas sensors are being operated in. Using an alternative source for RH data that are not onboard an individual U-Pod has the potential to increase uncertainty of quantified gas mole fractions. We used replacement RH data from the closest available U-Pod instead of ambient measurements in order to more closely approximate humidity at the operating temperature within a U-Pod enclosure. The closest U-Pod with good humidity sensors ranged from several feet, when U-Pods were co-located during deployments in the DJ Basin at the BAO and GRET sites, to approximately fifty miles during deployments in the San Juan Basin.

When the U-Pods were initially deployed at the GRET site, on August 23rd of 2016, the RH sensors in all ten U-Pods malfunctioned, logging an error code of -99 instead of the relative humidity. This malfunction seemed to coincide with the implementation of radio communication from each U-Pod to a central node in an effort to reduce trips to the field site to download data and to identify issues with data acquisition promptly. No other impacts to sensor systems were observed in connection with radio communications. RH signals in the U-Pods began logging correctly again in October when we stopped remote communication. We replaced RH values for the U-Pods during this time period by utilizing data

Joanna Gordon 8/16/18 1:50 PM
Deleted: a few

Joanna Gordon 8/16/18 1:53 PM
Deleted: nearby

Joanna Gordon 9/30/18 12:57 PM
Deleted: Pods

Joanna Gordon 9/30/18 12:58 PM
Deleted: sensorsensors

Joanna Gordon 8/16/18 1:52 PM
Deleted: -

Joanna Gordon 8/20/18 4:12 PM
Deleted: November

from the Picarro Cavity Ring-Down Spectrometer that was co-located at GRET with the U-Pods. Water mole fractions measured by the Picarro were converted into mass-based mixing ratios to match the units of the absolute humidity signal in the U-Pod data. We applied an adjustment to this absolute humidity signal so that it matched observations in U-Pods during the following month when good RH sensor data was available, to account for the fact that temperatures were higher in U-Pod enclosures than the ambient environment. We then replaced the relative humidity signal in each U-Pod from August 23rd through October 1st in 2016 with the mixing ratios derived from Picarro measurements. Using the temperature and pressure logged in each U-Pod along with the absolute humidity from the Picarro, relative humidity was calculated for each U-Pod during this period.

Joanna Gordon 9/2/18 6:40 PM

Deleted: absolute

To perform regressions toward field calibration of sensors, the reference and U-Pod data needed to be aligned. When reference measurements with minute time resolution were available for both training and corresponding testing periods, minute median data from the U-Pods were used. Medians were used as opposed to averages in order to reduce the potential influence of sensor noise as well as to remove short duration spikes in the reference and sensor data that resulted from air masses that may not have been well mixed across the reference instrument inlets and the U-Pod enclosures. When reference data were instead available with only 5-minute or 60-minute time resolution, U-Pod medians were calculated for to match that time step. In order to test models using the same time resolution they were trained with, the time resolution of reference and sensor measurements for corresponding training/testing datasets were matched, if necessary, by taking medians of the dataset with higher time resolution to match the data with the longer time resolution. The first 15 minutes of data after any period that the U-Pods had not recorded data for the previous 5 minutes was removed in order to filter transient behavior associated with sensor warm-up. During a given deployment, the data removed to avoid sensor warm-up transients constituted less than 1%.

Joanna Gordon 9/29/18 9:49 AM

Deleted: to match thatthe same time step

Joanna Gordon 9/29/18 10:45 AM

Deleted: Medians were also calculated for reference measurements with finer time resolution to match the time resolution of corresponding training/testing data

When time was included in a model as an input, the absolute time was used. Specifically, we used the datenum value from the MATLAB environment, which is defined by the number of days that have elapsed since the start of January 1st, in the year 0000. A model was extrapolated in time when ever

training data does not take place both before and after a given test deployment period. In several case studies we present, model training only took place after the test deployment period, comprising a ‘post only’ calibration. In Colorado, and more broadly in the western United States, ambient temperatures change significantly across the seasons throughout the year, so if a model is extrapolated in time, extrapolation in temperature often results as well.

2.4 Calibration Model Techniques

In this work, we explore how well field calibration methods hold up in new locations, a topic which has not yet been sufficiently addressed by the scientific community. As in (Casey et al., 2017), direct LMs and ANNs were trained with a number of different sensor input sets to map those inputs to target gas mole fractions measured by reference instruments. Direct LMs implemented were multiple linear regression models given by,

$$r = p_1 + p_2s_1 + p_3s_2 + \dots + p_ns_{n-1}, \quad (1)$$

where r is the target gas mole fraction (measured by a reference instrument) $s_1 - s_{n-1}$ are sensor signals from U-Pods that are included as model predictor variables, and $p_1 - p_n$ are corresponding predictor coefficients.

ANNs designed for regression tasks, like those employed in this work, generally consist of artificial neuron nodes that are connected with weights. Weights are initiated with randomly assigned values. An optimization algorithm is then employed to map a given set input values to corresponding target values. An example of a very simple feed forward neural network, and how weights are propagated through it are depicted in Fig. 3. In this work, ANNs were designed by assigning U-Pod sensor signals to artificial neurons in an input layer and assigning target gas mole fractions for an individual gas species, measured by a reference instrument to a single output neuron. Nonlinear, tansig, artificial neurons in one hidden layer for O_3 or two hidden layers for CO_2 (accordance with our earlier findings for each target gas species (Casey et al., 2017)) were then added between input layer and the network output neuron. Additionally, bias neurons, each assigned a value of 1, were connected to neurons in the hidden layer(s) so that individual connecting weights could be activated or deactivated during the

Joanna Gordon 9/30/18 12:58 PM

Deleted: .

Joanna Gordon 9/30/18 12:59 PM

Formatted: Indent: First line: 0.4"

Joanna Gordon 9/30/18 12:58 PM

Deleted: In this work, we explore how well field calibration methods hold up in new locations, a topic which has not yet been sufficiently addressed by the scientific community. As in (Casey et al., 2017), direct LMs and ANNs were trained with a number of different sensor input sets to map those inputs to target gas mole fractions measured by reference instruments. Direct LMs implemented were multiple linear regression models given by . Field calibration of low-cost gas sensors to quantify ambient trace gas mole fractions has been shown to work more effectively than lab calibration methods in several previous studies (Piedrahita et al., 2014; Sun et al., 2016; Zimmerman et al., 2017). Low-cost sensors are often sensitive to many aspects of the complex ambient environment in terms of temperature, humidity, and atmospheric chemistry variations. While trace gases that sensors are specifically designed to respond to elicit a sensor response, so can many environmental and chemical confounding influences. Field calibrations have been shown to more effectively mitigate the effects of these confounding influences in the quantification of atmospheric trace gases in the ambient environment relative to lab calibrations. In this work, we explore how well field calibration methods hold up in new locations, a topic which has not yet been sufficiently addressed by the scientific community. .

... [1]

optimization process. The number of neurons in each hidden layer was set equal to the number of inputs included in a given ANN. Fig. 4 shows a diagram of an ANN architecture employed in this work, when there were five inputs.

5 For ANN training we employed the Levenberg Marquardt optimization algorithm with Bayesian Regularization (Hagan et al., 1997). The Levenberg-Marquardt algorithm combines the Gauss-Newton and Gradient Decent methods, towards incremental minimization of a cost function, which is the summed squared error between the ANN output and target values as a function of all of the weights in the network. Training begins according to the Gauss-Newton method, in which the Hessian matrix, the second order Taylor series representation of the error surface, is approximated as a function of the Jacobian matrix and its transpose, significantly reducing required training time. Network weights are adjusted accordingly during each training step to reduce error. If the cost function is not reduced in a given training step, an algorithm parameter is adjusted so that optimization more closely approximates the gradient decent method (a first order Taylor series representation of the cost function), providing a guarantee of convergence on a cost function minimum. Since local minima may exist across the error surface, it is important to train the same network multiple times, with different randomly assigned starting weights, in order to access the stability of ANN performance. In this work, each ANN was trained 5 times.

20 In the implementation of Bayesian Regularization, a term is added to the sum of squared error cost function as a penalty for increased network complexity in order to guard against over fitting. A two level Bayesian inference framework is employed, operating on the assumptions that the noise in the training data is independent, normally distributed, and also that all of the weights in the ANN are small, normally distributed, and unbiased (Hagan et al., 1997). In preliminary ANN tests we found that over fitting occurred even when Bayesian Regularization was used, so we additionally implemented early stopping, which proved to be effective in the reduction of over fitting. To implement early stopping, a portion of training data is set aside as validation dataset, and during training. Training continues so long as the error associated with the validation dataset is reduced. When the error associated with the

Joanna Gordon 9/30/18 1:00 PM

Deleted: For ANN training we employed the Levenberg Marquardt optimization algorithm with Bayesian Regularization. Subsequently, the resulting quantification model performance was evaluated on test datasets. Generally, model inputs were normalized so that they ranged in magnitude from -1 to 1 since this practice is recommended for the ANN optimization algorithm used (Hagan et al., 1997). The LMs employed here consisted of multiple variable linear regressions, where the determined coefficients for a linear combination of sensor inputs minimized residuals between model predictions and reference instrument target gas mole fractions. We used the same set of inputs in ANNs, employing the Levenberg Marquardt optimization algorithm with Bayesian Regularization

Joanna Gordon 9/30/18 1:00 PM

Deleted: which is

Joanna Gordon 9/30/18 1:00 PM

Deleted: during

Joanna Gordon 9/30/18 1:00 PM

Deleted: that

Joanna Gordon 8/24/18 3:06 PM

Deleted: Bayesian Regularization is designed to guard against over fitting.

Joanna Gordon 8/25/18 4:19 PM

Deleted: showed successfully reduced

validation dataset is no longer being reduced, training stops early. For ANNs, training datasets were divided in half on an alternating 24-hr basis, with half used for training and half used as validation data for early stopping. Input signals for both LMs and ANNs were normalized so that they ranged in magnitude from -1 to 1 since this practice is recommended for the ANN optimization algorithm used (Hagan et al., 1997).

2.5 Calibration Model Evaluation and Testing

To evaluate the performance of each of the ANN and LM models that were generated using training data then applied to test datasets, we explored residuals, the coefficient of determination (r^2), root mean squared error (RMSE), mean bias error (MBE), and centered root mean squared error (CRMSE). The CRMSE is an indicator of the distribution of errors about the mean, or the random component of the error. The MBE, alternatively, is an indicator of the systematic component of the error. The sum of the squares of the CRMSE and the MBE is equal to the square of the total error, the square root of which is defined by the RMSE.

First, we generated and applied the best performing model, as determined in our previous work (presented in Table 3), to data from each new case study. Each new case study was selected to challenge models in different ways in order to evaluate the resiliency of the findings from our previous study when challenged by different circumstances. Then we tested LMs for CO₂ and O₃ that contained only the primary target gas sensor for each species, as well as temperature and absolute humidity as inputs. Finally, we generated, applied, and evaluated the performance of a number of LMs and ANNs with different sets of inputs for each case study in order to see which specific model performed the best for each individual case study. The r^2 , RMSE, and MBE for each of these alternative models when applied to test data are presented in the supplemental materials (SM) in Fig. S2 through Fig. S7, along with representative scatter plots and time series comparing the performance LMs and ANNs for a given set of inputs. In Fig. S2 through Fig. S7, the best performing model inputs for each train/test data pair are shaded in purple. The type of model that performed the best (ANN vs. LM) is indicated in the caption of each figure. We discuss both the performance of the previously determined best fitting

Joanna Gordon 8/24/18 3:38 PM

Deleted: Each ANN had one hidden neuron in each hidden layer. ANNs with two hidden layers were used for CO₂ and ANNs with one hidden layer were used for O₃, in accordance with our earlier findings for each target gas species (Casey et al., 2017).

Joanna Gordon 9/30/18 1:00 PM

Deleted: explored

Joanna Gordon 9/30/18 1:00 PM

Deleted: Then

Joanna Gordon 8/27/18 2:28 PM

Deleted: As in our previous work (Casey et al., 2017),

Joanna Gordon 8/27/18 2:33 PM

Deleted: compared performance

Joanna Gordon 8/27/18 2:33 PM

Deleted: a number of

Joanna Gordon 9/30/18 1:00 PM

Deleted: casetrain/test data pair

model (generated using data from the GRET Spring 2017 case study) when applied and generated to data from new case studies, and the performance of models that were tuned to perform the best for each individual case study. From these comparisons, we draw insight into circumstances that challenge model performance in terms of relative local emissions characteristics, location, and timing between model training and testing pairs. Table 4 lists the relative timing and parameter coverage between model training and testing periods for dataset pairs, highlighting instances of incomplete coverage during training that led to model extrapolation during testing.

3 Results and Discussion

3.1 BAO and SJ Basin Summer 2015

The set of deployments we conducted in the summer of 2015 is particularly useful to the objective of characterizing how well field calibration models can be extended to a new location relative to their performance where they were trained. During the testing period, two U-Pods were located at BAO, where training took place, while seven U-Pods were co-located with reference measurements for O₃ and two U-Pods were co-located with reference measurements for CO₂ in the SJ Basin, across Colorado and over the state line in New Mexico. Sampling sites at BAO, in the DJ Basin, and throughout the SJ Basin were all influenced by oil and gas production activities and their associated emissions to some extent, but the composition of the production stream is different in each basin. In the SJ Basin, particularly the northern portion of the basin where all our sampling sites were located production is dominated by coalbed methane. In contrast, many wells in the DJ Basin produce both oil and gas leading to greater relative abundance of heavier hydrocarbons in emissions. The DJ Basin air shed is also more strongly impacted by urban emissions than the SJ Basin air shed, and is more strongly influenced by mobile sources with Denver, Boulder, Fort Collins, Greeley, and many other smaller communities in its midst and along its borders. The Four Corners region, where the SJ Basin is situated, has a much smaller population density. Additionally, while there are some agricultural activities and associated emissions in and around the SJ Basin, there is a significantly larger agricultural industry in and around the DJ Basin. SJ Basin sampling sites spanned a range of elevations, including some that

Joanna Gordon 8/24/18 11:33 AM

Deleted:

Joanna Gordon 9/30/18 1:01 PM

Deleted: BAO and SJ Basin Summer 2015 ... [2]

Joanna Gordon 8/27/18 3:11 PM

Formatted: Font:Not Bold, English (UK)

Joanna Gordon 9/3/18 11:15 AM

Deleted: eight

Joanna Gordon 9/29/18 11:11 AM

Deleted: comoved to sampling sites

Joanna Gordon 9/2/18 7:04 PM

Deleted: Seven of these SJ Basin sites were equipped with O₃ reference measurements and two of the sites were equipped with CO₂ reference measurements. This testing period supports a direct comparison of model performance at the training location relative to a number of distributed sampling sites. With this dataset, we investigate how sensor field calibrations can hold up, both across a large geographic distance, as well as in an environment with different atmospheric chemistry of reducing and oxidizing gases that can influence sensor signals.

Joanna Gordon 9/29/18 11:12 AM

Deleted: compositionpredominant chemistry in the

Joanna Gordon 8/24/18 6:07 PM

Deleted: production stream

were higher and some that were lower than the BAO Tower, coinciding with a wide range of atmospheric pressure at the distributed sampling sites.

5 | [We began by testing the best-performing CO₂ model, as determined in our previous work](#) (Casey et al., 2017), [on this data, collected under a different set of circumstances, during the summer of 2015. ANNs that were trained for each U-Pod using data from the BAO Tower with the following inputs from each U-Pod: eltCO2 \(ELT S300 CO2 sensor\), temp \(temperature\), and absHum \(absolute humidity\).](#) These data [are](#) presented in Fig. [5](#) and Fig. [6](#). Fig. [5](#) shows scatter plots of U-Pod CO₂ vs. reference CO₂ during the test data period for sensors located at BAO as well as sensors that were located at distributed 10 | sampling sites throughout the SJ Basin. The scatter plots show that while there was generally a smaller dynamic range of CO₂ at the SJ Basin sites relative to BAO, model performance did not appear to be impacted or degraded by spatial extension to these locations in the SJ Basin. The line of best fit for Fort Lewis site (periwinkle) is even closer to the 1:1 than the lines of best fit for two U-Pods located at BAO (black and grey). Overlaid histograms of residuals in the bottom right corner of Fig. [5](#) show that CO₂ 15 | residuals from each of the SJ Basin U-Pods are generally centered and evenly distributed about zero with similar spread.

U-Pod CO₂ average residuals from the same data presented in Fig. [5](#), quantified using ANNs with eltCO2, temp, and absHum signals as inputs, are plotted according to time of day and date in Fig. [6](#). 20 | While the use of ANNs in place of LMs was shown to reduce U-Pod CO₂ residuals significantly with respect to temperature, some daily periodicity in the residuals for all four U-Pods is apparent in the upper plot in Fig. [6](#) that shows residuals by date. The lower plot in Fig. [6](#), showing residuals by time of day, demonstrates that CO₂ from [three of four U-Pods](#) was generally under predicted during early hours of the morning and generally over predicted during afternoon and evening hours. Interestingly, this 25 | trend in residuals by time of day is more pronounced for the two U-Pods that remained at BAO. [Upon examination of overlaid histograms showing distributions of parameters during model testing and training periods, in Fig. S12, and model time series and residuals plots in Fig. S3, there is no indication of model extrapolation at the BAO site, nor the sites in the SJ Basin \(with the exception of pressure due](#)

Joanna Gordon 9/29/18 11:13 AM
Deleted: on
Joanna Gordon 8/24/18 5:16 PM
Deleted: .
Joanna Gordon 9/30/18 1:01 PM
Deleted: These U-Pod CO ₂
Joanna Gordon 9/29/18 11:13 AM
Deleted: 5
Joanna Gordon 8/24/18 2:05 PM
Deleted: 3
Joanna Gordon 9/29/18 11:13 AM
Deleted: 6.
Joanna Gordon 8/24/18 2:06 PM
Deleted: 4
Joanna Gordon 8/24/18 5:21 PM
Deleted: were quantified using ANNs that were trained using data from the BAO Tower with the following inputs from each U-Pod: eltCO2, t... [3]
Joanna Gordon 9/29/18 11:13 AM
Deleted: 5
Joanna Gordon 8/24/18 2:06 PM
Deleted: 3
Joanna Gordon 9/30/18 1:01 PM
Deleted: 5
Joanna Gordon 8/24/18 2:06 PM
Deleted: 3
Joanna Gordon 9/29/18 11:13 AM
Deleted: 5
Joanna Gordon 8/24/18 2:06 PM
Deleted: 3
Joanna Gordon 9/29/18 11:13 AM
Deleted: 6
Joanna Gordon 8/24/18 2:06 PM
Deleted: 4
Joanna Gordon 9/29/18 11:13 AM
Deleted: 6
Joanna Gordon 8/24/18 2:06 PM
Deleted: 4
Joanna Gordon 9/29/18 11:14 AM
Deleted: 6
Joanna Gordon 8/24/18 2:06 PM
Deleted: 4
Joanna Gordon 8/27/18 1:59 PM
Deleted: all
Joanna Gordon 8/24/18 5:44 PM
Deleted:

the wide range altitudes encompassed) and no significant trends of concern with respect to residuals with respect to model inputs. Bias introduced to mole fraction estimates are likely attributable to differences in hydrocarbon mixtures in the SJ Basin relative to the DJ Basin.

5 | Next we evaluated the best model type and set of inputs for this specific case study. Differing from our previous findings, for this group of training and testing data pairs from the summer of 2015 at the BAO and SJ Basin sites, the inclusion of the e2vVOC (e2v MiCs-5521) and alphaCO (Alphasense CO-B4) sensor signals noticeably improved the RMSE in the quantification of CO₂. While the inclusion of these two secondary sensor signals didn't result in the best performance in our previous study, using data from the GRET site (Casey et al., 2017), their inclusion did not degrade performance relative to the models that included just eltCO₂, temp, and absHum signals as inputs, so including these sensor signals may be appropriate as a general rule, in areas that are strongly influenced by oil and gas production activities. Generally, using rh vs. absHum signals as ANN inputs did not have a measurable impact on model performance, though linear models were sometimes found to perform better when the absHum signal is used instead of the rh signal. From Fig. S2, it is apparent that inputs including e2vCO (e2v MiCs-5525), temp, rh, e2vVOC, and alphaCO sensor signals as model inputs resulted in the lowest RMSE for U-Pods at BAO as well as at the two SJ Basin sites. Plots analogous to those presented in Fig. 5 and Fig. 6, but with this best performing set of inputs for the present data set pairs are presented in the SM, in Fig. S24 and Fig. S25 respectively.

20 | O₃ was quantified for the two U-Pods deployed at BAO and seven of the U-Pods deployed at SJ Basin sampling sites using an ANN with the following inputs: e2vO₃ (e2v MiCs-2611), temp, absHum, e2vCO, e2vVOC, figCH₄ (Figaro TGS 2600), and figCxHy (Figaro TGS 2602). ANNs with this configuration were found to perform best in the quantification of O₃ in our previous study (Casey et al., 2017). These same inputs and model configuration were also found to be the best performing for the U-Pods at the BAO site and the majority of SJ Basin 2015 dataset pairs as noted in Fig. S2. Interestingly though, LMs with this same set of inputs performed competitively well for three of the seven U-Pods in the SJ Basin in terms of RMSE and r². The observation that LMs performed competitively well at a

- Joanna Gordon 9/30/18 5:02 PM
Deleted: The majority of U-Pods stopped logging data, unfortunately, at one point or another during these deployments. The periods of missing d... [5]
- Joanna Gordon 8/25/18 11:21 AM
Deleted: we found that
- Joanna Gordon 9/29/18 11:14 AM
Deleted: , .
- Joanna Gordon 9/2/18 9:05 PM
Deleted:
- Joanna Gordon 9/29/18 11:14 AM
Deleted: haveseem to make a big difference in
- Joanna Gordon 9/30/18 1:01 PM
Deleted: impact
- Joanna Gordon 9/30/18 1:01 PM
Deleted: to
- Joanna Gordon 9/29/18 11:15 AM
Deleted: 5
- Joanna Gordon 8/24/18 2:07 PM
Deleted: 3
- Joanna Gordon 9/29/18 11:15 AM
Deleted: 6
- Joanna Gordon 8/24/18 2:07 PM
Deleted: 4
- Joanna Gordon 8/25/18 12:25 PM
Deleted: all
- Joanna Gordon 9/30/18 1:01 PM
Deleted: two2the
- Joanna Gordon 9/30/18 1:01 PM
Deleted: seven7
- Joanna Gordon 9/1/18 9:09 PM
Deleted: among
- Joanna Gordon 9/1/18 9:09 PM
Deleted: others tested for
- Joanna Gordon 9/30/18 1:02 PM
Deleted: three
- Joanna Gordon 8/25/18 12:29 PM
Deleted: for a number
- Joanna Gordon 9/30/18 1:02 PM
Deleted: seven
- Joanna Gordon 8/25/18 12:30 PM
Deleted: . For three of seven U-Pods in
- Joanna Gordon 9/30/18 1:02 PM
Deleted: terms of RMSE and r².the SJ Basi... [5]
- Joanna Gordon 9/29/18 11:16 AM
Deleted: . When the BAO trained U-Pods ... [6]

subset of SJ Basin sites is likely connected to the relative abundance of hydrocarbons and other potentially interfering oxidizing and reducing gas species at individual sampling sites, diverging from conditions present during model training at the BAO site. ANNs can better represent the influence of these interfering species than LMs during training, but appear to have lost their ability to do so for this subset of microenvironments in the SJ Basin.

Scatter plots and trends in residuals are presented in Fig. 7 and Fig. 8 for O₃. These plots show the performance of U-Pods at BAO relative to those at SJ Basin sites in the quantification of O₃ during the test data period. U-Pod O₃ measurements at Fort Lewis, Navajo Dam, and the Sub Station did not agree with reference measurements as well as U-Pod O₃ measurements from the other four SJ Basin sites. As noted earlier, U-Pods at the Navajo Dam and Sub Station sites had faulty relative humidity sensor data, so humidity from the U-Pod located at the Ignacio site was used in place of their humidity signals. Since the Ignacio site was located approximately twenty-two and fifty miles away from the Navajo Dam and Sub Station sites respectively, this could have introduced some additional error into the application of a calibration equation, particularly since we showed earlier that O₃ ANNs like the ones we employed here are very sensitive to humidity inputs (Casey et al., 2017). Spatial variability in humidity across tens of miles could be significant as isolated storms (which are on average 15 miles in diameter) propagate throughout the region in the summer. At the Fort Lewis site, a 2b Technologies model 202 O₃ analyser was employed as a reference instrument, differing from the Thermo Scientific 49i, Thermo Scientific 49is, and Teledyne API T400 instruments utilized for reference measurements, elsewhere in the SJ Basin, and the Thermo Scientific 49c that was operated at the BAO site and used for model training. Of all the reference instruments, only the 2b Technologies model 202 O₃ at the Fort Lewis site was operated in a room that was not temperature controlled, as such, some bias may have been introduced to the Fort Lewis O₃ reference measurements. Different instruments, operators, calibration and data quality checking procedures could have contributed to observed discrepancies. It is also possible that the microenvironment at each of these three sites contributed to lower model performance.

Fig. S1 shows that differences among U-Pod O₃ performance during the test deployment period were larger than those observed during the training phase among the same U-Pods; therefore, the incongruous

- Joanna Gordon 9/29/18 11:16 AM
Deleted: 7
- Joanna Gordon 8/24/18 2:07 PM
Deleted: 5
- Joanna Gordon 9/29/18 11:16 AM
Deleted: 8
- Joanna Gordon 8/24/18 2:07 PM
Deleted: 6
- Joanna Gordon 9/29/18 11:16 AM
Deleted: DamLake
- Joanna Gordon 9/30/18 1:02 PM
Deleted: PodsP
- Joanna Gordon 8/25/18 12:49 PM
Deleted: Lake
- Joanna Gordon 8/25/18 12:51 PM
Deleted: bad humidity
- Joanna Gordon 9/29/18 11:17 AM
Deleted: data, as noted in section 6.2.3 and Table S1, so
- Joanna Gordon 9/29/18 11:16 AM
Deleted: approximatelyrelatively far
- Joanna Gordon 9/2/18 6:42 PM
Deleted:
- Joanna Gordon 9/29/18 11:17 AM
Deleted: instrument
- Joanna Gordon 8/25/18 1:19 PM
Deleted: The Fort Lewis site
- Joanna Gordon 9/30/18 1:02 PM
Deleted: , as such, sSome
- Joanna Gordon 9/30/18 1:02 PM
Deleted: .
- Joanna Gordon 9/29/18 11:17 AM
Deleted: could had a different reference instrument than those used at other sites, which may
- Joanna Gordon 9/30/18 1:02 PM
Deleted: to

field calibration performance phenomena we observed seems to be connected to [unique characteristics associated with individual](#) sampling sites; [possibly relative abundance of oxidizing and reducing molecules in the local atmosphere, which could interfere with sensor responses to their target gas species](#), as opposed to the quality of individual sensors in each of those U-Pods.

All SJ Basin U-Pod O₃ measurements systematically over estimate lower levels of O₃ each night, a trend apparent in the scatter plots in Fig. [7](#) and in the residuals by time of day plot in Fig. [8](#). Upon examination of the scatter plots in Fig. [7](#), U-Pods at some sampling sites had positive bias for higher O₃ measurements as well (Shiprock, Ignacio, Sub Station, and Bloomfield), while for others, bias at the higher end of O₃ distributions did not appear to be present (Navajo Dam, Fort Lewis, and Bondad). The residuals by time of day plot in Fig. [8](#) shows that the two U-Pods at BAO did not have significant trends in their residuals according to the time of day, but that U-Pods deployed at SJ Basin sites consistently over estimated nighttime O₃. The residuals are also plotted with respect to temperature in Fig. [8](#), where all U-Pods, even those at BAO to a lesser extent, appear to over predict O₃ at lower temperatures, which generally occurred at night. [In general, the](#) times of day that correspond to the highest O₃ levels had the lowest residuals, with some exceptions at the Fort Lewis and Navajo Dam sites.

Fig. [8](#) includes a plot of the residuals across the duration of the deployment period, showing no significant sensor drift in measurements for any of the U-Pods. This plot also shows that the highest residuals observed generally occurred over short periods in time, particularly for the Fort Lewis (blue) and Sub Station (magenta) sites. In order to further explore the performance of field calibration models for O₃ at SJ Basin sites relative to BAO, the combined parameter space of temperature with O₃ reference mole fractions and temperature with absolute humidity are presented in Fig. [9](#). The combined temperature and reference O₃ parameter space appears to be similar for all of the U-Pods, both at BAO and the SJ Basin sites. However, there appears to be some outlying combined temperature and humidity parameter space at the Sub Station site and at the Navajo Dam site. Brief excursions, [lasting approximately 2 – 4 hours](#), of high humidity ([up to 0.025 kg/kg, relative to the upper bound of absolute](#)

Joanna Gordon 8/25/18 1:33 PM
Deleted: the
Joanna Gordon 8/25/18 1:34 PM
Deleted: or the circumstances discussed previously

Joanna Gordon 9/29/18 11:18 AM
Deleted: 7
Joanna Gordon 8/24/18 2:07 PM
Deleted: 5
Joanna Gordon 9/29/18 11:18 AM
Deleted: 8
Joanna Gordon 8/24/18 2:08 PM
Deleted: 6
Joanna Gordon 9/29/18 11:18 AM
Deleted: 7
Joanna Gordon 8/24/18 2:08 PM
Deleted: 5
Joanna Gordon 9/29/18 11:18 AM
Deleted: 8
Joanna Gordon 8/24/18 2:08 PM
Deleted: 6
Joanna Gordon 9/29/18 11:18 AM
Deleted: 8
Joanna Gordon 8/24/18 2:08 PM
Deleted: 6
Joanna Gordon 9/30/18 1:02 PM
Deleted: In general, tTheThe
Joanna Gordon 9/29/18 11:18 AM
Deleted: 8
Joanna Gordon 8/24/18 2:08 PM
Deleted: 6
Joanna Gordon 9/29/18 11:18 AM
Deleted: 9
Joanna Gordon 8/24/18 2:08 PM
Deleted: 7

[humidity observed at other sampling sites of 0.013 kg/kg](#)) may be connected to some of the large short-term residuals observed at these two sampling sites.

The majority of U-Pods stopped logging data, unfortunately, at one point or another during these deployments. Periods of missed data during the month-long deployment included approximately one day at the Shiprock site, two days at the Bloomfield site, four days at the Sub Station site, nine days at the Fort Lewis site, and seventeen days at the Navajo Dam site. We carried out frequent sampling site visits (on a weekly or biweekly basis as logistics and travel to remote locations in some cases allowed) in order to identify and fix problems as they arose during field deployments. Operational issues were predominantly attributable to power supply problems associated with BNC bulkhead fittings and corrupted micro SD cards. The periods of missing data are reflected in the plots of residuals by date in Fig. 6 for CO₂ and in Fig. 8 for O₃. Fortunately, no drift over the course of the deployment period was observed in these plots.

3.2 Insight from Additional Case Studies of Field Calibration Extension to New Locations

3.2.1 Urban calibration moved to rural/peri-urban setting: Dawson Summer 2014

The Boulder County deployment in the summer of 2014 was used to test how well a field calibration for sensors in one U-Pod, generated in a busy urban area (at CAMP in downtown Denver), could be extended to a peri-urban setting (at Dawson School in eastern Boulder County). Training took place at CAMP for several days each month, before and after each approximately month-long deployment period at Dawson School over the course of four months. Fig. S7 shows the performance of a number of ANN and LM-based CAMP field calibrations with different sets of inputs at this Dawson School test site. In this case study, LMs performed better than ANNs across all sets of sensor inputs. Unlike findings from our previous study (Casey et al., 2017), including secondary metal oxide type sensors as inputs didn't help to improve model performance. The best performing set of inputs included just e2vO3, temp, and absHum signals. The very different relative abundance of various oxidizing and reducing compounds in downtown Denver relative to the Dawson School site, surrounded by open grassy fields, and in closer proximity to oil and gas production activities, may be the reason why

Joanna Gordon 9/30/18 1:03 PM

Deleted: ,

Joanna Gordon 9/30/18 1:03 PM

Deleted:),

including additional gas sensors as model inputs and the use of ANNs failed to improve the quantification of U-Pod O₃ in this case. Relatively short training durations could also contribute to this finding, based on findings from our previous work (Casey et al., 2017).

5 [The fact that LMs performed better than ANNs in this case \(with an r2 of .95 and RMSE of 0.35 ppb for LMs, as opposed to an r2 of .9 and an RMSE of 5.1 ppb for ANNs\) may have to do with the general expectation that LMs be more resilient to extrapolation than ANNs.](#) Notably though, neither ANNs nor the LMs captured the highest levels of O₃ at Dawson School well. We attribute the poor performance at high levels of O₃ at this site, those in exceedance of about 70 ppb, to extrapolation of the O₃ mole fractions encompassed during the training period. The [LM](#) generally performed well within the O₃ levels covered during the training period. [Across applications, ANNs have been found to be unreliable when extrapolated, due to the nonlinear nature and complexity of the relationships they represent. Though they are generally expected to be more robust to extrapolation than ANNs, increased uncertainty in measurements can also be introduced to LMs when parameters are extrapolated.](#) In order to have [high confidence in measurements of uncommonly high mole fractions of a target gas, the model training period has to encompass the full possible range. Combining both field calibration and lab calibration data together in a training dataset could accomplish this type of coverage. If extrapolation is expected to occur with respect to the target gas mole fraction, as in this case study, the use of an inverted LM may yield better results than LMs or ANNs. We describe inverted LMs and their potential advantages](#) [in our previous work \(Casey et al., 2017\).](#) Keeping in mind this finding about O₃ extrapolation, for ambient measurements in the DJ Basin, for subsequent deployments, we selected field calibration sites that were more representative of distributed sampling site locations, outside of the dense urban environment in downtown Denver, where O₃ did not get as high, likely due to increased titration of O₃ at night in connection with abundant NO_x compounds.

25 3.2.2 Post only calibration moved across the state: [SJ Basin Spring 2015](#)

We also examined model performance that was subject to extrapolation in time and temperature. [We present O₃ model performance data from four U-Pods that were co-located with reference instruments](#)

Joanna Gordon 9/1/18 8:26 PM

Deleted: models

Joanna Gordon 9/1/18 8:27 PM

Deleted:

Joanna Gordon 8/25/18 2:51 PM

Deleted: (when training data consisted of only post test deployment co-locations)

Joanna Gordon 9/29/18 11:19 AM

Deleted: We In the

in the SJ Basin in the spring of 2015, at the Navajo Dam, Sub Station, and Bloomfield sites. Two U-Pods at the Bloomfield site provide a set of duplicate measures. Fig. S4 shows the performance of a number of ANN and LM-based BAO field calibrations with different sets of inputs at this SJ Basin test sites in the spring of 2015, just prior to the summer 2015 BAO training period. U-Pod O₃ was quantified for these deployments using training data from the same co-location period at BAO that was used toward quantification of the summer 2015 SJ Basin deployment, described in section 3.1.

The addition of time as a model input didn't seem to improve the performance of either ANNs or LMs and ANNs generally outperformed LMs. Gas sensor manufactures don't clearly define sensor lifetimes, but sensors are generally expected to loose sensitivity over time. For example, Alphasense CO-B4 electrochemical sensors are expected to have 50% of their original sensitivity after two years (Alphasense, 2015). The heater resistance in a given metal oxide type sensor is expected to drift over time, influencing sensor measurements (e2v Technologies Ltd., 2007). Masson and colleagues observed a significant drift in a metal oxide sensor heater resistance over the course of a 250 day sampling period in a laboratory setting (Masson et al., 2015). While we did not measure and record metal oxide sensor heater resistance for sensors included in U-Pods, we have investigated eltCO₂ and e2vO₃ sensor signal drift from the summer of 2015 through the summer of 2017. These data are presented in Fig. S26. Systematic downward drift in all eltCO₂ sensor signals is apparent over this time frame. A clear and consistent pattern of systematic drift over this time period is less apparent for e2vO₃ sensors. Since the training data was collected immediately after, the test data period, and since the test data period was relatively short (approximately one month) sensor drift could be negligible across the combined training/testing time frame. U-Pods experienced colder temperatures during this spring deployment than were encompassed subsequently in the BAO training period. Linear models generally resulted in more bias than ANNs. Again the model for O₃ that was found to perform best in our previous (Casey et al., 2017), an ANN with temp, absHum and all metal oxide sensor signals as inputs, performed the best at sites included in this case study, with one exception. At the Sub Station site the inclusion of the figCxHy sensor signal decreased model performance. Additionally, the performance of all models tested at the Sub Station site during the SJ Basin Spring 2015 deployment was significantly

Joanna Gordon 8/25/18 2:21 PM

Deleted: we co-located U-Pods with reference instruments in the SJ Basin,

Joanna Gordon 9/30/18 1:03 PM

Deleted: TheInterestingly, t

Joanna Gordon 9/30/18 1:03 PM

Deleted: given

Joanna Gordon 9/30/18 1:03 PM

Deleted: LinearLinear

Joanna Gordon 9/3/18 11:39 AM

Deleted: Unlike our previous findings,

Joanna Gordon 9/3/18 11:40 AM

Deleted: with this training/testing dataset pair,

Joanna Gordon 9/3/18 11:55 AM

Deleted: we found that including some secondary metal oxide type sensors augmented model performance, but not

worse in terms of MBE than model performance at other sites, both LMs and ANNs with different sets of inputs. Since this sensor signal input augmented model performance at the same sampling location during the summer deployment period, this finding could be attributable to the extrapolation with respect to temperature that occurred during the test period of this case study. As discussed in the introduction, metal oxide sensor sensitivity to different gas species can vary along with sensor surface temperature. Models were trained to use the figCxHy sensor signal, across the ambient temperatures in encompassed by the training data, to help account for the influence of confounding gas species at the BAO site. We think it is possible that the different temperatures in combination with the unique mix of gas species present at the Sub Station site, which the figCxHy sensors are highly sensitive to, caused the ANN to perform worse. The Sub Station site is close to two large coal-fired power plants, indicated in Fig. 4 by orange markers in the SJ Basin pane. It is possible that emissions from the San Juan Generating Station (north) and/or the Four Corners Power Plant (south) uniquely influenced the response of this particular Figaro sensor in ways that are not well represented at BAO in the DJ Basin, or present at other SJ Basin sampling sites. Several-hour long enhancements or spikes are apparent in the raw eltCO2 and alphaCO sensor signals in the U-Pod deployed at the Sub Station site, indicating the presence of a near-by combustion-related emissions source. Another indication of a near-field power plant plume across the deployment area is apparent, in the form of several-hour long enhancements of reference measurements of NO and NO₂ at the site.

3.2.3 Post only calibration moved 40 miles across the DJ Basin: BAO Summer 2016

In testing the performance of field calibrations that were generated using data collected at the GRET site in 2017 and applied for the quantification of O₃ at BAO in the 2016, across the DJ Basin, we were interested to find that again, the inclusion of time as a model input did not yield any improvements in calibration equation performance, even though model training took place several months after the test period. Fig. S5 shows the performance of a number of ANN and LM-based GRET field calibrations with different sets of inputs at this BAO test site the previous summer. Another interesting finding from this training/testing dataset pair was that the addition of secondary metal oxide type gas sensors, didn't seem to help improve the performance of field calibration equations either. Fig. S5 shows that ANNs

Joanna Gordon 9/30/18 1:03 PM

Deleted: 1

Joanna Gordon 8/24/18 2:08 PM

Deleted: 2

Joanna Gordon 9/3/18 12:07 PM

Deleted: at other

Joanna Gordon 8/25/18 3:16 PM

Deleted: Temperature or time extrapolation could also be responsible for negating the utility of the figCxHy sensor signal in model performance.

Joanna Gordon 9/30/18 1:04 PM

Deleted: the performance of

Joanna Gordon 9/30/18 1:04 PM

Deleted: and applied forinform

Joanna Gordon 9/29/18 11:20 AM

Deleted: severalso long

Joanna Gordon 8/22/18 3:50 PM

Deleted:

performed better than LMs and that the most useful set of inputs included just e2vO3, temp, and absHum. Similarly, the performance of field calibration equations for CO₂ generated at GRET in 2017 and applied to data from BAO in the summer of 2016, did not seem to be augmented by the inclusion of additional gas sensor signals, though the inclusion of time as a predictor was useful. In the case of CO₂, LMs outperformed ANNs, which could be largely attributable to notable instability associated with the performance of ANNs when time was included as an input. For CO₂, we expected the inclusion of time as an input to be a useful to model performance across this time frame, owing to observed trends of decreased CO₂ sensor sensitivity in time. To keep the power requirements for the U-Pod sensor systems low, and to keep systems quiet, fans were used to exchange air in the enclosures as opposed to pumps. As a result, the air entering the enclosures was not filtered, and sensors were exposed to some dust over time. This dust exposure is likely largely responsible for observed decreases in CO₂ sensors sensitivity over time, shown in Fig. S26. Decreases in infrared lamp intensity over time may also play a role. In the case of CO₂ sensors, the implementation of pumps to draw new, filtered air into sensor enclosures could likely significantly reduce lose rates in the sensitivity of an individual sensor over periods of continuous deployment in ambient environment. While we are not sure why ANN performance tended not to benefit from the addition of a time input, while LM performance did, it is likely attributable to the extrapolation of the time input, since only data that was collected significantly subsequent to the test data period was used for training. ANNs are expected to be able to better represent time decay trends if data from measurements both prior and subsequent to the test period are used in training, so that there is no extrapolation with respect to the time input.

3.2.4 Post only calibration applied to the same location: GRET Fall 2016

To investigate if reduced performance from these GRET to BAO field calibration tests were more connected to the new deployment location or to the significant extrapolation with respect to time of the calibration models, we generated calibration equations based on similarly long training periods at GRET and applied them to data collected prior to the training period at GRET in the fall of 2016. We couldn't draw strong conclusions from this comparison, unfortunately, because of an issue with humidity sensors, described in the methods section and below. Fig. S6 shows the performance of a

Joanna Gordon 8/22/18 4:19 PM

Deleted: with

Joanna Gordon 8/22/18 3:50 PM

Deleted:

number of ANN and LM-based GRET field calibrations with different sets of inputs at the GRET test site during fall of the previous year. For O₃ models, the best performing ANN inputs for this dataset pair were the same ones that we found in our previous study (Casey et al., 2017), with the exception of the humidity signal. The fall 2016 GRET test period coincided with the time period U-Pod absolute humidity was replaced using mixing ratios from a co-located Picarro due to missing humidity sensor data. Interestingly, when this ‘borrowed’ humidity signal was not included as an input, the model performance markedly increased and became competitive with other ‘same location’ test deployment case studies. In our previous work, we showed that O₃ models were very sensitive to the humidity signal input (Casey et al., 2017). In this case study, it seems that replacing actual humidity signals with closely approximated humidity signals, negatively influenced model performance. In order to investigate this observation further, we tested the influence of replacing humidity data in the same manner, using mixing ratios from the same co-located Picarro, on test data from the GRET Spring 2017 case study. A comparison of model performance under normal and this ‘borrowed RH’ circumstance is presented in Fig. S27 in the SM. O₃ model performance was negatively impacted when ‘borrowed’ RH values based on Picarro data replaced U-Pod RH sensor signals. From these findings, it seems likely that the inclusion of multiple metal oxide type sensors as inputs in the model, which all respond strongly to humidity fluctuations, helped the ANN to effectively represent the influence of humidity in the system, more so than including a ‘borrowed RH’ signal from another instrument. We tested models with multiple gas sensor signals and no humidity signal as inputs for a number of other case studies as well (as seen in Fig. S2, Fig. S4, and Fig. S5), when good humidity data from U-Pod enclosures was available, but they did not turn out to be the best performing model in any of these other tests.

3.3 Evaluation of models across training/testing dataset pairs

For each of the case studies, we compare the relative model performance under three governing model-training paradigms. The first of these paradigms includes linear models with only the primary gas sensor signal, along with temperature, and absolute humidity signals as inputs. Performance of these models is shown in Fig. 10. The next paradigm includes models that were found to perform best for each trace gas in our previous work. Performance of these models is shown in Fig. 11. The third

Joanna Gordon 9/30/18 1:04 PM

Deleted: theThe

Joanna Gordon 9/3/18 1:23 PM

Deleted: It seems likely that the inclusion of multiple metal oxide type sensors as inputs in the model, which all respond strongly to humidity fluctuations, helped the ANN to effectively represent the influence of humidity in the system.

Joanna Gordon 9/1/18 7:54 PM

Formatted: Not Highlight

Joanna Gordon 9/3/18 1:25 PM

Deleted: with

Joanna Gordon 9/3/18 1:25 PM

Deleted: humidity

Joanna Gordon 8/27/18 4:41 PM

Deleted: removed

Joanna Gordon 9/30/18 1:05 PM

Deleted: is

paradigm includes models that were optimized for each case study specifically. Performance of these models is shown in Fig. 12. Table 5 and Table 6 show the mean and standard deviation of model performance metrics for each of the case studies presented. Table 7 shows the percent change in model performance metrics when one model-training paradigm is used in place of another, highlighting relative benefits associated with the implementation of different models for O₃ and CO₂.

Fig. 10, Fig. 11, and Fig. 12 contain target plots showing the MBE and CRMSE of models from each dataset pair in terms of absolute mole fractions and mole fractions normalized uniformly by the standard deviation of reference data during the spring 2017 GRET deployment. In the SM, Fig. S23 contains target diagrams equivalent to those presented in Fig. 12, but with individually normalized MBE and CRMSE, according to the standard deviation of reference measurements during each individual test period. The outer circle's radius in each of these target diagrams denotes an error-to-signal ratio of 1. The inner circle's radius in each of these target diagrams encompasses the performance of models when they were tested at the same location that they were trained and when training data bookended the test period, so that there was no extrapolation of the model across time or deployment location. We present our findings in the form of these target diagrams in order to compare our findings with those presented in several particularly relevant previous studies focused on the field calibration of low-cost sensors (Spinelle et al., 2015, 2017; Zimmerman et al., 2017).

Fig. 10 and Fig. 11 show that for CO₂, ANN models generally performed slightly better than LM models with the same set of inputs, though models that were extrapolated more than several months in time were the exception. For O₃, ANNs that included multiple secondary metal oxide sensor signals as inputs were also found to generally perform slightly better than the relatively simple LMs that didn't include any secondary gas sensors as inputs over all (with exceptions for individual case studies). This can be seen in Table 7 and in Fig. 10 and Fig. 11, with all plot markers falling within the outer radius in Fig. 11 (ANNs) but some plot markers falling outside the outer radius in Fig. 10 (LMs). Models that were not moved to a new location for the test period gained the most benefit in their performance when ANNs were used instead of LMs, resulting in a smaller inner radius in the target plots in Fig. 11 relative

- Joanna Gordon 9/3/18 1:49 PM
Deleted: Fig.
- Joanna Gordon 9/29/18 11:43 AM
Deleted: 10,8 is a graphical summary of the performance of quantification models for O₃ and CO₂ for each of the training/testing dataset pairs included in this study to explore how these models hold up when applied to data collected at new locations and in new time frames. The top panels in Fig. 8 show the performance of relatively simple linear models with only three inputs, including just the primary gas sensor for each species, temperature, and humidity. The bottom panels in Fig. 8 show the performance of the ANN regression models and inputs that were found to perform the best on data collected at the GRET site in Greeley, Colorado for O₃ and CO₂ in our previous study (Casey et al., 2017). Fig. 9 is a similar graphical summary of model performance for each dataset pair exar... [7]
- Joanna Gordon 8/24/18 2:09 PM
Deleted: 8
- Joanna Gordon 9/29/18 11:43 AM
Deleted: 12
- Joanna Gordon 8/24/18 2:09 PM
Deleted: 9
- Joanna Gordon 9/29/18 11:44 AM
Deleted: 12
- Joanna Gordon 8/24/18 2:09 PM
Deleted: 8
- Joanna Gordon 9/3/18 1:59 PM
Deleted: - ... [8]
- Joanna Gordon 9/3/18 1:51 PM
Deleted: Table 3 lists the relative circumst... [9]
- Joanna Gordon 9/30/18 1:06 PM
Deleted: 11 show8 shows
- Joanna Gordon 8/27/18 4:48 PM
Deleted: significantly
- Joanna Gordon 9/29/18 11:44 AM
Deleted: 11
- Joanna Gordon 8/24/18 2:10 PM
Deleted: 8
- Joanna Gordon 9/29/18 11:44 AM
Deleted: 11
- Joanna Gordon 8/27/18 6:43 PM
Deleted: in the bottom panel
- Joanna Gordon 9/29/18 11:44 AM
Deleted: Fig.the top panel
- Joanna Gordon 9/29/18 11:45 AM
Deleted: inon the bottom panel

to Fig. 10 for both O₃ and CO₂. The target diagrams in Fig. 10 and Fig. 11 show some degradation in performance when models were applied to data in new locations and when training data took place only after the test period. The of the target plots in Fig. 10 and Fig. 11 demonstrate that bias was introduced when field calibration models were extrapolated in terms of time, when training periods only encompassed data after the test data period and not prior. Interestingly, there are noticeable similarities between the target plots for CO₂ in Fig. 10 and 11 and the target plots for O₃ in Fig. 10 and 11.

The relative performance of models among each training/test dataset pair remained fairly consistent across the different models employed in data quantification. These systematic trends highlight the importance of model training and testing circumstances relative to specific field calibration model types and inputs. For the BAO Summer 2016 case study, when time was extrapolated significantly, and when models were moved across the DJ Basin, CO₂ and O₃ were both better represented by LMs than ANNs. CO₂ and O₃ models did not benefit from additional gas sensors added as inputs either for this case study. In Fig. 11, of models that performed best for each species as determined in our previous study, models that were not extrapolated in time for CO₂, and all O₃ models, plot markers fall within the outer radius, meeting performance standards framed by previous studies (Spinelle et al., 2015, 2017; Zimmerman et al., 2017). In Fig. 12 the best field calibration model performances for each case study all fall within the outer radius, showing good performance, and indicating that incomplete coverage of parameter space in terms of atmospheric chemistry, weather patterns, sampling location, and sampling timing, can be addressed to some extent by tailoring field calibration models and their inputs to specific training/testing datasets pairs.

For CO₂ we found that field calibration models generally extended with good performance to new locations. ANNs outperformed LMs when training took place pre and post of a test deployment. When training only took place after a test deployment LMs performed better. LMs seem to be better at extrapolating in time. Over time, ELT NDIR CO₂ sensors seem to lose sensitivity and/or drift. When CO₂ models were extended back in time, significant bias resulted when time was not included as an input. ANNs were not able to extrapolate in time with any success and their performance became very

- Joanna Gordon 9/3/18 2:02 PM
Deleted: the top panel
- Joanna Gordon 9/29/18 11:45 AM
Deleted: 11
- Joanna Gordon 8/24/18 2:10 PM
Deleted: 8
- Joanna Gordon 9/3/18 2:03 PM
Deleted: before or
- Joanna Gordon 9/3/18 2:03 PM
Deleted: All
- Joanna Gordon 8/27/18 6:44 PM
Deleted: four
- Joanna Gordon 9/29/18 11:45 AM
Deleted: 11
- Joanna Gordon 8/24/18 2:10 PM
Deleted: 8
- Joanna Gordon 9/29/18 11:45 AM
Deleted: Fig. the left panels
- Joanna Gordon 9/3/18 2:04 PM
Deleted: in the right panels

- Joanna Gordon 9/3/18 2:37 PM
Deleted: In Fig.
- Joanna Gordon 9/30/18 1:06 PM
Deleted: 12 the best field calibration model performances for each8, all models that were not extrapolated in time all result in error that falls within the outer circle radius, meeting performance standards framed by previous studies (Spinelle et al., 2015, 2017; Zimmerman et al., 2017). ... [10]
- Joanna Gordon 9/3/18 2:44 PM
Deleted: . Fig. 9 shows
- Joanna Gordon 9/1/18 6:28 PM
Deleted:

unstable when time was added as an input, an occurrence that is apparent in Fig. S5 and Fig. S6. Models performed better when they were extended spatially, all the way across Colorado from the DJ Basin to the SJ Basin, than they did when they were extended back in time. Extension of a LM back in time and across the DJ Basin (from GRET in 2017 to BAO in 2016) resulted in significant MBE
5 relative to the other case studies. The inclusion of multiple additional gas sensors augmented model performance when extended back in time at the same location as training took place, but not at a new location.

For O₃ we found that ANNs with the same set of inputs worked best across a number of case studies,
10 informed by all the metal oxide sensor signals as well as temperature and humidity. The extension of models to new locations often resulted in increased MBE or systematic error, and in some cases increased CRMSE or random error. Some observed bias in the extension of models to new locations could be attributable to different reference instruments with different operators and/or different calibration and data quality measures employed. O₃ model extension to new locations seemed to be
15 more impactful than extension back in time. Interestingly, additional metal oxide sensor signals remained helpful when models were extended all the way across Colorado, from BAO to the SJ Basin, but these additional gas sensor signals did not remain helpful when O₃ models were extended across a county line, from Adams County (CAMP) to Boulder County (Dawson) or from Weld County (GRET) to Boulder County (BAO). ANNs generally performed better than LMs for O₃, with the exception of
20 these two Front Range case studies ([Dawson Summer 2014](#) and [BAO Summer 2016](#)). We found in our previous study that shorter training times led to decreased performance in ANNs and sometimes increased performance in LMs. The training time used in the CAMP to Dawson case study was relatively short, which could have contributed to the superior performance of LMs over ANNs. [For the BAO Summer 2016 case study, both ANN and LM markers are included \(each with the same inputs: e2vO3, temp, and absHum\). LMs had smaller random error but ANNs had smaller bias, highlighting an important consideration in the application of one or the other to inform specific research or measurement goals.](#)
25

4 Conclusions

Several previous studies have shown that multiple gas sensor signals and the implementation of supervised learning techniques can improve the performance of field calibration of low-cost sensors in the quantification of a number of atmospheric trace gas mole fractions. We investigated how well a supervised learning technique (ANNs) held up when sensors were moved to a new location, different from where calibration model training took place. We tested the spatial and temporal transferability of field calibration models for O₃ and CO₂ under a number of different circumstances using data from multiple reference instrument co-locations, using the same sensors over the course of several years, when sensors were deployed in two oil and gas production basins, along with urban and peri-urban sites. We found that the best performing field calibration models for both O₃ and CO₂ were not consistent across all training and testing deployment pairs, though some patterns emerged in terms of model type and inputs in association with the spatial and temporal extension of calibration equations, from training to testing performed in oil and gas production areas. The performance of O₃ models generally benefited from the inclusion of multiple metal oxide sensor signals in addition to the primary e2vO3 sensor signal, while the performance of CO₂ models relied more heavily on temperature and humidity inputs. CO₂ model performance was impacted more by temporal extension than spatial extension. In contrast, O₃ model performance was impacted more by spatial extension than temporal extension.

While ANNs and other supervised learning techniques have been shown to consistently out perform linear models in previous studies when training and testing took place in the same location, we find that this trend does not always hold when field calibration models are applied in a new location, with significantly different local emissions source signatures for O₃ models, or when model training data takes place more than several months subsequent to the model application period for CO₂ models. We found that the implementation of calibration models that were well suited to specific training and test data pairs resulted in generally good test performance in terms of centered root mean squared error and mean biased error, scaled by reference measurement standard deviation, reported in target diagrams in previous studies. For example, when models were significantly extrapolated in time and transferred to a

Joanna Gordon 9/3/18 2:54 PM
Deleted: these
Joanna Gordon 9/30/18 1:06 PM
Deleted: heolds hold
Joanna Gordon 9/30/18 1:06 PM
Deleted: weareare

Joanna Gordon 9/29/18 11:50 AM
Deleted: appliedmoved in terms of
Joanna Gordon 8/22/18 12:20 PM
Deleted: spatial
Joanna Gordon 8/22/18 12:20 PM
Deleted: and
Joanna Gordon 9/30/18 1:06 PM
Deleted: place
Joanna Gordon 9/29/18 11:50 AM
Deleted: models.
Joanna Gordon 8/22/18 12:21 PM
Deleted: temporal coverage

[new location](#), a well-suited set of sensor inputs would generally not include secondary gas sensor signals.

LMs with just one primary gas sensor signal as well as temperature and humidity were found to outperform ANNs when models were applied to a location with different dominating sources of pollution [in the case of O₃](#), like Downtown Denver relative to eastern Boulder County. These three-input LMs also outperformed ANNs when models were significantly extrapolated in time. While these LMs seemed to be more stable under circumstances of significant extrapolation in terms of local air chemistry and timing, we found that they [did not extrapolate well in terms of the O₃ mole fraction, resulting in underproduction of O₃ values during the test period that exceeded those encompassed in the training data](#).

Joanna Gordon 8/22/18 12:12 PM

Deleted: were not able to fully represent some of the complex nonlinear response behavior exhibited by the arrays of sensors.

Field calibration models tested in new locations often resulted in the introduction of additional bias relative to field calibration models that were tested in the same location they were trained in. [As seen in Fig. 12, plot markers from all case studies have very similar CRMSE values, but plot markers from case studies in which models were tested in new locations have larger MBE values than models that were tested in the same location as they were trained](#). Finding ways to effectively mitigate bias associated with new field deployment locations would further improve the performance of field calibrations toward quantification of atmospheric trace gases using arrays of low-cost sensors. Such improvements in the field of low-cost sensors will help to enable dense distributed networks of low-cost sensors to inform air quality in oil and gas production basins. [The following findings from this work, and associated recommendations, are made to help inform the logistics of future studies that employ field calibration methods of low-cost gas sensors](#).

Joanna Gordon 9/29/18 11:50 AM

Deleted: AsThis trend is nicely show

Joanna Gordon 9/30/18 1:06 PM

Deleted: 12, plot9as the data is almost a band running vertically in a range of CRMSEs

1. **Finding:** For O₃ models, LMs perform better than ANNs when the chemical composition of local emissions sources is significantly different in the model-training location relative to the model-application location. We found that when models were trained in an urban area with significant mobile sources, then tested in a peri-urban area, more strongly influenced by oil and

Joanna Gordon 8/22/18 10:56 AM

Deleted: .

gas emissions, the differences in local sources of pollution were significantly different enough that LMs outperformed ANNs. Alternatively, when models were trained in one oil and gas production region and tested in another the different composition of local emissions (lighter vs. heavier hydrocarbons) was not significant enough for LM performance to surpass the performance of ANNs, though some positive bias was evident in predicted O₃ mole fractions.

Explanation: ANNs are very effective at compensating for the influence of interfering gas species through pattern recognition of a training dataset. However, if different patterns, in terms of the relative abundance of various oxidizing and reducing compounds in the air, are present in the testing location relative to the training location, ANNs may not be able to compensate for the influence of interfering gas species as effectively. The relative abundance of interfering oxidizing and reducing compounds are not included as model parameters, but ANN performance is challenged by these circumstances.

Recommendation: When measuring O₃ or other gas species with a metal oxide type sensor, if the nature of dominant emissions sources at the model training location is significantly different than the nature of dominant emissions sources in the model application location, use a LM instead of an ANN. For the best performance, try to train models in locations with similar emissions sources to a desired sampling location. If the nature of dominant emissions sources at the model training and application locations are similar, signals from an array of multiple unique metal oxide sensors will likely augment model performance.

2. **Finding:** For CO₂ models, LMs perform better than ANNs when model training occurs significantly (more than several months) prior to or subsequent to the model application period.

Explanation: CO₂ sensors drift over time in terms of sensitivity and baseline response. When models are extrapolated in time (when training takes place more than several months prior or subsequent to the model application period), ANN performance can be compromised to a greater extent than LM performance. ANNs are able to represent relationships during training very effectively, and with significant more complexity and nonlinear relationships among time and other model inputs than LMs. The more complex the model, the less likely it can be extrapolate

effectively. LMs, with no interaction terms like we employ in this work, are not able to fit data and potentially complex patterns inherent in sensor drift over time during training as closely as an ANN, but the simple linear relationships they represent between the time input and the target gas mole fraction over the course of training are more likely to hold prior or subsequent to the training period.

Recommendation: When measuring CO₂ with a NDIR sensor, if model-training data is only available more than several months prior or subsequent to the model application period, use a LM instead of an ANN. For the best model performance, use training data that is collected directly pre or post of the model application period, and preferably data from both pre and post of the model application period. Training models using data from both pre and post of a given model application period helps models to encompass sensor drift over time as well as increases the likelihood of covering the full range of environmental parameter space that occurs during the model application period so that extrapolation of these parameters is avoided.

3. **Finding:** Extrapolation of an O₃ or CO₂ model in time, and especially significant extrapolation in time, can change both the type of model that is most effective, as well as the specific model input signals that are most effective.

Explanation: Low-cost sensors change over time, both in terms of their baseline response and in terms of their sensitivity to target and interfering gas species. Different sensor types drift due to different physical phenomenon so further a generalization across sensor types is difficult.

Recommendation: Use training data collected directly pre and post of the model application period in order to implement a 'best performing model' for each gas species that can be applied using data from different model training and application pairs.

4. **Finding:** ANNs yield less bias and more accurate gas mole fraction quantification than LMs, even when transferred to a new location under the following circumstances: (a) extrapolation of training parameters is avoided during the model application period, (b) training takes place for

several weeks to a month prior and subsequent to the model application period, and (c) the dominant local emissions sources are similar in the model training and application locations.

Explanation: Our previous study and multiple other ambient and laboratory based experiments have shown, arrays of low-cost sensors in combination with ANN regression models can support useful quantification of gases in mixtures and in the ambient environment because ANNs can more effectively represent complex nonlinear relationships among environmental variables and signals in a sensor system like a U-Pod than LMs. With this work, we have explored limitations associated with these methods when challenged in different ways, as we present with a number of case studies.

Recommendation: If minimizing error and bias in measurements of gas mole fractions using low-cost sensors systems is a primary goal, design sensor system training and field deployment experiments so that extrapolation of model training parameters is avoided during the model application period, so that training takes place for several weeks to a month directly prior and directly subsequent to the model application period, and so that the dominant local emissions sources are similar in the model training and application locations. When these conditions are satisfied, ANNs can be robustly implemented, with better performance than LMs.

It is also imperative that sensor users keep in mind the primary importance of minimizing extrapolation of temperature, humidity and sensor signal from model training to application. We show that field

normalization trace gas quantification models can more readily be transferred across a large state from one oil and gas production to another, than from an urban to oil and gas production basin that are in closer proximity to each other. We also show that pre and post model training, directly prior to and after field site deployment, is generally much more effective than pre or post model training alone, especially when the training takes place significantly before or after the deployment period. Along with

these findings and general guidelines for future studies, we recommend further validation efforts in the extension of quantification of atmospheric trace gases using low-cost gas sensor arrays in oil and gas production basins and toward other ambient measurement applications. The findings presented here may be applicable and generalizable in the use of low-cost metal oxide, electrochemical, and non-

dispersive infrared sensor arrays in various configurations and sampling regions to characterize mole fractions of a number of atmospheric trace gases. Future studies exploring the sensitivity of our findings to these factors are recommended. [In order to account for unique variations in sensor responses, in each individual sensor system, due to variations in manufacturing along with elapsed time and specific exposure subsequent to manufacturing, we present models that are generated for each sensor system on an individual basis. Future studies exploring the potential for universal calibration models would be very useful to the field.](#)

The authors declare that they have no conflict of interest.

Acknowledgements

The many low-cost sensor and reference instrument measurements that facilitated this study were made possible with the gracious help of a number of agencies and individuals. We'd like to thank Bradley Rink and Erick Mattson of CDPHE, Katherine Benedict and Jeff Collett of CSU, Detlev Helmig and Jacques Hueber of INSTAAR, Gaby Pétron, Jon Kofler, Audra McClure, Bruce Batram, and Daniel Wolfe of NOAA, Michael King of NEPA, Christopher Ellis and Andrew Switzer of the SUIP AQP, along with Joe Cotie and Roman Szkoda of the NM AQB for sharing reference instrument data with us and facilitating our U-Pod measurements. We thank Jana Milford for assistance in the preliminary analysis and reporting of results from the Boulder County Study. Thanks to John Ortega for preparing U-Pods and arranging permissions and logistics for the Boulder County project in 2014, including the Dawson School site we use data from in this work. Thanks are also due to include Brianna Yepa, Victoria Danner, Tasha Nez, Rebecca Bullard, Madeline Polmear, and Bryce Goldstien for helping to maintain U-Pods in the field as well as downloading and organizing data. Bryce Goldstien made some maps in ArcMap of the SJ and DJ Basins with data from the Colorado Oil and Gas Conservation Commission that were adapted and presented here. It was a pleasure working with all these interesting and helpful people. We thank Jana Milford for assistance in the preliminary analysis and reporting of results from the Boulder County Study as well as useful feedback and suggestions in review of this manuscript. Many thanks also to Shelly Miller, Marina Vance, and Christopher Ellis for

kindly reviewing this work which was funded by Boulder County and the National Science Foundation Air Water Gas Sustainability Research Network under grant number CBET-1240584.

References

- Abeleira, A. J. and Farmer, D. K.: Summer ozone in the northern Front Range metropolitan area: Weekend-weekday effects, temperature dependences, and the impact of drought, *Atmos. Chem. Phys.*, 17(11), 6517–6529, doi:10.5194/acp-17-6517-2017, 2017.
- Adgate, J. L., Goldstein, B. D. and McKenzie, L. M.: Potential Public Health Hazards, Exposures and Health Effects from Unconventional Natural Gas Development, *Environ. Sci. Technol.*, A-N, doi:dx.doi.org/10.1021/es404621d, 2014.
- 10 Ahmadov, R., McKeen, S., Trainer, M., Banta, R., Brewer, A., Brown, S., Edwards, P. M., De Gouw, J. A., Frost, G. J., Gilman, J., Helmig, D., Johnson, B., Karion, A., Koss, A., Langford, A., Lerner, B., Olson, J., Oltmans, S., Peischl, J., Pétron, G., Pichugina, Y., Roberts, J. M., Ryerson, T., Schnell, R., Senff, C., Sweeney, C., Thompson, C., Veres, P. R., Warneke, C., Wild, R., Williams, E. J., Yuan, B. and Zamora, R.: Understanding high wintertime ozone pollution events in an oil- and natural gas-
- 15 producing region of the western US, *Atmos. Chem. Phys.*, 15(1), 411–429, doi:10.5194/acp-15-411-2015, 2015.
- Alphasense: Technical Specification CO-B4, , 44(0), 0–1, 2015.
- Brudzewski, K.: Gas analysis system composed of a solid-state sensor array and hybrid neural network structure, *Sensors And Actuators*, 55, 38–46, doi:10.1016/S0925-4005(99)00040-4, 1999.
- 20 Casey, J. G., Collier-Oxandale, A. and Hannigan, M. P.: Performance of Artificial Neural Networks and Linear Models To Quantify 4 Trace Gas Species In an Oil and Gas Production Region with Low-Cost Sensors, *Sensor. Actuat. B-Chem.*, 2017.
- Cheadle, L. C., Oltmans, S. J., Pétron, G., Schnell, R. C., Mattson, E. J., Herndon, S. C., Thompson, A. M., Blake, D. R. and McClure-begley, A.: Surface ozone in the Northern Front Range and the influence
- 25 of oil and gas development on ozone production during FRAPPE / DISCOVER-AQ, , 6, 2017.
- Clements, A. L., Griswold, W. G., RS, A., Johnston, J. E., Herting, M. M., Thorson, J., Collier-Oxandale, A. and Hannigan, M.: Low-Cost Air Quality Monitoring Tools: From Research to Practice

- (A Workshop Summary), *Sensors*, 17(11), 2478, doi:10.3390/s17112478, 2017.
- Cross, E. S., Lewis, D. K., Williams, L. R., Magoon, G. R., Kaminsky, M. L., Worsnop, D. R. and Jayne, J. T.: Use of electrochemical sensors for measurement of air pollution : correcting interference response and validating measurements, *Atmos. Meas. Tech. Discuss.*, 2017–138(May), 1–17, doi:10.5194/amt-2017-138, 2017.
- 5 e2v Technologies Ltd.: MiCS-2611 O3 Sensor, Datasheet, 44(August), 4–6, 2007.
- Edwards, P. M., Young, C. J., Aikin, K., deGouw, J. a., Dubé, W. P., Geiger, F., Gilman, J. B., Helmig, D., Holloway, J. S., Kercher, J., Lerner, B., Martin, R., McLaren, R., Parrish, D. D., Peischl, J., Roberts, J. M., Ryerson, T. B., Thornton, J., Warneke, C., Williams, E. J. and Brown, S. S.: Ozone
- 10 photochemistry in an oil and natural gas extraction region during winter: simulations of a snow-free season in the Uintah Basin, Utah, *Atmos. Chem. Phys. Discuss.*, 13(3), 7503–7552, doi:10.5194/acpd-13-7503-2013, 2013.
- Edwards, P. M., Brown, S. S., Roberts, J. M., Ahmadov, R., Banta, R. M., deGouw, J. a., Dubé, W. P., Field, R. a., Flynn, J. H., Gilman, J. B., Graus, M., Helmig, D., Koss, A., Langford, A. O., Lefter, B. L.,
- 15 Lerner, B. M., Li, R., Li, S.-M., McKeen, S. a., Murphy, S. M., Parrish, D. D., Senff, C. J., Soltis, J., Stutz, J., Sweeney, C., Thompson, C. R., Trainer, M. K., Tsai, C., Veres, P. R., Washenfelder, R. a., Warneke, C., Wild, R. J., Young, C. J., Yuan, B. and Zamora, R.: High winter ozone pollution from carbonyl photolysis in an oil and gas basin, *Nature*, doi:10.1038/nature13767, 2014.
- Field, R. A., Soltis, J., McCarthy, M. C., Murphy, S. and Montague, D. C.: Influence of oil and gas field
- 20 operations on spatial and temporal distributions of atmospheric non-methane hydrocarbons and their effect on ozone formation in winter, *Atmos. Chem. Phys.*, 15(6), 3527–3542, doi:10.5194/acp-15-3527-2015, 2015.
- Frankenberg, C., Thorpe, A. K., Thompson, D. R., Hulley, G., Kort, E. A., Vance, N., Borchardt, J., Krings, T., Gerilowski, K., Sweeney, C., Conley, S., Bue, B. D., Aubrey, A. D., Hook, S. and Green, R.
- 25 O.: Airborne methane remote measurements reveal heavy-tail flux distribution in Four Corners region, *Proc. Natl. Acad. Sci.*, 113(35), 201605617, doi:10.1073/pnas.1605617113, 2016.
- Gilman, J. B., Lerner, B. M., Kuster, W. C. and De Gouw, J. A.: Source signature of volatile organic compounds from oil and natural gas operations in northeastern Colorado, *Environ. Sci. Technol.*, 47(3),

- 1297–1305, doi:10.1021/es304119a, 2013.
- Gulbag, A. and Temurtas, F.: A study on quantitative classification of binary gas mixture using neural networks and adaptive neuro-fuzzy inference systems, *Sensors Actuators, B Chem.*, 115(1), 252–262, doi:10.1016/j.snb.2005.09.009, 2006.
- 5 Hagan, M. T., Demuth, H. B., Beale, M. H. and De Jesus, O.: *Neural Network Design*, PWS Publishing Co., Boston, MA., 1997.
- Huyberegts, G. and Szeco, P.: Simultaneous quantification of carbon monoxide and methane in humid air using a sensor array and an artificial neural network, *Sensor. Actuat. B-Chem.*, 45, 123–130, 1997.
- Kort, E. a., Frakenberg, C., Costigan, K. R., Lindenmaier, R., Dubey, M. K. and Wunch, D.: Four
10 corners: The largest US Methane anomaly viewed from space, , 6898–6903, doi:10.1002/2014GL061503.Received, 2014.
- Kumar, P., Morawska, L., Martani, C., Biskos, G., Neophytou, M., Di Sabatino, S., Bell, M., Norford, L. and Britter, R.: The rise of low-cost sensing for managing air pollution in cities, *Environ. Int.*, 75, 199–205, doi:10.1016/j.envint.2014.11.019, 2015.
- 15 Larasati, A., DeYong, C. and Slevitch, L.: Comparing Neural Network and Ordinal Logistic Regression to Analyze Attitude Responses, *Serv. Sci.*, 3, 304–312, doi:10.1007/978-3-642-19193-0, 2011.
- Licor: The Importance of Water Vapor Measurements and Corrections, *Licor Appl. Note #129*, 1–8 [online] Available from: <https://www.licor.com/documents/igs56gijkc4ftks30pci>, 2010.
- Martín, M. A., Santos, J. P. and Agapito, J. A.: Application of artificial neural networks to calculate the
20 partial gas concentrations in a mixture, *Sensors Actuators, B Chem.*, 77(1–2), 468–471, doi:10.1016/S0925-4005(01)00736-5, 2001.
- Masson, N., Piedrahita, R. and Hannigan, M.: Approach for quantification of metal oxide type semiconductor gas sensors used for ambient air quality monitoring, *Sensors Actuators, B Chem.*, 208, 339–345, doi:10.1016/j.snb.2014.11.032, 2015.
- 25 McClure-Begley, A., Petropavlovskikh, I. and Oltmans, S.: NOAA Global Monitoring Surface Ozone Network: BAO Tower, , doi:http://dx.doi.org/10.7289/V57P8WBF, n.d.
- McDuffie, E. E., Edwards, P. M., Gilman, J. B., Lerner, B. M., Dubé, W. P., Trainer, M., Wolfe, D. E., Angevine, W. M., DeGouw, J., Williams, E. J., Tevlin, A. G., Murphy, J. G., Fischer, E. V., McKeen,

- S., Ryerson, T. B., Peischl, J., Holloway, J. S., Aikin, K., Langford, A. O., Senff, C. J., Alvarez, R. J., Hall, S. R., Ullmann, K., Lantz, K. O. and Brown, S. S.: Influence of oil and gas emissions on summertime ozone in the Colorado Northern Front Range, *J. Geophys. Res.*, 121(14), 8712–8729, doi:10.1002/2016JD025265, 2016.
- 5 McKenzie, L. M., Guo, R., Witter, R. Z., Savitz, D. A., Newman, L. S. and Adgate, J. L.: Research | Children ' s Health Birth Outcomes and Maternal Residential Proximity to Natural Gas Development in Rural Colorado, *Environ. Health Perspect.*, (4), 412–417, doi:10.1289/ehp.1306722, 2014.
- McKenzie, L. M., Witter, R. Z., Newman, L. S. and Adgate, J. L.: Human health risk assessment of air emissions from development of unconventional natural gas resources., *Sci. Total Environ.*, 424, 79–87, doi:10.1016/j.scitotenv.2012.02.018, 2012.
- 10 McKenzie, L. M., Allshouse, W. B., Byers, T. E., Bedrick, E. J., Serdar, B. and Adgate, J. L.: Childhood hematologic cancer and residential proximity to oil and gas development, *PLoS One*, 12(2), e0170423, doi:10.1371/journal.pone.0170423, 2017.
- Niebling, G.: Identification of gases with classical pattern-recognition methods and artificial neural networks, *Sensors Actuators B. Chem.*, 18(1–3), 259–263, doi:10.1016/0925-4005(94)87091-8, 1994.
- 15 Niebling, G. and Schlachter, A.: Qualitative and quantitative gas analysis with non-linear interdigital sensor arrays and artificial neural networks, *Sensors Actuators B. Chem.*, 27(1–3), 289–292, doi:10.1016/0925-4005(94)01603-F, 1995.
- Olaguer, E. P.: The potential near-source ozone impacts of upstream oil and gas industry emissions, *J. Air Waste Manage. Assoc.*, 62(8), 966–977, doi:10.1080/10962247.2012.688923, 2012.
- 20 Oltmans, S. J., Karion, A., Schnell, R. C., Pétron, G., Helmig, D., Montzka, S. A., Wolter, S., Neff, D., Miller, B. R., Hueber, J., Conley, S., Johnson, B. J. and Sweeney, C.: O₃, CH₄, CO₂, CO, NO₂ and NMHC aircraft measurements in the Uinta Basin oil and gas region under low and high ozone conditions in winter 2012 and 2013, *Elem. Sci. Anthr.*, (2), 12, doi:10.12952/journal.elementa.000132, 25 2016.
- Penza, M. and Cassano, G.: Application of principal component analysis and artificial neural networks to recognize the individual VOCs of methanol/2-propanol in a binary mixture by SAW multi-sensor array, *Sensors Actuators, B Chem.*, 89(3), 269–284, doi:10.1016/S0925-4005(03)00002-9, 2003.

- Piedrahita, R., Xiang, Y., Masson, N., Ortega, J., Collier, A., Jiang, Y., Li, K., Dick, R. P., Lv, Q., Hannigan, M. and Shang, L.: The next generation of low-cost personal air quality sensors for quantitative exposure monitoring, *Atmos. Meas. Tech.*, 7(10), 3325–3336, doi:10.5194/amt-7-3325-2014, 2014.
- 5 Reza Nadafi, D. B., Nejad, S. N., Kabganian, M. and Barazandeh, F.: Neural network calibration of a semiconductor metal oxide micro smell sensor, *Symp. Des. Test, Integr. Packag. MEMS/MOEMS, DTIP 2010*, 5–8, 2010.
- Schnell, R. C., Oltmans, S. J., Neely, R. R., Endres, M. S., Molenar, J. V. and White, A. B.: Rapid photochemical production of ozone at high concentrations in a rural site during winter, *Nat. Geosci.*, 10 2(2), 120–122, doi:10.1038/ngeo415, 2009.
- Spinelle, L., Gerboles, M., Villani, M. G., Aleixandre, M. and Bonavitacola, F.: Field calibration of a cluster of low-cost available sensors for air quality monitoring. Part A: Ozone and nitrogen dioxide, *Sensors Actuators, B Chem.*, 215, 249–257, doi:10.1016/j.snb.2015.03.031, 2015.
- Spinelle, L., Gerboles, M., Villani, M. G., Aleixandre, M. and Bonavitacola, F.: Field calibration of a cluster of low-cost commercially available sensors for air quality monitoring. Part B: NO, CO and CO₂, 15 *Sensors Actuators, B Chem.*, 238(2), 706–715, doi:10.1016/j.snb.2016.07.036, 2017.
- Srivastava, A. K.: Detection of volatile organic compounds (VOCs) using SnO₂ gas-sensor array and artificial neural network, *Sensors Actuators, B Chem.*, 96(1–2), 24–37, doi:10.1016/S0925-4005(03)00477-5, 2003.
- 20 Sun, L., Wong, K. C., Wei, P., Ye, S., Huang, H., Yang, F., Westerdahl, D., Louie, P. K. K., Luk, C. W. Y. and Ning, Z.: Development and application of a next generation air sensor network for the Hong Kong marathon 2015 air quality monitoring, *Sensors (Switzerland)*, 16(2), doi:10.3390/s16020211, 2016.
- Sundgren, H., Winquist, F., Lukkar, I. and Lundstrom, I.: Artificial neural networks and gas sensor arrays: quantification of individual components in a gas mixture, *Meas. Sci. Technol.*, 2, 464–469, 25 1991.
- De Vito, S., Massera, E., Piga, M., Martinotto, L. and Di Francia, G.: On field calibration of an electronic nose for benzene estimation in an urban pollution monitoring scenario, *Sensors Actuators, B*

Chem., 129(2), 750–757, doi:10.1016/j.snb.2007.09.060, 2008.

De Vito, S., Piga, M., Martinotto, L. and Di Francia, G.: CO, NO₂ and NO_x urban pollution monitoring with on-field calibrated electronic nose by automatic bayesian regularization, *Sensors Actuators, B Chem.*, 143(1), 182–191, doi:10.1016/j.snb.2009.08.041, 2009.

5 Wang, C., Yin, L., Zhang, L., Xiang, D. and Gao, R.: Metal oxide gas sensors: Sensitivity and influencing factors, *Sensors*, 10(3), 2088–2106, doi:10.3390/s100302088, 2010.

Xiong, L. and Compton, R. G.: Amperometric gas detection: A review, *Int. J. Electrochem. Sci.*, 9(12), 7152–7181, 2014.

10 Zimmerman, N., Presto, A. A., Kumar, S. P. N., Gu, J., Hauryliuk, A., Robinson, E. S., Robinson, A. L. and Subramanian, R.: Closing the gap on lower cost air quality monitoring: machine learning calibration models to improve low-cost sensor performance, *Atmos. Meas. Tech. Discuss.*, (2), 1–36, doi:10.5194/amt-2017-260, 2017.

15

20

25

Joanna Gordon 9/29/18 11:56 AM
Deleted: - ... [11]
Joanna Gordon 9/29/18 11:56 AM
Formatted: Normal, Indent: Left: 0"
Joanna Gordon 9/29/18 11:56 AM
Formatted: Font:11 pt

Table 1: Gas sensors included in U-Pods along with the model input codes for each. The input code is an abbreviation for the make of the sensor, followed by the target gas species(s).

Sensor Type	NDIR	Metal Oxide					Electrochemical
Target Gas(s)	CO ₂	CH ₄ *	CxHy **	O ₃	VOCs	CO	CO
Model	S300	TGS 2600	TGS 2602	MiCs-2611	MiCs-5521	MiCs-5525	CO-B4
Make	ELT	Figaro	Figaro	e2v/SGX	e2v/SGX	e2v/SGX	Alphasense
Code	eltCO2	figCH4	figCxHy	e2vO3	e2vVOC	e2vCO	alphaCO

*Light hydrocarbons

**Heavy Hydrocarbons

Joanna Gordon 9/30/18 1:08 PM
 Deleted: for
 Joanna Gordon 9/30/18 1:08 PM
 Deleted: is

5

10

15

20

Table 2: Reference instrument measurements at U-Pod sampling sites

Deployment	Reference Instrument	Calibration	Operator	Res
Ozone				
CAMP	Teledyne API 400E	quarterly cal/nightly quality checks	CDPHE	1
Dawson	Thermo Electron 49	pre cal/post cal check	INSTAAR	5
BAO*	Thermo Scientific 49c	annual cal/monthly quality checks	NOAA	60
Navajo Dam	Thermo Scientific 49i	quarterly cal/weekly quality checks	NM AQB	1
Bloomfield	Thermo Scientific 49i	quarterly cal/weekly quality checks	NM AQB	1
Sub Station	Thermo Scientific 49i	quarterly cal/weekly quality checks	NM AQB	1
Ignacio	Thermo Scientific 49is	monthly cal/weekly quality checks	SUIT AQP	1
Bondad	Thermo Scientific 49is	monthly cal/weekly quality checks	SUIT AQP	1
Shiprock	Teledyne API T400	quarterly cal/monthly quality checks	NEPA	60
Fort Lewis	2b Technologies 202	factory cal/post cal check	CU Boulder	1
GRET	Teledyne API T400E	quarterly cal/nightly quality checks	CDPHE	1
Carbon Dioxide				
BAO	Picarro G2401		NOAA	1
SJ Basin	LI-COR LI-840A	pre + post cal: zero precision span	CU Boulder	1
GRET	Picarro G2508	periodic zero stability checks	CSU	1

*(McClure-Begley et al., n.d.) Res = Time resolution of measurements in minutes

5

10

Table 3: Best performing models, as determined for each gas species, in the previous study (Casey et al., 2017)


Gas Species	Model Type	Sensor Signal Model Inputs
CO ₂	ANN	eltCO2 (ELT S300 CO2 Sensor) temp (temperature) absHum (absolute humidity)
O ₃	ANN	e2vO3 (e2v MiCs-2611) e2vCO (e2v MiCs-5525) e2vVOC (e2v MiCs-5521) figCH4 (Figaro TGS 2600) figCxHy (Figaro TGS 2602) temp (temperature) absHum (absolute humidity)

5

10

15

20



5 Table 4: Relative timing and parameter coverage between model training and test deployment dataset pairs. Incomplete coverage of time occurred if training only took place before or after the test data period and not before and after (pre and post). Incomplete coverage of location occurred when training took place in one location and testing took place in another. Incomplete coverage of parameters, or extrapolation of models, including the target gas mole fraction, temperature, time, and pressure occurred when the values observed during training did not encompass the values observed during testing. Extrapolation in time occurred when training only took place after the test period (post model training timing). Extrapolation in location occurred when a model was trained in one location then applied to data collected in a new location.

Case Study	Summary	Training Timing	Extrapolation During Test
Dawson Summer 2014	Urban calibration moved to rural/peri-urban setting	Pre/Post	Location, O ₂
SJ Basin Spring 2015	DJ Basin calibration moved across the state to SJ Basin sampling sites	Post	Location, Pressure, Time
SJ Basin Summer 2015	DJ Basin calibration moved across the state to SJ Basin sampling sites	Pre/Post	Location, Pressure
BAO Summer 2015	DJ Basin calibration applied to same location	Pre/Post	None
BAO Summer 2016	DJ Basin calibration moved 40 miles across the DJ Basin	Post	Location, Time
GRET Fall 2016	DJ Basin calibration applied to same location	Post	Time
GRET Spring 2017	DJ Basin calibration applied to same location	Pre/Post	None

Microsoft Office User 9/21/18 5:47 AM
 Deleted: across the

10

15

20

Table 5: O₃ model performance metrics.

Case Study	N	R ²	RMSE (ppb)	MBE (ppb)	Standard Deviation R ²	Standard Deviation n RMSE	Standard Deviation MBE
O₃ Models							
Best O₃ Model (Casey et al., 2017)							
ANN with inputs: e2vO3 temp absHum e2vVOC e2vCO FigCH4 FigCxHy							
Dawson Summer 2014	1	0.83	6.46	-0.91	0.00	0.00	0.00
SJ Basin Spring 2015	4	0.86	7.74	3.69	0.05	3.82	5.78
SJ Basin Summer 2015	7	0.85	7.03	4.89	0.10	1.10	1.73
BAO Summer 2015	2	0.93	4.26	1.45	0.00	0.31	0.07
BAO Summer 2016	2	0.92	12.21	-11.14	0.00	0.31	0.07
GRET Fall 2016	2	0.96	12.87	12.02	0.01	2.30	2.35
GRET Spring 2017	2	0.98	2.59	1.49	0.00	0.69	1.02
Simple Model (Single Gas Sensor)							
LM with inputs: e2vO3 temp absHum							
Dawson Summer 2014	1	0.95	3.59	-0.46	0.00	0.00	0.00
SJ Basin Spring 2015	4	0.83	17.95	16.09	0.06	6.10	5.83
SJ Basin Summer 2015	7	0.86	6.30	3.53	0.06	1.40	2.06
BAO Summer 2015	2	0.87	5.50	0.94	0.00	0.78	1.56
BAO Summer 2016	2	0.89	5.78	-2.71	0.00	0.78	1.56
GRET Fall 2016	2	0.93	12.73	11.92	0.01	0.62	0.88
GRET Spring 2017	2	0.89	6.00	-3.19	0.00	0.73	1.38
Models Optimized For Case Studies							
Dawson Summer 2014	1	0.95	3.59	-0.46	0.00	0.00	0.00
SJ Basin Spring 2015	4	0.86	7.74	3.69	0.05	3.82	5.78
SJ Basin Summer 2015	7	0.85	7.03	4.89	0.10	1.10	1.73
BAO Summer 2015	2	0.93	4.26	1.45	0.02	0.51	1.54
BAO Summer 2016	2	0.87	6.25	-0.20	0.02	0.51	1.54
GRET Fall 2016	2	0.95	3.99	2.14	0.00	0.28	0.89
GRET Spring 2017	2	0.98	2.59	1.49	0.00	0.69	1.02

Table 6: CO₂ model performance metrics.

Case Study	N	R ²	RMSE (ppm)	MBE (ppm)	Standard Deviation R ²	Standard Deviation RMSE	Standard Deviation MBE
CO₂ Models							
Best CO₂ Model from (Casey et al., 2017)							
ANN with inputs: eltCO2 temp absHum							
SJ Basin Summer 2015	2	0.65	8.42	-0.62	0.00	1.81	1.41
BAO Summer 2015	2	0.75	9.98	-2.60	0.05	13.00	13.89
BAO Summer 2016	2	0.69	54.38	48.37	0.05	13.00	13.89
GRET Fall 2016	2	0.74	42.37	39.58	0.02	2.44	2.57
GRET Spring 2017	2	0.83	6.31	0.59	0.03	0.13	2.61
Simple Model (Single Gas Sensor)							
LM with inputs: eltCO2 temp absHum							
SJ Basin Summer 2015	2	0.71	7.84	0.27	0.01	1.43	0.42
BAO Summer 2015	2	0.69	10.62	-1.26	0.06	1.52	10.67
BAO Summer 2016	2	0.73	11.82	0.73	0.06	1.52	10.67
GRET Fall 2016	2	0.82	8.62	-3.46	0.00	0.69	1.45
GRET Spring 2017	2	0.55	9.88	-0.33	0.03	0.29	1.91
Models Optimized For Case Studies							
SJ Basin Summer 2015	2	0.72	7.45	-0.11	0.04	2.06	0.31
BAO Summer 2015	2	0.80	8.85	-2.29	0.10	6.47	7.08
BAO Summer 2016	2	0.73	11.82	0.73	0.06	1.52	10.67
GRET Fall 2016	2	0.82	8.62	-3.46	0.00	0.69	1.45
GRET Spring 2017	2	0.83	6.31	0.59	0.03	0.13	2.61

5

10

Table 7: Relative benefits associated with the implementation of different models for O₃ and CO₂.

Case Study	Mean % Increase in R ²	Mean % Decrease in RMSE	Mean % Decrease in MBE	Mean % Increase in R ²	Mean % Decrease in RMSE	Mean % Decrease in MBE
CO ₂ Models			O ₃ Models			
Benefit of Models Optimized for Case Studies Over the Best Models from (Casey et al., 2017)						
Dawson Summer 2014				14.51	44.42	50.00
SJ Basin Spring 2015				0.00	0.00	0.00
SJ Basin Summer 2015	10.56	11.52	82.60	0.00	0.00	0.00
BAO Summer 2015	5.84	11.27	11.95	0.00	0.00	0.00
BAO Summer 2016	5.72	78.27	98.49	-5.01	48.82	98.19
GRET Fall 2016	11.17	79.66	108.73	-0.54	68.99	82.22
GRET Spring 2017	0.00	0.00	0.00	0.00	0.00	0.00
Benefit of the Best Models from (Casey et al., 2017) Over Simple Linear Models						
Dawson Summer 2014				-12.67	-79.92	-99.99
SJ Basin Spring 2015				3.20	56.88	77.09
SJ Basin Summer 2015	-8.41	-7.29	331.39	-1.34	-11.53	-38.41
BAO Summer 2015	8.70	6.05	-106.48	6.79	22.48	-53.85
BAO Summer 2016	-5.41	-360.09	-6543.84	2.57	-111.22	-310.71
GRET Fall 2016	-10.05	-391.73	1244.99	2.88	-1.12	-0.86
GRET Spring 2017	51.92	36.13	278.55	10.00	56.90	146.65
Benefit of Models Optimized for Case Studies Over Simple Linear Models						
Dawson Summer 2014				0.00	0.00	0.00
SJ Basin Spring 2015				3.20	56.88	77.09
SJ Basin Summer 2015	1.26	5.06	140.25	-1.34	-11.53	-38.41
BAO Summer 2015	15.04	16.64	-81.80	6.79	22.48	-53.85
BAO Summer 2016	0.00	0.00	0.00	-2.57	-8.10	92.59
GRET Fall 2016	0.00	0.00	0.00	2.33	68.64	82.07
GRET Spring 2017	51.92	36.13	278.55	10.00	56.90	146.65

Microsoft Office User 9/24/18 6:38 AM
Deleted: F
Microsoft Office User 9/24/18 6:38 AM
Deleted: T

Microsoft Office User 9/24/18 6:38 AM
Deleted: T

Microsoft Office User 9/24/18 6:38 AM
Deleted: F

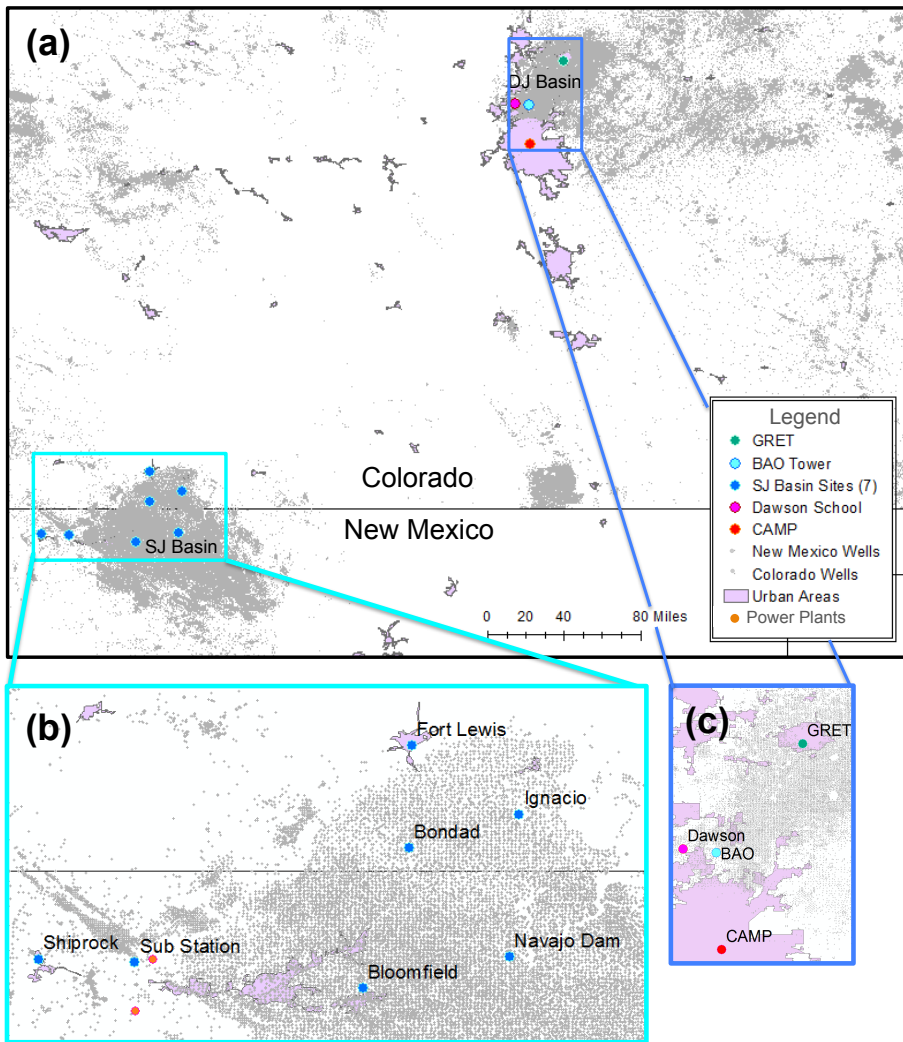



Figure 1: (a) Training and test deployment locations are identified in the SJ and DJ Basins in context with urban centers and oil and gas production wells. (b) Panel zoomed in on the SJ Basin, covering approximately 4,250 square miles (85x50 miles). (c) Panel zoomed in on the DJ Basin covering approximately 1,540 square miles (28x55 miles).



(b) Case Study	Training Location	Test Location	O ₃ # U-Pods	O ₃ U-Pod Names	CO ₂ # U-Pods	CO ₂ U-Pod Names
Dawson Summer 2014	CAMP	Dawson	1	BE	NA	NA
SJ Basin Spring 2015	BAO	SJ Basin	4	BB, BD, BF, BJ	NA	NA
SJ Basin Summer 2015	BAO	SJ Basin	7	BA, BB, BD, BE, BF, BH, BI	2	BB, BD
BAO Summer 2015	BAO	BAO	2	BC, BJ	2	BC, BJ
BAO Summer 2016	GRET	BAO	2	BH, BI	2	BH, BI
GRET Fall 2016	GRET	GRET	2	BH, BI	2	BH, BI
GRET Spring 2017	GRET	GRET	2	BH, BI	2	BF, BI

5 Figure 2: (a) ANN and LM training and test deployment timelines. The Dawson, BAO, and GRET sampling sites are all located in the DJ Basin. Model training periods for each test deployment are shown in blue, and model test periods are shown in magenta. For the BAO Summer 2016 case study, the period outlined in blue shows data that was used to train O₃ model, but not CO₂ models since CO₂ reference data was not available during winter months. (b) Information about each of the case studies presented in the above timelines, including model training and testing locations, as well as the number and names of U-Pods included in each case study for both O₃ and CO₂ models. The U-Pods with names shown in grey were constructed and deployed starting in May of 2014. The U-Pods with names shown in black were constructed and deployed starting in April of 2015.

10

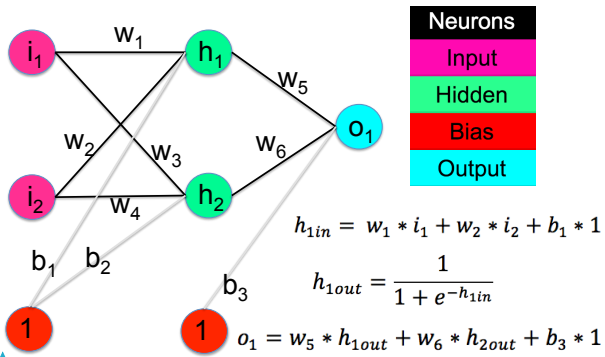


Figure 3. Example of a simple feed forward neural network, showing how inputs are propagated through the network during each of the training iterations (Casey et al., 2017)

Unknown
Formatted: Font:11 pt, Bold

5

10

15

20




Figure 4. Diagram of an example ANN with the same color-coded components as are presented in Figure SM3 in section 2.2 of the SM. This ANN has 5 inputs, 1 hidden layer with 5 tansig hidden neurons, and 1 linear output layer leading to 1 output. The network is fully connected with weights and biases (Casey et al., 2017).

5

10

15

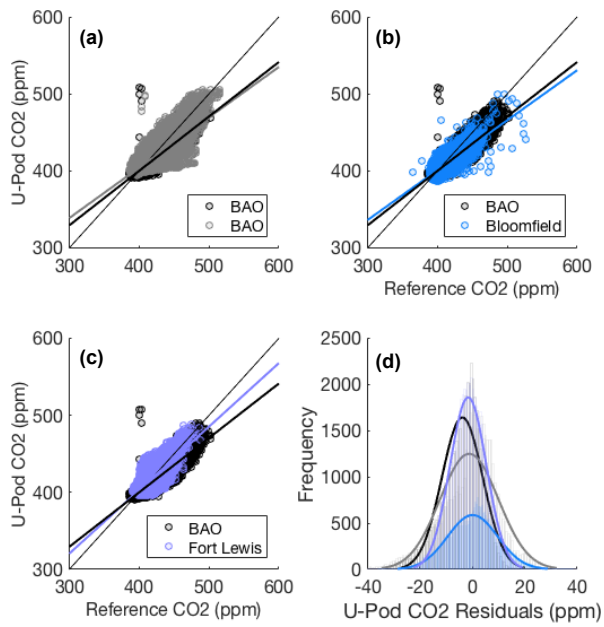


Figure 5: Scatter plots of U-Pod CO₂ vs. reference CO₂ and overlaid histograms of U-Pod CO₂ residuals for (a) BAO and BAO (b) BAO and Bloomfield (c) BAO and Fort Lewis. A 1:1 single-weight reference line is included in each scatter plot along with double-weight lines of best fit for U-Pods at each sampling location. Data from U-Pod BC at BAO is plotted in black along with U-Pods BJ, BB, and BD at BAO, Fort Lewis, and Bloomfield, respectively. Sensor signal inputs include eltCO₂, temp, and absHum. (d) Overlaid histograms of model residuals with respect to reference CO₂.

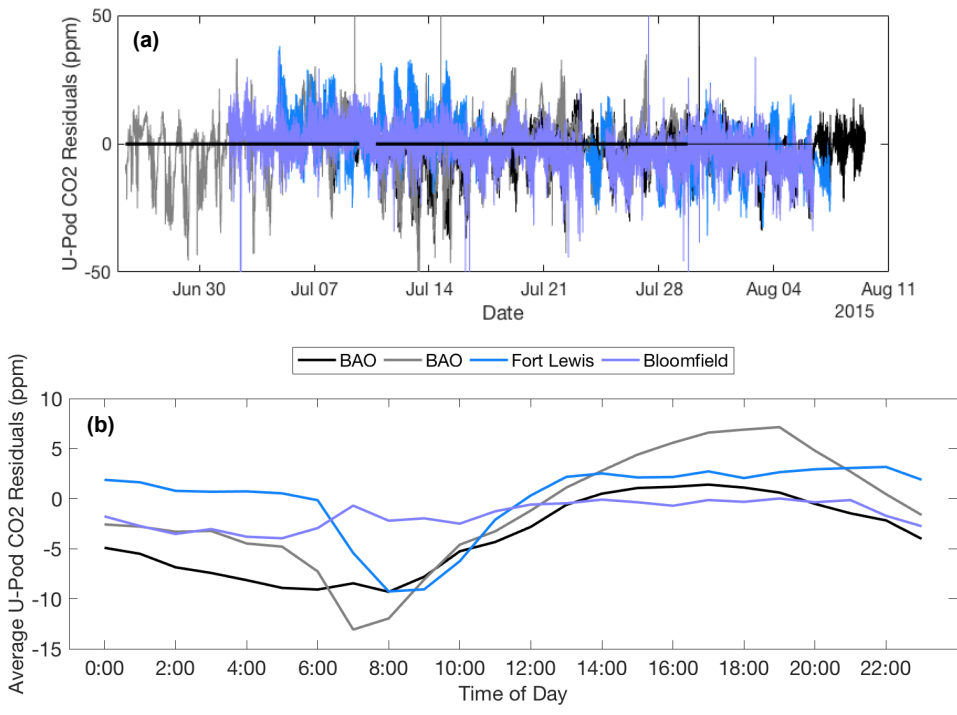


Figure 6: U-Pod CO₂ residuals by (a) data and (b) time of day and throughout the duration of the deployment period. Sensor signal inputs include eltCO₂, temp, and absHum.

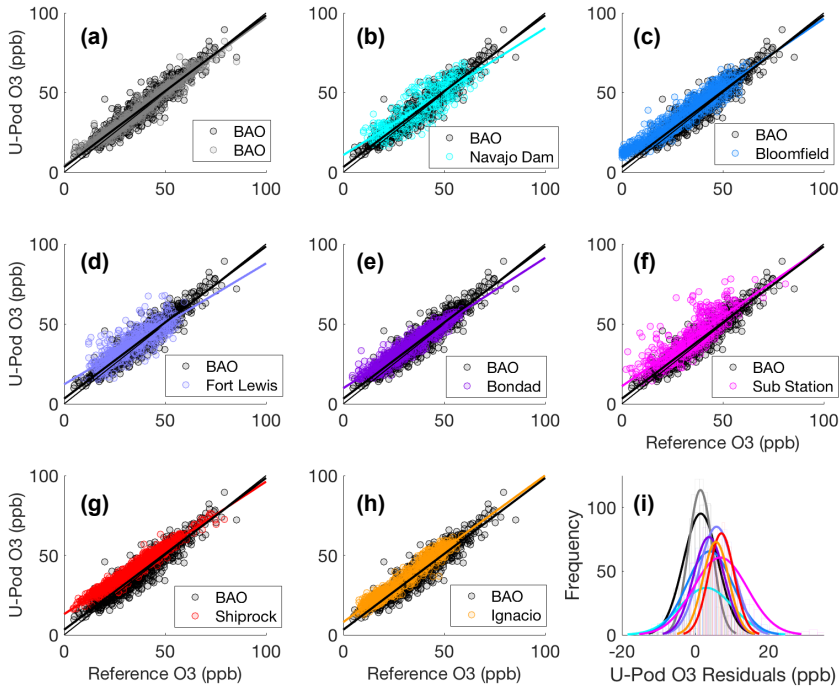


Figure 7: Scatter plots of U-Pod vs reference O₃, comparing U-Pod BC at BAO, in black, with (a) U-Pod BJ at BAO (b) U-Pod BA at Navajo Dam (c) U-Pod BB at Fort Lewis (d) U-Pod BD at Bloomfield (e) U-Pod BE at Bondad (f) U-Pod BF at the Sub Station (g) U-Pod BH at Shiprock and (h) U-Pod BI at Ignacio. (i) Overlaid histograms of model residuals with respect to reference O₃.

5

Microsoft Office User 9/24/18 6:41 AM
Deleted: in cyan

10

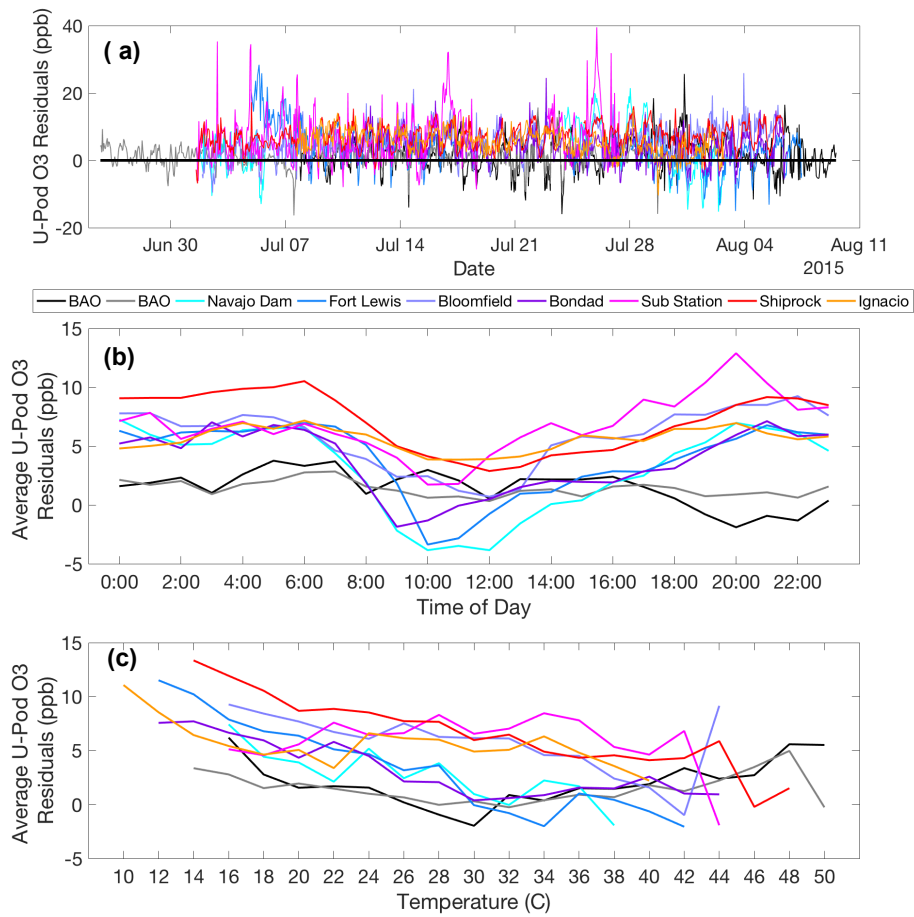


Figure 8: Residuals of U-Pod O₃ spanning of the deployment period, by (a) date (b) time of day and (c) temperature.

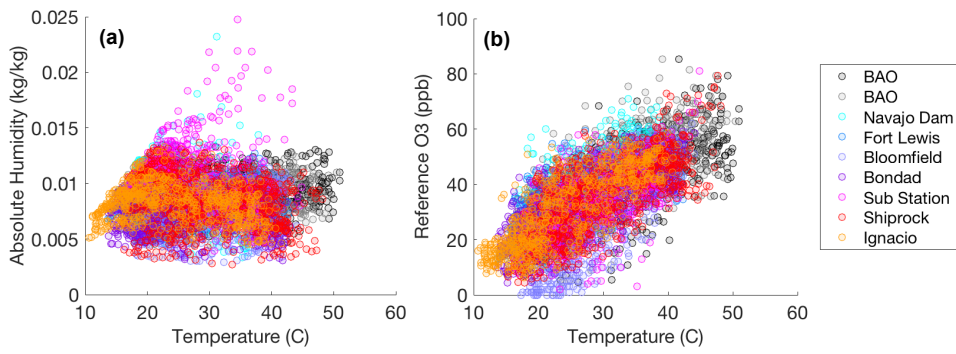


Figure 9: Scatter plots showing the combined parameter space of (a) absolute humidity with temperature and (b) reference O₃ with temperature for each of the U-Pod sampling sites at BAO and the SJ Basin.

5

10

15

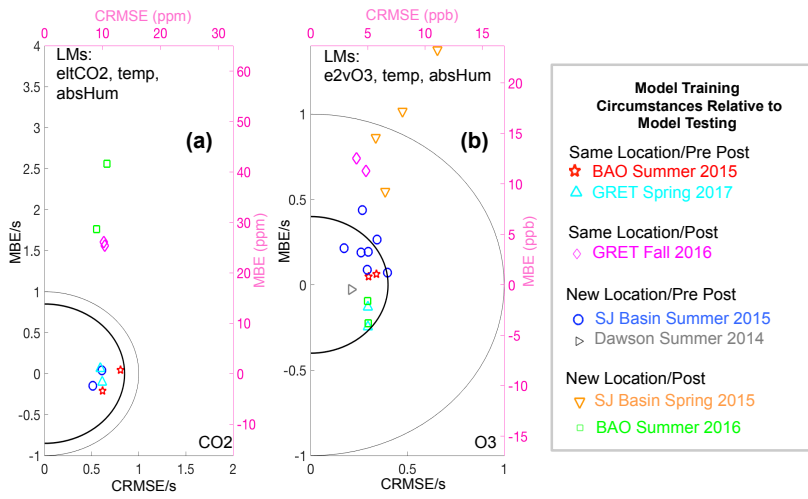


Figure 10: Target diagrams demonstrating performance of a previously determined best-performing model across all new test datasets. (a) CO₂ and (b) O₃ LM performance when only the primary gas sensor, temperature and humidity are inputs. (c) CO₂ and (d) O₃ ANN performance with inputs that were found to perform best at the GRET site in the spring of 2017 (Casey et al., 2017). Model input definitions: eltCO2 (ELT S300 CO₂ sensor), e2vO3 (e2v MiCs-2611 sensor), temp (temperature), and absHum (absolute humidity).

5

10

- Microsoft Office User 9/24/18 6:44 AM
Formatted: Subscript
- Microsoft Office User 9/24/18 6:45 AM
Formatted: Subscript
- Microsoft Office User 9/24/18 6:45 AM
Formatted: Subscript
- Microsoft Office User 9/24/18 6:45 AM
Formatted: Subscript
- Microsoft Office User 9/24/18 6:45 AM
Formatted: Subscript
- Microsoft Office User 9/24/18 6:44 AM
Deleted:
- Microsoft Office User 9/24/18 6:45 AM
Formatted: Subscript
- Microsoft Office User 9/24/18 6:44 AM
Deleted:

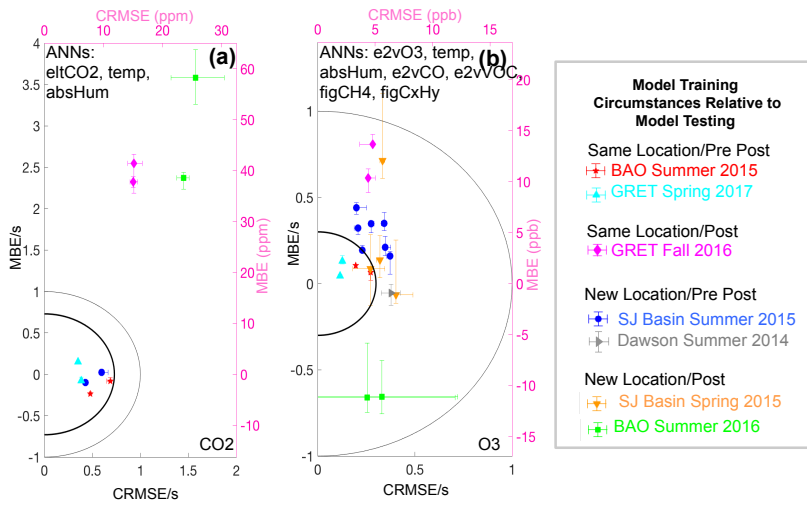
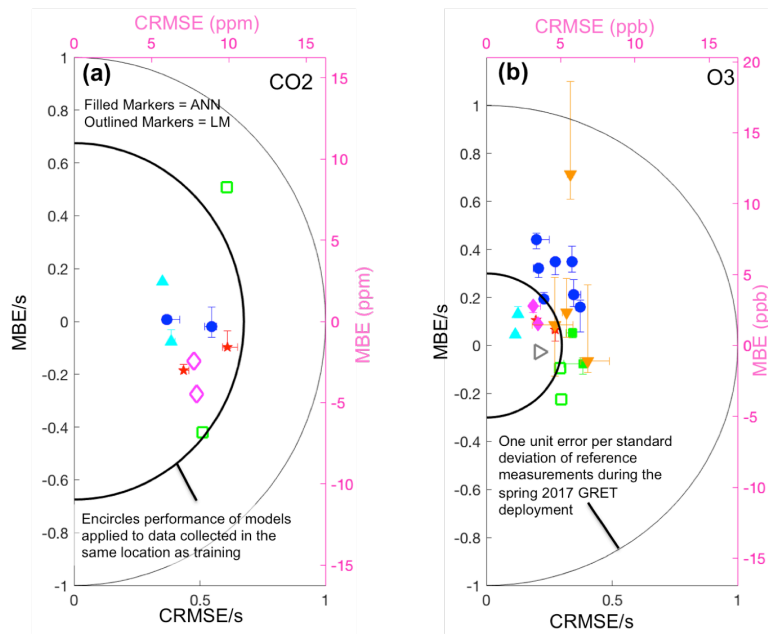


Figure 11: Target diagrams demonstrating performance of a previously determined best-performing model across all new test datasets (a) CO₂ and (b) O₃ ANN performance with inputs that were found to perform best at the GRET site in the spring of 2017 (Casey et al., 2017). Model input definitions: eltCO₂ (ELT S300 CO₂ sensor), e2vCO (e2v MiCs-5525 sensor), e2vVOC (e2v MiCs-5521 sensor), e2vO₃ (e2v MiCs-2611 sensor), figCH₄ (Figaro TGS 2600 sensor), figCxHy (Figaro TGS 2602 sensor), temp (temperature), and absHum (absolute humidity).

5

- Microsoft Office User 9/24/18 6:47 AM
Formatted: Subscript
- Microsoft Office User 9/24/18 6:47 AM
Formatted: Subscript
- Microsoft Office User 9/24/18 6:47 AM
Formatted: Subscript
- Microsoft Office User 9/24/18 6:47 AM
Deleted:



Case Study	CO2 Model Inputs	O3 Model Inputs	Description
BAO Summer 2015	eltCO2 temp rh e2VOC alphaCO	e3VO3 temp absHum e2VOC e2vCO figCH4 figCxHy	Same Location/Pre Post
SJ Basin Summer 2015	eltCO2 temp rh e2VOC alphaCO	e3VO3 temp absHum e2VOC e2vCO figCH4 figCxHy	New Location/Pre Post
SJ Basin Spring 2015		e3VO3 temp absHum e2VOC e2vCO figCH4 figCxHy	New Location/Post
GRET Spring 2017	eltCO2 temp absHum	e3VO3 temp absHum e2VOC e2vCO figCH4 figCxHy	Same Location/Pre Post
BAO Summer 2016	eltCO2 temp absHum time	e3VO3 temp absHum	New Location/Post
GRET Fall 2016	eltCO2 temp absHum rh e2VOC alphaCO figCH4 figCxHy e2vCO time	e3VO3 temp e2VOC e2vCO figCH4 figCxHy	Same Location/Post
Dawson Summer 2014		e3VO3 temp absHum	New Location/Pre Post

Figure 12: Target diagrams for (a) CO₂ and (b) O₃ calibration model performance for the best performing model for each particular case when tested on data from a number of field deployments. Model input definitions: eltCO2 (ELT S300 CO2 sensor), e2vCO (e2v MiCs-5525 sensor), e2vVOC (e2v MiCs-5521 sensor), e2vO3 (e2v MiCs-2611 sensor), figCH4 (Figaro TGS 2600 sensor), figCxHy (Figaro TGS 2602 sensor), alphaCO (Alphasense CO-B4 sensor) temp (temperature), absHum (absolute humidity), rh (relative humidity), and time (absolute time).

5