

Evaluation and environmental correction of ambient CO₂ measurements from a low-cost NDIR sensor

Cory R. Martin¹, Ning Zeng^{1,2}, Anna Karion³, Russell R. Dickerson^{1,2}, Xinrong Ren^{1,4}, Bari N. Turpie¹, Kristy J. Weber^{1,5}

5 ¹Department of Atmospheric and Oceanic Science, University of Maryland, College Park, MD 20742, USA

²Earth System Science Interdisciplinary Center, University of Maryland, College Park, MD 20742, USA

³National Institute of Standards and Technology, Gaithersburg, MD 20899, USA

⁴Air Resources Laboratory, National Oceanic and Atmospheric Administration, College Park, MD 20740, USA

10 ⁵Now at: Department of Geography, University of Colorado at Boulder, Boulder, CO 80309, USA

Correspondence to: Ning Zeng (zeng@umd.edu)

Abstract. Non-dispersive infrared (NDIR) sensors are a low-cost way to observe carbon dioxide concentrations in air, but their specified accuracy and precision are not sufficient for some scientific applications. An initial evaluation of six SenseAir K30 carbon dioxide NDIR sensors in a lab setting
15 showed that without any calibration or correction, the sensors have an individual root mean square error (RMSE) between ~5 to 21 parts per million (ppm) compared to a research-grade greenhouse gas analyzer using cavity enhanced laser absorption spectroscopy. Through further evaluation, after correcting for environmental variables with coefficients determined through a multivariate linear regression analysis, the calculated difference between the each of six individual K30 NDIR sensors and the higher-precision
20 instrument had an RMSE of between 1.7 ppm and 4.3 ppm for one minute data. The median RMSE improved from 9.6 for off the shelf sensors to 1.9 ppm after correction and calibration, demonstrating the potential to provide useful information for ambient air monitoring.

1 Introduction

Carbon dioxide (CO₂) is a major greenhouse gas, with fundamental importance to Earth's climate. Since
25 measurements started at the Mauna Loa Observatory in the 1950s (Keeling et al., 2005), the global mean concentration of CO₂ has steadily risen from the preindustrial mole fraction of approximately 280 μmol mol⁻¹ of dry air (parts per million, or ppm) to today's level exceeding 400 ppm. These observations, both from flask samples and state-of-the-art continuous measurement instruments, have a typical compatibility goal of ~0.1 ppm, recommended for observations at background global network sites (World
30 Meteorological Organization, 2013). Flask-based measurements require observers to collect samples, which are subsequently transported to a lab for analysis, at significant cost. Continuous in-situ CO₂ analyzers located at towers do not suffer from these regular costs, but these high-precision analyzers can cost

upwards of \$100,000 per site, plus any additional costs for calibration gases and installation of equipment and inlet lines. High-accuracy CO₂ observations are thus relatively sparse compared to other climatological variables such as temperature and precipitation.

5 Recent research efforts have focused more locally and on the use of networks of observing sites that use instrumented towers similar to what is used for global monitoring, but applied to the urban environment (Pataki et al., 2003; Briber et al., 2013; Kort et al., 2013; McKain et al., 2012; Turnbull et al., 2015). High-accuracy observations from these tower sites are then used to create inversions to estimate the total greenhouse gas flux from the urban area in question (McKain et al., 2012; Bréon et al., 2015; Lauvaux et al., 2016). However, due to the cost of these networks being comparable to ones at the global scale, the observation towers are still sited at a relatively low density of typically 3 to 12 sites in a single metropolitan area (McKain et al., 2012; Kort et al., 2013; Turnbull et al., 2015; Bréon et al., 2015). Observing system simulation experiments have found that, depending on the methodology used, a higher spatial density of observations in these urban regions has been shown to better constrain the inversion estimates, even if the absolute uncertainty of the observations is higher (Turner et al., 2016; Wu et al., 2016; Lopez-Coto et al., in press), but a trade-off between total network cost and inversion constraint must be balanced.

Recently, a wave of small, low-cost sensors, some of which measure trace gases or particulate matter, in addition to traditional meteorological variables, using various technologies have become commercially available. Evaluation and implementation of some of these new low-cost sensors demonstrate their promise for ambient air monitoring. (Eugster and Kling, 2012; Holstius et al., 2014; Piedrahita et al., 2014; Young et al., 2014; Wang et al., 2015; Shusterman et al., 2016). Many of these instruments are based on electrochemical reactions to measure the concentrations of trace gases. With the advent of widely available and low cost mid-IR light sources and detectors, a small group of non-dispersive infrared (NDIR) CO₂ sensors have also become commercially available. They are designed for use in a number of applications including ventilation control, agricultural and industrial applications, and inclusion in stand-alone commercial products. Additionally, with the high volume of possible applications, these small NDIR CO₂ sensors are affordably priced on the order of \$100 to \$200 per sensor. Previous studies have compared some of these NDIR CO₂ devices and concluded that after application of some type of calibration procedure, some of these devices can provide reasonably accurate measurements (± 3 -5ppm) of ambient CO₂ concentrations (Hurst et al., 2011; Yasuda et al., 2012; Shusterman et al., 2016).

In this paper, one of these small NDIR CO₂ devices is assessed by determining its accuracy with and without environmental corrections. Section 2 describes the CO₂ sensor and its Allan variance, the other

instruments included in the system, and the data collection and processing methodology. Section 3 describes the calibration and shows the stability of the reference high-precision gas analyzer, and the initial results from the NDIR sensor are shown in Sect. 4. In Section 5, two methods are described to determine functional relationships and coefficient values to correct the observed values of the instrument for environmental variables and Sect. 6 discusses the potential utility of observations from this sensor after correction and temporal averaging.

2 Instruments and methods

To test the validity of using low-cost sensors for scientific applications, a sensor package was implemented consisting of various off-the-shelf components. The K30 sensor module (K30) from SenseAir (Sweden), is the low-cost NDIR CO₂ observing instrument used in this study¹. The K30 is a microprocessor-controlled device with on-board signal averaging, has a measurement range of 0 to 10,000 ppm, observation frequency of 0.5 Hz, and resolution of 1 ppm. The manufacturer's stated accuracy of the K30 sensor is ± 30 ppm ± 3 % of reading (SenseAir, 2007) for the 0.5Hz raw output. Additional NDIR sensors were initially evaluated before selecting the K30, including the COZIR ambient sensor and Telaire T6615, which have manufacturer specified accuracies of ± 50 ppm ± 3 % and ± 75 ppm respectively (Gas Sensing Solutions, 2014; General Electric, 2011). The K30 was chosen not only because it has the highest manufacturer-specified accuracy, but also because initial testing showed reliability and consistency when compared to higher-quality observations. In addition to CO₂, temperature, relative humidity, and pressure readings are recorded using a breakout board purchased from Adafruit. This board features a Bosch Sensortec BME280, which according to the manufacturer's datasheet has an average absolute accuracy of ± 1 °C, ± 3 %, and ± 1 hPa, and an output resolution of 0.1 °C, 0.008 %, and 0.01 hPa for temperature, relative humidity, and pressure, respectively (Bosch Sensortec, 2015).

To compare the performance of the K30 to better-performing research instrumentation, a greenhouse gas analyzer based on cavity enhanced absorption spectrometry (CEAS) was used as the control. The LGR-24A-FGGA fast greenhouse gas analyzer from Los Gatos Research (LGR, San Jose, CA) provides

¹Certain commercial equipment, instruments, or materials are identified in this paper in order to specify the experimental procedure adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the materials or equipment identified are necessarily the best available for the purpose.

CO₂, CH₄, as well as water vapor mixing ratios at a frequency of 0.5 Hz and has an un-calibrated uncertainty of < 1 % (Los Gatos Research, 2013). The LGR was connected to a tee connection, to allow either ambient air or a calibration source (during calibrations) to be sampled continuously by the analyzer at a flow rate of 400 standard mL min⁻¹. Calibrations for CH₄ and CO₂ were conducted using several NIST-
5 certified standard mixtures every 23 to 47 hours for a period of one month with molar mixing ratios ranging from 1869.6 parts per billion (ppb) to 2159.4 ppb for CH₄ and from 369.19 ppm to 429.68 ppm for CO₂. See Sect. 3 for details and results of this calibration period.

It is important to note that there are differences in how CEAS works compared to NDIR, most notably that
10 the LGR and other CEAS instruments have a controlled cavity where pressure and temperature are kept nearly constant (with a standard deviation of under 0.5 torr and 0.1 °C for 2-second data), removing potential environmental interference and the need for corrections, whereas the NDIR K30 works in the ambient environment without any mechanism for keeping temperature or pressure constant. Additionally, the LGR implements a water vapor correction on its greenhouse gas concentrations to estimate the dry gas
15 mixing ratio, while the K30 makes no water vapor corrections. A difference between the two analyzers with regard to their sensitivity to the isotopes of CO₂ is expected to be small because the standards used to calibrate the LGR account for all CO₂ isotopes. To increase the effective path length, both the K30 and LGR use mirrors, but the LGR system uses highly reflective mirrors that allow for an effective path length that is many times longer than that of the K30. Additionally, the CEAS instrument determines the
20 concentration of a gas by how long it takes for the signal to degrade inside the cavity (the e-folding time), whereas an NDIR sensor merely measures the intensity of the signal received relative to the total intensity emitted.

For data collection, a Raspberry Pi (RPi) computer is used (Raspberry Pi Foundation, 2015). The RPi is a
25 credit card sized (approximately 6 x 9 cm) computer running a full Linux distribution, allowing for easy customization and usability, that is priced at around \$25. The K30 is connected to the RPi over Universal Asynchronous Receiver/Transmitter (UART) Serial, and the BME280 over Inter-Integrated Circuit (I²C) serial. An image of the complete sensor package is available in Fig. 1. Data is archived on the RPi and uploaded to a centralized data storage and processing server. The LGR collects and archives its own data,
30 but an RPi is used here as well to collect the data from the LGR over a local area network and transfer it to the same centralized server. The added computational power of a Raspberry Pi over traditional data loggers allows for the ability to archive two levels of data: the raw data collected every two seconds, and one-minute averages.

Archiving and comparing multiple datasets proved to be challenging, so steps are taken to ensure that each compared value is at the same observed time. All of the RPi's use an internet server to synchronize their time, and the LGR uses an internal clock with battery that was set to the same time as the RPi's at the beginning of the experiment. Because of various complications including the exact LGR start time and the potential for delays in the RPi's Linux operating system, the data collection times of each K30 sensor package and the LGR are asynchronous. Additionally, power issues can corrupt parts of the plain text data files stored on the RPi's SD card with random characters. Thus, a post-processing procedure has been developed that filters extraneous characters, and then each dataset is synchronized based on recorded time stamps and averaged over selected time periods. These new datasets can then be directly compared without missing or out of phase data points.

2.1 K30 Allan variance

Allan variance (Allan, 1966) is a measure of the time-averaged stability between consecutive measurements or observations, often applied to clocks and oscillators. In addition, an Allan variance analysis can be used to determine the optimum averaging interval for a dataset to minimize noise without sacrificing signal. Figure 2 shows the Allan deviation (the square root of the variance) for one K30's raw two-second data when exposed to a known reference gas. The original two-second data shows the maximum noise, with a standard deviation comparable to the manufacturer's specifications of ± 30 ppm, but averaging for even ten seconds drops the variance significantly. According to this analysis, the optimum averaging time, when the Allan variance is at a minimum (Langridge et al., 2008), is approximately three minutes; longer averaging times do not reduce the noise. The other sensors were found to perform similarly. For the subsequent analysis, an averaging time of one minute is used, as the Allan variance is only slightly higher than for three minutes, and one minute observations allow for resolution of atmospheric variability at shorter time scales.

2.2 Experiment

The need to quickly and effectively evaluate a relatively large number of sensors under conditions with relatively stable CO₂ led to the use of a rooftop observation room on the University of Maryland campus in College Park, Maryland. Because this rooftop room had limited access, and it was not part of the building's HVAC system, it served as an ambient evaluation chamber with minimal influence from human respiration. The room was slightly ventilated for the entire evaluation period to allow outside air to slowly diffuse into the room, with a small household box fan also in the room to ensure that the air was well mixed. The room

also features a small, independent heating and cooling unit, but it was only used to keep the room from exceeding a certain temperature, thus the room was not fully temperature controlled. Even with this control, the diurnal fluctuations of temperature in the room were similar to that of the outdoor environment. This ventilation strategy was intentional so that the room then mimicked the ambient CO₂ concentration of the surrounding atmosphere, and approximated the outdoor temperature and humidity, while protecting instruments from direct sunlight, extreme temperatures, and inclement weather. This provided an advantage over controlled tests in a laboratory setting in that rather than just a multi-point calibration, comparing datasets over ambient concentrations and environmental conditions allowed for a realistic evaluation of these instruments in more real world scenarios.

10

For a continuous period of approximately four weeks in spring 2016, six K30 sensor packages as described in Sect. 2 were deployed alongside the LGR in the rooftop room, all sampling room air. The LGR was also connected to a mass flow controller and standard tank to periodically provide a reference for stability (details in Sect. 3). For the reference dataset, the dry CO₂ (CO_{2,dry}) output calculated by the LGR was used. This output includes an applied correction to the mole fraction of CO₂ to give the dry air mole fraction in ppm. The raw CO₂ values were recorded from each K30, temperature and pressure were recorded from each BME280 sensor, and water vapor mole fraction was also recorded by the LGR. All of the observations were recorded every two seconds, and averaged into one minute values. The next two sections describe the stability of the LGR as well as the initial comparison between the K30 and LGR observations.

20 **3 Los Gatos evaluation and correction**

To evaluate the K30 NDIR sensor performance compared to a research-grade analyzer, first the control dataset needs to be calibrated and corrected for drift. To calibrate the LGR, after the experiment concluded the dataset was corrected using a two-point calibration curve derived from using two NIST-traceable gas standards, one with a CO₂ mole fraction of 369.19 ppm, and the other with a mole fraction of 429.68 ppm. A linear fit was then assumed between the two calibration points, with the recorded values as the dependent variable and the NIST-assigned tank values as the independent variable. In addition, three cylinders of breathing air with higher CO₂ mole fractions of 449.73, 486.53, and 516.41ppm (that are NIST-traceable) were also previously used to calibrate the LGR and showed its linearity. Once the coefficients were determined, the entire LGR dataset was then corrected for further analysis.

30

In addition to the calibration described above, there was a need to quantify any drift in the LGR analyzer. During the experiment period, the LGR was attached to a tee connector, which pulled ambient air from the

aforementioned evaluation chamber using its included pump most of the time, but received periodic calibration every 23 to 47 hours for a period of one hour, initially, and later, ten minutes, to conserve the tank, using a reference tank of breathing air connected to a Dasibi Model 5008 calibrator, which was used to schedule the input of calibration gas. This breathing air tank is assumed to have a fixed CO₂ mole fraction, which was estimated by using the LGR to be 463.7 ppm and was used to quantify and subtract the drift of the LGR over the comparison period.

In Fig. 3, the ambient data from the LGR has been filtered out to show only each calibration period performed during the month long experiment. The data during each calibration period was averaged (either a total of 10 minutes or one hour depending on the calibration period) and the averages are plotted on Fig. 3. While there is some small variation in the mean mole fraction observed during each calibration from day-to-day, there was an upward trend in the recorded value, by over 1.2 ppm over a 30-day period. This observed drift, while not insignificant, is well within the manufacturer's specifications for this analyzer. However, the observed standard deviation of the two-second points used in each average (the error bars on Fig. 3) remained relatively constant throughout the period with a mean standard deviation of ± 0.3 ppm, which is the manufacturer's specified repeatability for 2-second data. This high-frequency noise is not a problem for the analysis with the K30 sensor because both datasets are averaged to one minute values, which removes most, if not all, of this noise. For comparisons between the K30s in the remainder of this paper, the LGR drift is corrected by first computing a linear fit to the calibration points in time (red line, Fig. 3) and then subtracting from the LGR dataset the difference of this fit line from the tank's assigned value of 463.7 ppm. After this linear correction, the means of each calibration had an RMSE of 0.2 ppm from the fit line.

4 Initial K30 results

Figure 4 shows the original time series of data recorded during the evaluation experiment described in Sect. 2.2. The top panel shows raw CO₂ mole fractions reported by six K30 sensors as well as the LGR analyzer, each of which is located in the same rooftop evaluation chamber. The middle panels show the reported atmospheric pressure and temperature values from one BME280 sensor, and the water vapor mole fraction from the LGR. Then, the bottom panel is the difference between the original recorded K30 value and the corrected LGR recorded CO₂ mole fraction with the calibration periods removed.

30

Over this four-week period, the LGR observed an ambient variation of CO₂ with an average value of just over 423 ppm, and a standard deviation of just under 21 ppm. There is distinct synoptic variation in the

diurnal cycle observed, with the magnitude varying from as little as 10 ppm over 24 hours to more than 100 ppm. Each of the K30s was successfully able to resolve the ambient variations in CO₂ over this evaluation period, although none of the K30s matched the LGR perfectly in both absolute concentration and relative change. However, without any correction or calibration, each K30 was well within the manufacturer's
5 stated uncertainty of ± 30 ppm ± 3 % of the reading for 1-minute values.

From the difference plot (Fig. 4, bottom panel), there are some important things to note. First and foremost, each individual K30 sensor has a distinct zero offset. A few of the sensors are approximately the same as the LGR, but many can have an offset that is as much as 5 % (20 ppm) from the LGR. The differences
10 between each K30 and the LGR all have standard deviations between 4 ppm to 6 ppm and root mean square errors (RMSE) between 5 ppm to 21 ppm. This means that after accounting for the offset of each individual K30, the practical accuracy of the K30 CO₂ sensor can be within 1 % of the observed concentration. Secondly, each K30 difference time series appears to feature two wave patterns, one with a period of around one week, and another with a period of approximately one day. Given that the cycles seem fairly
15 consistent and are present in each K30, this suggests that the difference between the recorded values from the LGR and each K30 is not random, but instead that there are external factors that can be assessed for potential compensation in the K30 response.

5 Environmental correction

In Fig. 4, the difference between the LGR and each K30 is shown in the bottom panel below time series of
20 environmental data from the evaluation chamber. Just like in the difference plot, each of the environmental variables features two distinct time scales of variability. There is a diurnal cycle of each variable, as well as synoptic-scale variability attributed to weather systems that occurs on the order of one week. Because the observed CO₂ differences and the environmental variables are correlated on both short and long time scales, statistical regression methods were used to correct the observed concentration of CO₂ from the K30
25 sensor to a value approximately that of the concentration determined from the calibration-corrected LGR measurements. Generally, a multivariate linear regression is of the form shown in Eq. (1):

$$y = a_1x_1 + a_2x_2 + \dots a_nx_n + \varepsilon_n \quad (1)$$

30 In this case, the measured value y is influenced by: the 'true' CO₂ value (taken as the value from the LGR instrument), pressure, and other environmental variables as the dependent variables x_1, x_2, x_n , respectively. A multivariate regression analysis can then be used to find the corresponding coefficients. In addition, in

order to better identify the contribution from each individual factor, the data were also analyzed in a successive regression analysis, as described below.

5.1 Successive regression method

Each individual K30 sensor's original observed CO₂ dataset is first regressed to the LGR dry CO₂ dataset.

5 This regression accounts for the traditional zero and span corrections made during an instrument calibration. The calibration curve of one K30 for just zero and span is shown in Fig. 5. But to include biases due to environmental factors, then the residual, epsilon (ε), is calculated in Eq. (2) as:

$$\varepsilon = y - ax - b \quad (2)$$

10

where in this instance x , the independent variable, is the LGR dataset and y , the dependent variable, is the K30 dataset.

This process is repeated for each environmental variable pressure (P), temperature (T), and water vapor (q),

15 where (P, T, q) is the independent variable, x , and the ε from the previous step is the dependent variable, y .

This linear regression method leads to eight correction coefficients, of the form a_n and b_n , where n is from 0 to 3 representing each of the independent variables included in the regression. These coefficients can then be used in Eq. (3) along with the environmental variables to correct K30 CO₂ observations for environmental influences.

$$20 \quad y_{corrected} = \frac{y - b_0 - (a_1x_1 + b_1) - (a_2x_2 + b_2) - \dots - (a_nx_n + b_n)}{a_0} \quad (3)$$

For one typical K30, the initial standard deviation of the difference between the K30 and LGR, the RMSE

of the data was 6.9 ppm. Using the cumulative univariate regression method described above for the entire evaluation period, the RMSE decreased after each step. After the span and offset regression, it dropped

25 significantly to 3.3 ppm. Then after correcting for atmospheric pressure, the RMSE dropped even lower to

2.7 ppm. Furthermore, including air temperature and water vapor mixing ratio resulted in a RMSE of 2.7 ppm and 2.1 ppm respectively. It is important to note that the temperature regression did slightly reduce the

RMSE, but not significantly enough to be resolved with only two significant figures. Therefore, using the successive regression method, the RMSE of the observed difference dropped from 6.9 ppm to 2.1 ppm, a

30 reduction of the error by over a factor of three. Fig. 6 shows the results and scatter plots for each step of the

correction for this K30; Fig. 7 shows a difference plot at each step for this same K30 unit. Similar results were observed for each K30 sensor evaluated and a summary can be found in Table 1.

5.2 Multivariate linear regression method

Alternatively, a multivariate linear regression statistical method can be used to calculate the regression coefficients for each K30 sensor. This results in five correction coefficients a_n and b where n represents each independent variable, the dry CO₂ from the LGR, pressure P , temperature T , and water vapor mixing ratio q . Like the successive method above, these coefficients can be used in Eq. (4) along with the original K30 data, y , and the environmental variables to predict the true CO₂ concentration observed.

$$y_{corrected} = \frac{y - b - (a_1x_1) - (a_2x_2) - \dots - (a_nx_n)}{a_0} \quad (4)$$

Using the multivariate regression function provided by Python-SciPy-Stats (Jones et al., 2001), differences from the LGR of the same K30 described in Sect. 5.1 were reduced to an RMSE of 2.1 ppm, slightly better than the iterative method. This consistently better performance from the multivariate method is shown in the other K30 sensors evaluated. Figure 8 shows the final results of the multivariate regression for the same K30 as in Fig. 6 and Fig. 7, as well as the difference between the corrected K30 dataset and the LGR. As with the univariate method, similar results were observed from each K30 sensor evaluated and a summary can also be found in Table 1.

6 Discussion

6.1 Time averaging

There are two observations to note based on the evaluation and analysis. First, both before and after the multivariate regressions, there are frequent shifts in the sign of the difference between each K30 and the LGR; these sudden changes occur at or around sunrise most days. Because of the rapid change in atmospheric CO₂ concentration at this time, the ambient calibration chamber may not be well mixed during this time period. Each K30 is located in a slightly different location in the ambient calibration chamber, and are all approximately 1 to 2 meters away from the LGR inlet. This effect, combined with the different response time of the K30s compared to the LGR, can lead to dramatic differences between what each K30 observes and what the LGR observes at the same timestamp for a short period of time each day.

Atmospheric inversion methods often use hourly averaged data from tower observations (McKain et al., 2012; Bréon et al., 2015; Lauvaux et al., 2016), so after the multivariate regression was applied, the K30 and LGR datasets were further averaged to 10 minute and hourly datasets. The average RMSE for the six K30s with the one-minute data is 2.3 ppm, 2.0 ppm for 10-minute averages, and 1.8 ppm for hourly-

averaged data. Throughout this analysis period, one of the six K30s evaluated performed consistently worse than the others, and after removing it from the averages, the RMSE values dropped to 1.9 ppm, 1.6 ppm, and 1.5 ppm, for 1-minute, 10-minute, and hourly averages, respectively. Thus, by using hourly averages and discarding underperforming sensors, the average RMSE of the difference between the LGR and a K30
5 NDIR sensor can be reduced to approximately 1.5 ppm.

6.2 Regression period

The RMSE described above and in Table 1 are for regressions calculated over the entire experiment period of approximately four weeks. One goal of this work is to develop a methodology to evaluate individual sensors quickly so that they can be used in scientific applications. In Fig. 9 the average RMSE calculated
10 over the entire month of all six K30s is plotted with respect to the number of days used in the multivariate regression from Sect. 5.2. While the RMSE is generally minimized with increasing regression length, after a regression period of just a few days, the RMSE drops significantly from its initial values. Once a few diurnal cycles of varying amplitude have been incorporated, as well as the synoptic scale variations in the atmosphere (with a time scale of around one week), the regression stabilizes. Thus, a regression length of
15 around two weeks is recommended to maximize correction while minimizing the required amount of time the sensor needs to run concurrently with the LGR.

In Fig. 10, a multivariate regression is applied to the same K30 as described in the aforementioned sections and shown in Figs. 6, 7 and 8, but the coefficients are calculated using only data from the first 15 days. The
20 change in the RMSE between the two regressions is 0.1 ppm, going from 1.8 ppm when using all data points to 1.9 ppm when using only approximately the first half. This small, but not insignificant change is most likely attributed to the fact that during the first half of the evaluation period, the ambient CO₂ concentrations do not vary significantly, especially relative to the second half, where both the minimum and maximum values occur. In fact, when instead regressing for the last 15 days of the period, the RMSE is
25 1.8 ppm, a difference not distinguishable with only one decimal place. So as stated above, the diurnal cycles act as a range of calibration points, but values above and below what is included in the regression period may cause the corrected data to still have large errors during these periods, increasing the RMSE for the entire evaluation cycle. Based on these results, it is reasonable to assume that there is either no noticeable baseline drift or that it is assumed to be linear and removed by the multivariate regression in the
30 sensors observed on the weekly to monthly timescales. The longer-term drift of the sensors for periods greater than one month is not known at this time, however, and would require a longer evaluation period of at least six months.

6.3 Generalized Regression Coefficients

All of the final RMSEs calculated in this analysis are from using individual regression coefficients for each K30 sensor. However, it would be beneficial to determine if a generalized set of regression coefficients could be applied to any K30 sensor, and what the RMSE over the evaluation period would be. To calculate the generalized coefficients, the four slopes for each variable as well as the intercepts for each of the five remaining sensors were averaged together, K30-3 was omitted due to the fact that it was the poorest performing sensor, and that its coefficients were significantly different from the other five. After correction using the same set of coefficients, the RMSEs of the six sensors ranged from 3.1 ppm to as high as 23.9 ppm. The final RMSEs in some cases were higher than with the original, uncorrected data. Similar results were observed when the multivariate regression coefficients were calculated using the mean concentration of the five sensors. Thus, it appears that for each K30 sensor, an independent evaluation must be completed to provide observations with a sufficient level of quality.

7. Conclusions and future work

The K30 is a small, low-cost NDIR CO₂ sensor designed for industrial OEM applications. Each of the sensors tested falls within the manufacturer's stated accuracy range of ± 30 ppm ± 3 % of the reading when compared to a high-precision CEAS analyzer, but these ranges are not particularly useful for scientific applications aimed at measuring ambient atmospheric CO₂. If these sensors are individually calibrated, selected for stability, and corrected for sensitivity to temperature, pressure, and RH, the practical error of these sensors is < 5 ppm, or approximately 1 % of the observed value. The final RMSE of the six K30 ranged between 1.7 ppm and 4.3 ppm for 60 s averaging times. Averaging for 200 s further reduces the noise by about 30 %, but longer times did not further improve precision. With errors in this range, these instruments could be used in a variety of scientific applications, including observations at high spatial density to better represent the range and distribution of an urban or natural region's CO₂ concentration.

In the future, further analysis will be performed evaluating the K30 as well as other low-cost CO₂ sensors in a laboratory setting with controlled temperature, pressure and relative humidity. A Picarro cavity ring-down spectroscopic greenhouse gas analyzer will be used as a high-precision control and the various instruments will be subjected to ambient air as well as periodic reference gases. From this lab analysis, we hope to determine the theoretical maximum performance of these sensors in a controlled environment. This subsequent study will additionally attempt to quantify any long-term drift over the course of multiple months.

Acknowledgements. We acknowledge support for this project from the FLAGG-MD grant from NIST's Greenhouse Gas and Climate Science Measurements program (Cooperative Agreement #70NANB14H333). The authors wish to thank the undergraduate and graduate students at the University of Maryland who helped with this analysis. Additionally, we would like to thank all of the members of the NIST Greenhouse Gas and Climate Science Measurements program including Subhomoy Ghosh, Israel Lopez-Coto, Kimberly Mueller, Kuldeep Prasad, James Whetstone, and Tamae Wong for their help.

References

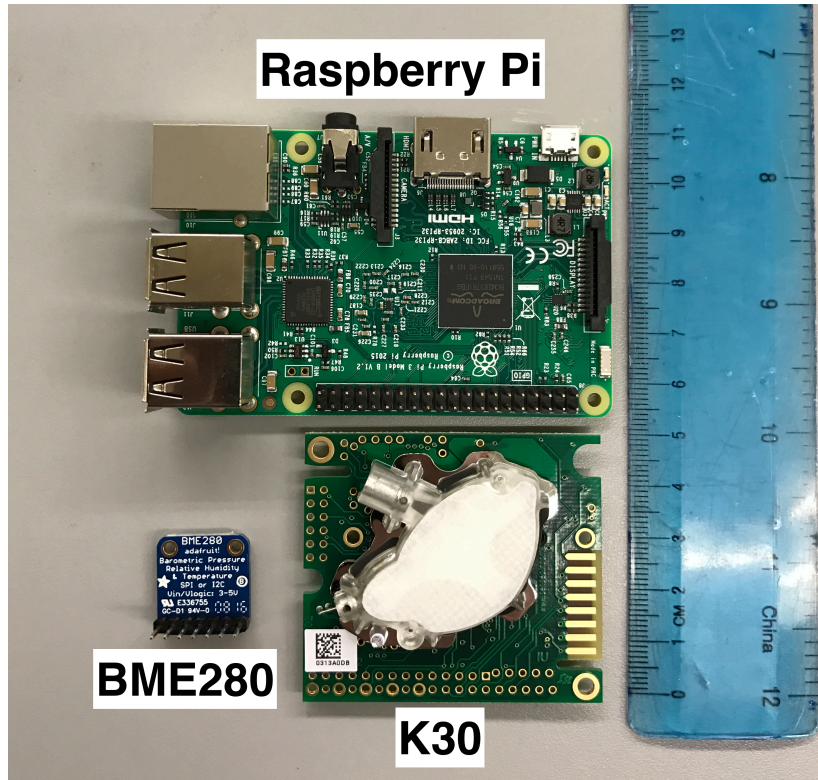
- Allan, D. W.: STATISTICS OF ATOMIC FREQUENCY STANDARDS, Proceedings of the Institute of Electrical and Electronics Engineers, 54, 221-&, 10.1109/proc.1966.4634, 1966.
- 5 Bosch Sensortec: BME280 Digital Pressure Sensor Datasheet, available at: https://cdn-shop.adafruit.com/datasheets/BST-BME280_DS001-10.pdf (last access: 7 June 2016), 2015.
- Bréon, F. M., Broquet, G., Puygrenier, V., Chevallier, F., Xueref-Remy, I., Ramonet, M., Dieudonne, E., Lopez, M., Schmidt, M., Perrussel, O., and Ciais, P.: An attempt at estimating Paris area CO₂ emissions from atmospheric concentration measurements, *Atmospheric Chemistry and Physics*, 15, 1707-1724, doi:(?)
10 10.5194/acp-15-1707-2015, 2015.
- Briber, B., Hutyra, L., Dunn, A., Raciti, S., and Munger, J.: Variations in Atmospheric CO₂ Mixing Ratios across a Boston, MA Urban to Rural Gradient, *Land*, 2, 304, 2013.
- Eugster, W., and Kling, G. W.: Performance of a low-cost methane sensor for ambient concentration measurements in preliminary studies, *Atmospheric Measurement Techniques*, 5, 1925-1934, 10.5194/amt-5-
15 1925-2012, 2012.
- Gas Sensing Solutions: COZIR Ultra Low Power Carbon Dioxide Sensor, available at: http://www.gassensing.co.uk/media/1050/cozir_ambient_datasheet_gss.pdf (last access: 29 December 2015), 2014.
- General Electric: Telaire T6615 Sensor Dual Channel Module, available at: http://www.avnet-abacus.eu/fileadmin/user_upload/Products_Menu/Amphenol/AmphenolAdvancedSensors_CO2_double_channel_module.pdf (last access: 29 December 2015), 2011.
20
- Holstius, D. M., Pillarisetti, A., Smith, K. R., and Seto, E.: Field calibrations of a low-cost aerosol sensor at a regulatory monitoring site in California, *Atmospheric Measurement Techniques*, 7, 1121-1131, doi:(?)
10.5194/amt-7-1121-2014, 2014.
- 25 Hurst, S., Durant, A. J., and Jones, R. L.: A low cost, disposable instrument for vertical profile measurements of atmospheric CO₂, Centre for Atmospheric Science, Department of Chemistry, University of Cambridge, 2011.
- Jones E, Oliphant E, Peterson P, et al.: SciPy: Open Source Scientific Tools for Python, <http://www.scipy.org/> (last access: 10 October 2016), 2001.
- Keeling, C. D., Piper, S. C., Bacastow, R. B., Wahlen, M., Whorf, T. P., Heimann, M., and Meijer, H. A.:
30 Atmospheric CO₂ and ¹³CO₂ exchange with the terrestrial biosphere and oceans from 1978 to 2000: Observations and carbon cycle implications, *History of Atmospheric CO₂ and Its Effects on Plants, Animals, and Ecosystems*, 177, 83-113, 2005.

- Kort, E. A., Angevine, W. M., Duren, R., and Miller, C. E.: Surface observations for monitoring urban fossil fuel CO₂ emissions: Minimum site location requirements for the Los Angeles megacity, *Journal of Geophysical Research-Atmospheres*, 118, 1-8, 10.1002/jgrd.50135, 2013.
- Lopez-Coto, I., Ghosh, S., Prasad, K., and Whetstone, J.: Tower-Based Greenhouse Gas Measurement Network Design – the NIST North East Corridor Testbed, *Advances in Atmospheric Sciences*, 10.1007/s00376-017-6094-6, in press, 2017.
- Los Gatos Research: Fast Greenhouse Gas Analyzer (Enhanced Performance Model) Datasheet, available at: http://www.lgrinc.com/documents/LGR_FGGA_Datasheet.pdf (last access: 29 December 2015), 2013.
- Langridge, J. M., Ball, S. M., Shillings, A. J. L., and Jones, R. L.: A broadband absorption spectrometer using light emitting diodes for ultrasensitive, in situ trace gas detection, *Review of Scientific Instruments*, 79, 14, 10.1063/1.3046282, 2008.
- Lauvaux, T., Miles, N. L., Deng, A., Richardson, S. J., Cambaliza, M. O., Davis, K. J., Gaudet, B., Gurney, K. R., Huang, J., and O'Keefe, D.: High-resolution atmospheric inversion of urban CO₂ emissions during the dormant season of the Indianapolis Flux Experiment (INFLUX), *Journal of Geophysical Research: Atmospheres*, 2016.
- McKain, K., Wofsy, S. C., Nehrkorn, T., Eluszkiewicz, J., Ehleringer, J. R., and Stephens, B. B.: Assessment of ground-based atmospheric observations for verification of greenhouse gas emissions from an urban region, *Proceedings of the National Academy of Sciences of the United States of America*, 109, 8423-8428, 10.1073/pnas.1116645109, 2012.
- Pataki, D. E., Bowling, D. R., and Ehleringer, J. R.: Seasonal cycle of carbon dioxide and its isotopic composition in an urban atmosphere: Anthropogenic and biogenic effects, *Journal of Geophysical Research-Atmospheres*, 108, 8, 10.1029/2003jd003865, 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, *Atmospheric Measurement Techniques*, 7, 3325-3336, 10.5194/amt-7-3325-2014, 2014.
- Raspberry Pi Foundation: Raspberry Pi Hardware Documentation, available at: <https://www.raspberrypi.org/documentation/hardware/raspberrypi/> (last access: 29 December 2015), 2015.
- SenseAir : CO₂ Engine K30 Specification, available at: http://www.senseair.com/wp-content/uploads/2015/03/CO2-Engine-K30_PSP110-R7.pdf (last access: 29 December 2015), 2007
- Shusterman, A. A., Teige, V., Turner, A. J., Newman, C., Kim, J., and Cohen, R. C.: The BERkeley Atmospheric CO₂ Observation Network: initial evaluation, *Atmos. Chem. Phys. Discuss.*, 2016, 1-23, 10.5194/acp-2016-530, 2016.
- Turnbull, J. C., Sweeney, C., Karion, A., Newberger, T., Lehman, S. J., Tans, P. P., Davis, K. J., Lauvaux, T.,

- Miles, N. L., Richardson, S. J., Cambaliza, M. O., Shepson, P. B., Gurney, K., Patarasuk, R., and Razlivanov, I.: Toward quantification and source sector identification of fossil fuel CO₂ emissions from an urban area: Results from the INFLUX experiment, *Journal of Geophysical Research-Atmospheres*, 120, 292-312, 10.1002/2014jd022555, 2015.
- 5 Turner, A. J., Shusterman, A. A., McDonald, B. C., Teige, V., Harley, R. A., and Cohen, R. C.: Network design for quantifying urban CO₂ emissions: assessing trade-offs between precision and network density, *Atmos. Chem. Phys.*, 16, 13465-13475, 10.5194/acp-16-13465-2016, 2016.
- Wang, Y., Li, J. Y., Jing, H., Zhang, Q., Jiang, J. K., and Biswas, P.: Laboratory Evaluation and Calibration of Three Low- Cost Particle Sensors for Particulate Matter Measurement, *Aerosol Science and Technology*, 49, 1063-1077, 10.1080/02786826.2015.1100710, 2015.
- 10 World Meteorological Organization: GAW Report No. 213: 17th WMO/IAEA Meeting on Carbon Dioxide, Other Greenhouse Gases and Related Tracers Measurement Techniques (GGMT-2013), http://www.wmo.int/pages/prog/arep/gaw/documents/Final_GAW_213_web.pdf (last access: July 11, 2016), 2013.
- 15 Wu, L., Broquet, G., Ciais, P., Bellassen, V., Vogel, F., Chevallier, F., Xueref-Remy, I., and Wang, Y.: What would dense atmospheric observation networks bring to the quantification of city CO₂ emissions?, *Atmos. Chem. Phys.*, 16, 7743-7771, 10.5194/acp-16-7743-2016, 2016.
- Yasuda, T., Yonemura, S., and Tani, A.: Comparison of the Characteristics of Small Commercial NDIR CO₂ Sensor Models and Development of a Portable CO₂ Measurement Device, *Sensors*, 12, 3641-3655, 10.3390/s120303641, 2012.
- 20 Young, D. T., Chapman, L., Muller, C. L., Cai, X. M., and Grimmond, C. S. B.: A Low-Cost Wireless Temperature Sensor: Evaluation for Use in Environmental Monitoring Applications, *Journal of Atmospheric and Oceanic Technology*, 31, 938-944, 10.1175/jtech-d-13-00217.1, 2014.

Table 1. Root mean square error in ppm between the CEAS LGR and each K30 NDIR sensor's one-minute averaged data for: the original dataset before correction, at each step of the successive regression correction (correcting for 1. zero/span, 2. atmospheric pressure, 3. temperature, and 4. water vapor mixing ratio), and 5 after the multivariate regression correction. Each value shown is for a regression calculated using data from the entire evaluation period.

	Original	Zero/Span	Pressure	Temp	q (final)	Multivariate
K30 # 1	6.9	3.3	2.7	2.7	2.1	1.8
K30 # 2	5.4	3.5	2.2	2.2	1.9	1.7
K30 # 3	10.9	6.0	5.0	4.9	4.5	4.3
K30 # 4	20.8	3.7	2.5	2.4	1.9	1.7
K30 # 5	8.3	3.7	2.6	2.6	2.2	2.0
K30 # 6	15.2	4.9	3.6	3.5	2.7	2.2



5

Figure 1. Photograph of a Raspberry Pi computer (top), a SenseAir K30 (NDIR) CO₂ sensor (bottom center), a Bosch BME280 temperature and pressure sensor (bottom left), and a ruler for size reference.

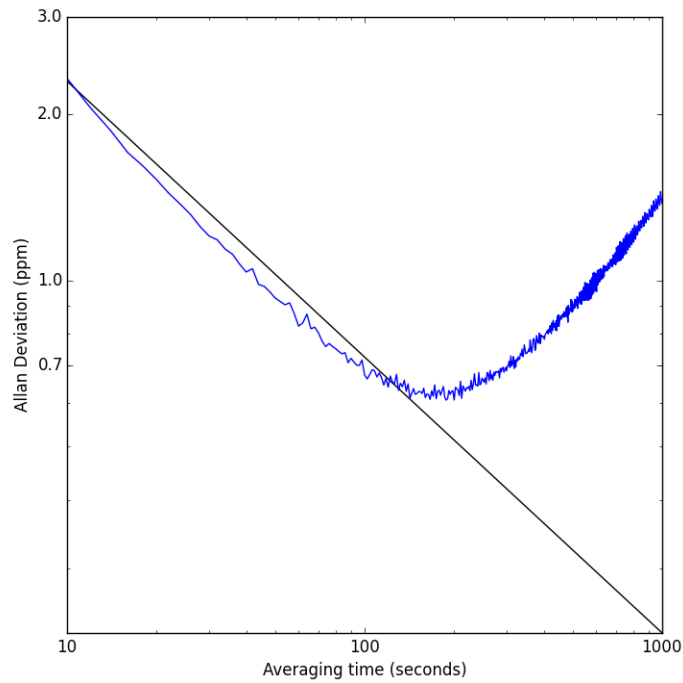


Figure 2. Allan variance analysis for an NDIR (K30) CO₂ sensor when introduced to breathing air from a high-pressure cylinder of a constant and known CO₂ concentration. Averaging times between 10 and 1,000 seconds are shown. The black line (slope -0.5) shows where the noise is white or Gaussian. Averaging times greater than about 200 s produce no improvement.

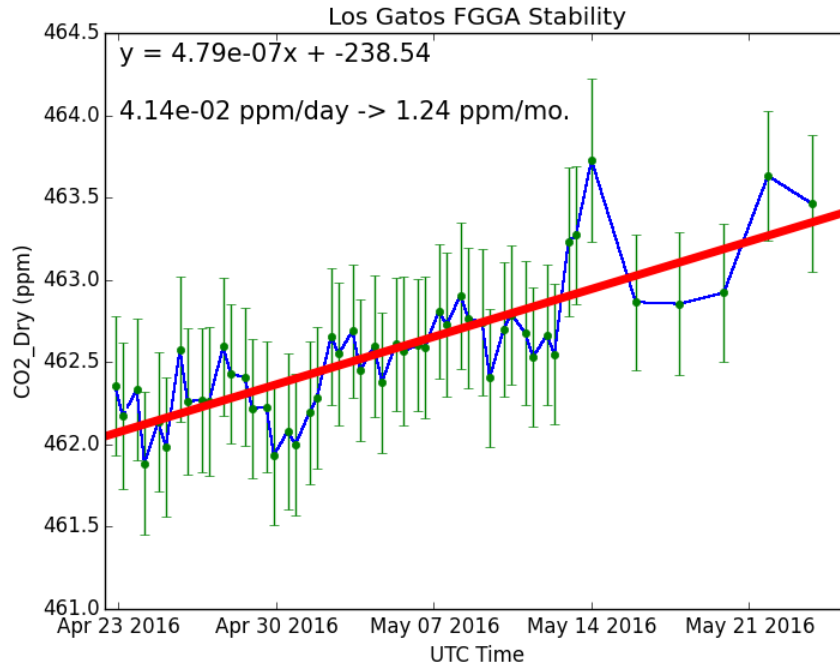


Figure 3. Stability of the Los Gatos Fast Greenhouse Gas Analyzer shown over a 30-day period. Excess breathing air with a fixed CO₂ concentration was introduced periodically using a mass flow controller. The mean of each calibration period is plotted in green with the standard deviation as error bars. The blue line is the linear interpolation between each calibration point, and the red line is a linear fit of each calibration point over the entire time series. The red line is subtracted from the dataset to account for the drift of the analyzer over this period.

K30 vs Los Gatos in Evaluation Chamber

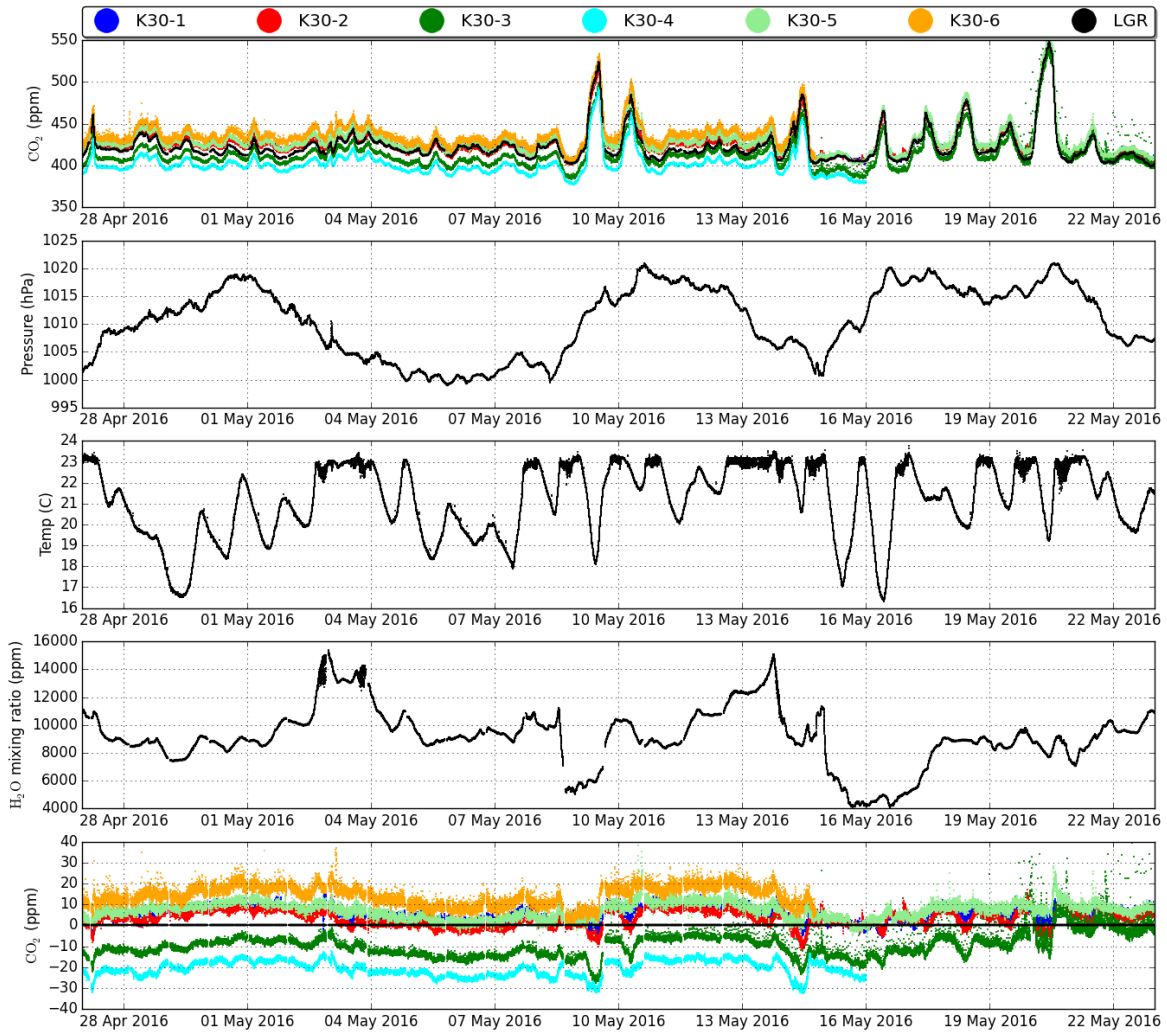


Figure 4. Continuous 1-minute time series data during the evaluation experiment. Top panel: CO₂ observed by six K30 sensors as well as the Los Gatos Research Fast Greenhouse Gas Analyzer. Middle panels: observed atmospheric pressure, temperature, and water vapor mixing ratio, respectively. Bottom panel: difference of each K30 from the Los Gatos instrument.

K30-1 vs Los Gatos Calibration Curve

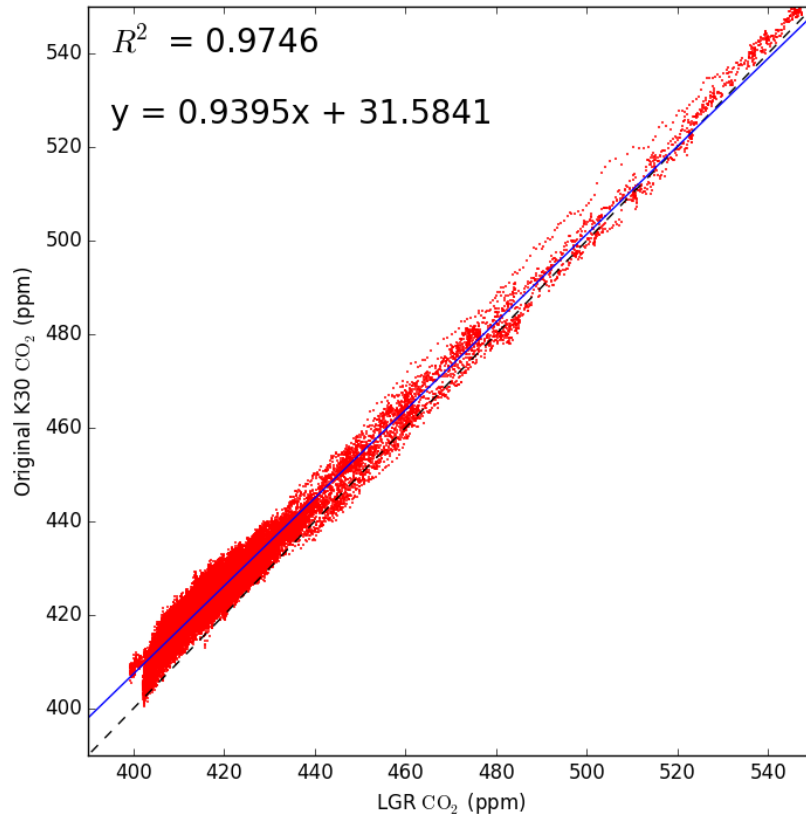


Figure 5. Calibration curve of K30-1 vs LGR for 1-minute averages without any environmental correction, only span and zero offset are corrected. Solid line is the best fit; dashes represent the 1:1 line

5

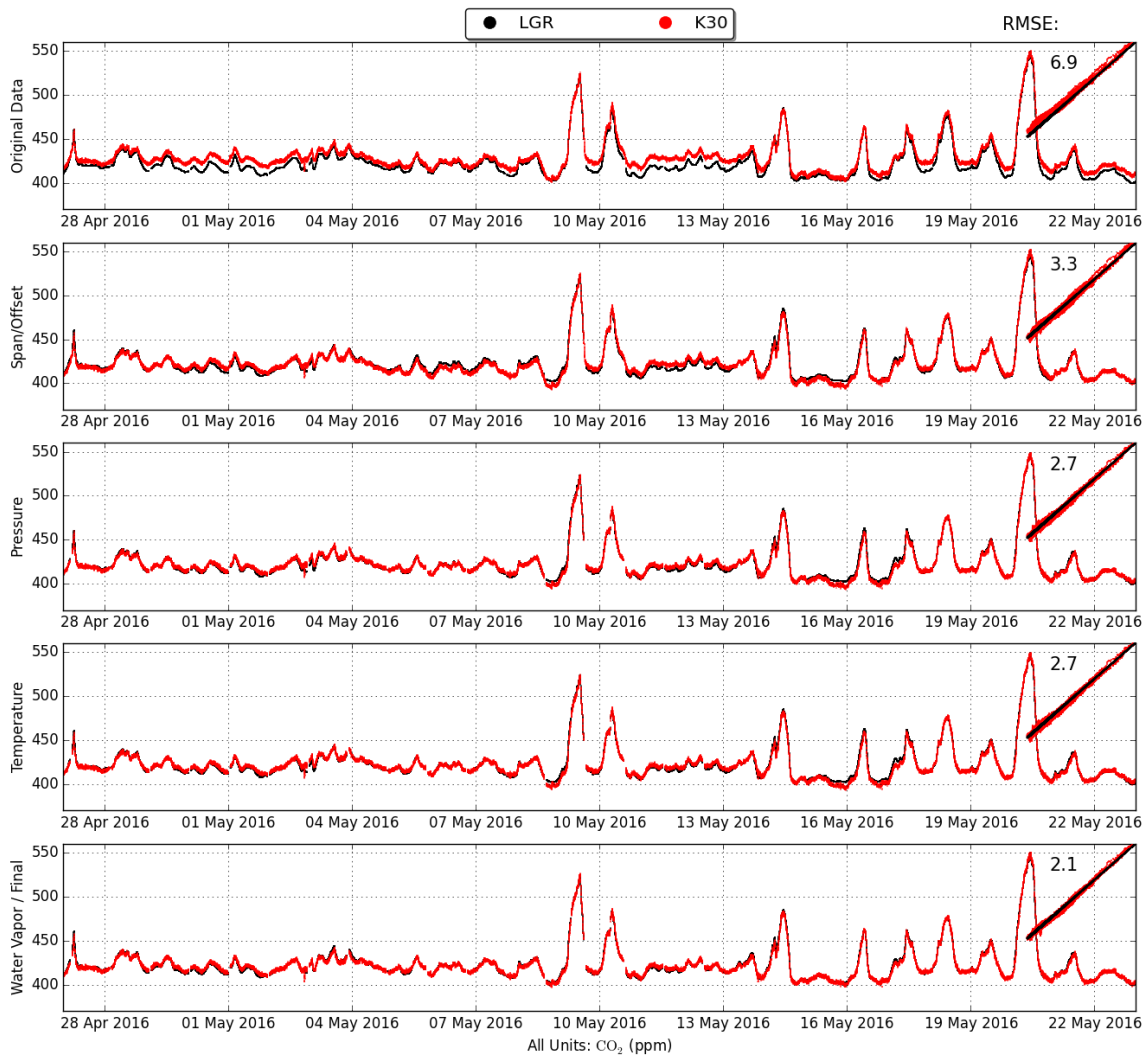


Figure 6. A continuous time series of 1-minute averages as well as scatter plots for K30 #1 compared to the LGR instrument during each step of the successive regression described in Sect. 5.1. Cumulative, in order from top to bottom: the original dataset, after correcting for span and offset, after correcting for pressure, after correcting for temperature, and finally, after correcting for water vapor. The root mean square error (RMSE) of the K30 data compared to the LGR at each step is annotated to the upper left of the scatter plot. This regression contains all data points observed in the evaluation period.

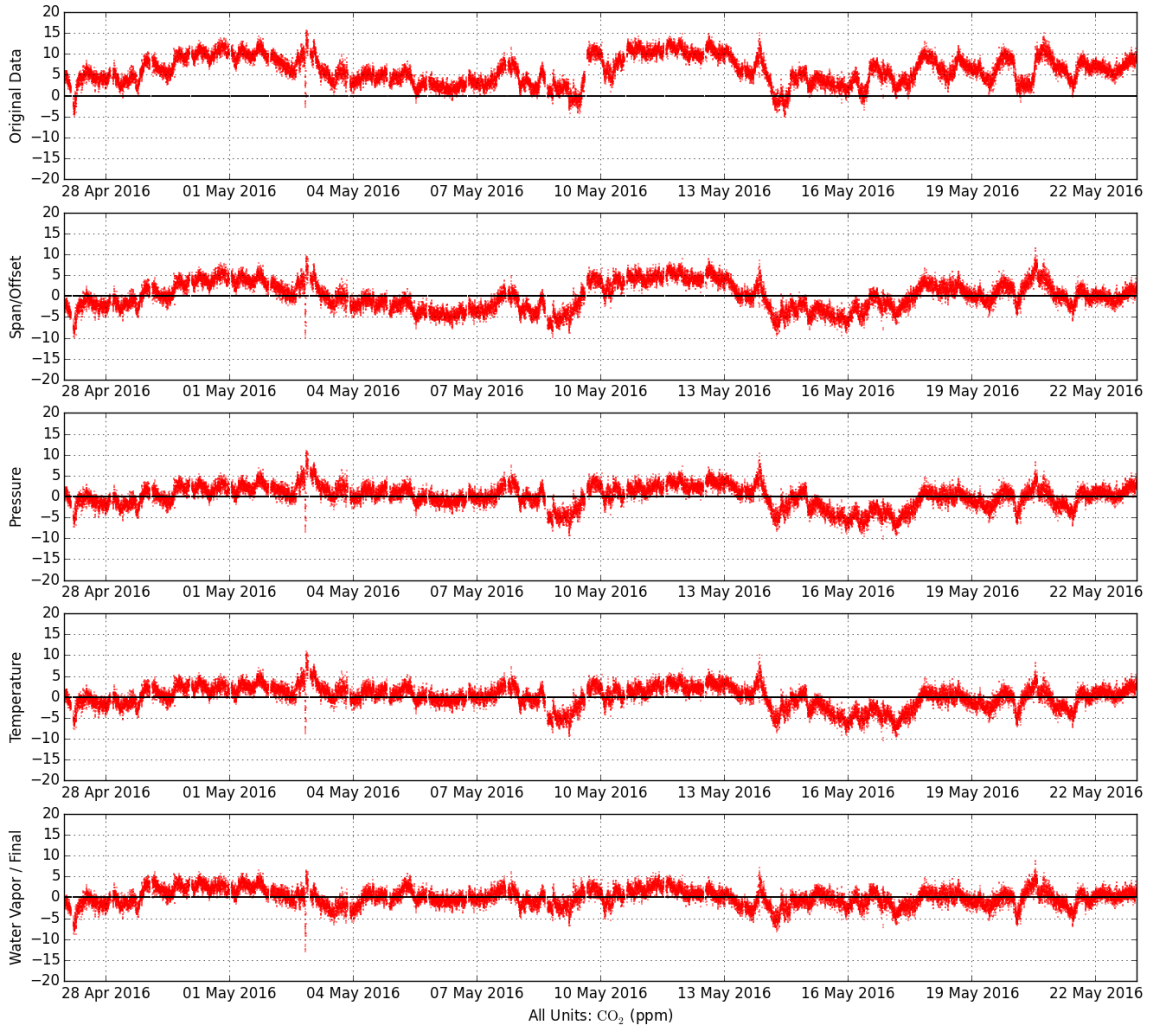


Figure 7. Difference plots for K30 #1 compared to the LGR during each step of the successive regression described in Sect. 5.1 and shown in Fig. 6 for 1-minute averages. Cumulative, in order from top to bottom: the original dataset, after correcting for span and offset, after correcting for pressure, after correcting for temperature, and finally, after correcting for water vapor.

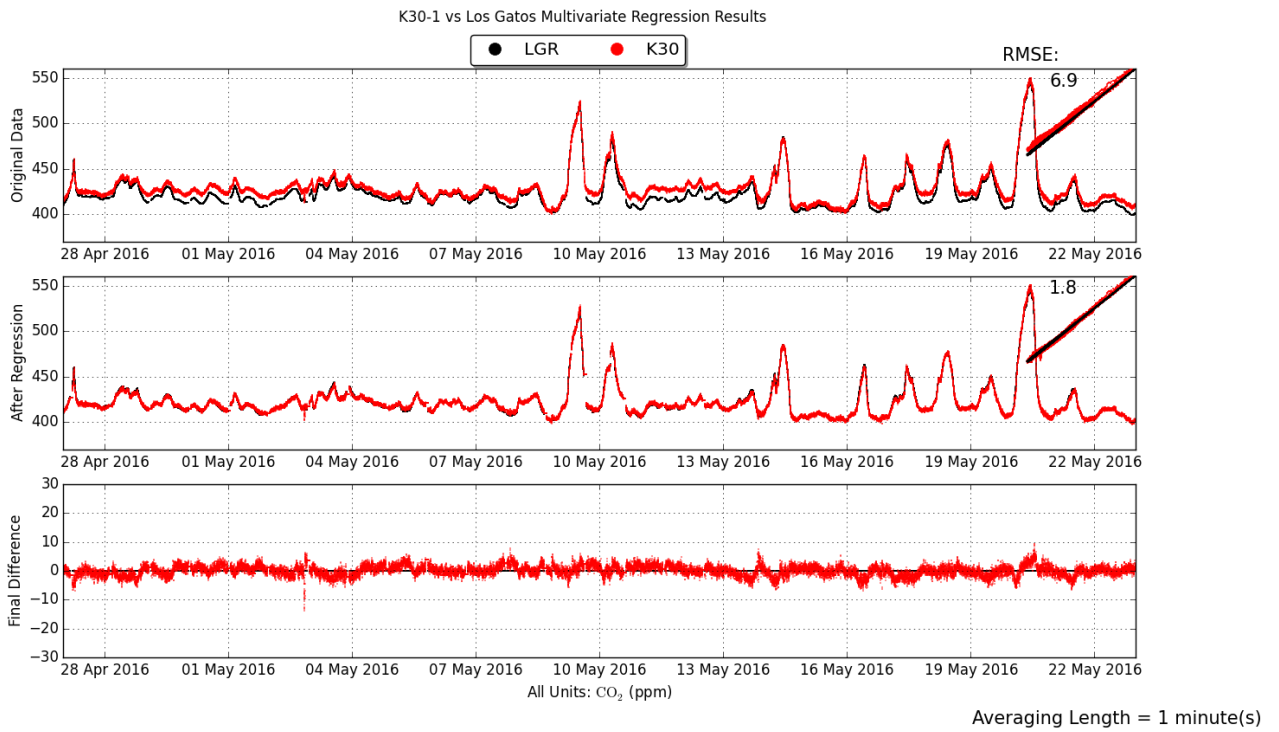


Figure 8. A continuous time series of 1-minute averages as well as scatter plots for K30 #1 compared to the LGR for the multivariate regression described in Sect. 5.2. Top panel: the original data, middle panel: final time series after correction, and the bottom panel: difference plot between the corrected K30 dataset and the original LGR dataset. The root mean square error (RMSE) of the K30 data compared to the LGR before and after the regression is annotated to the upper left of the scatter plot.

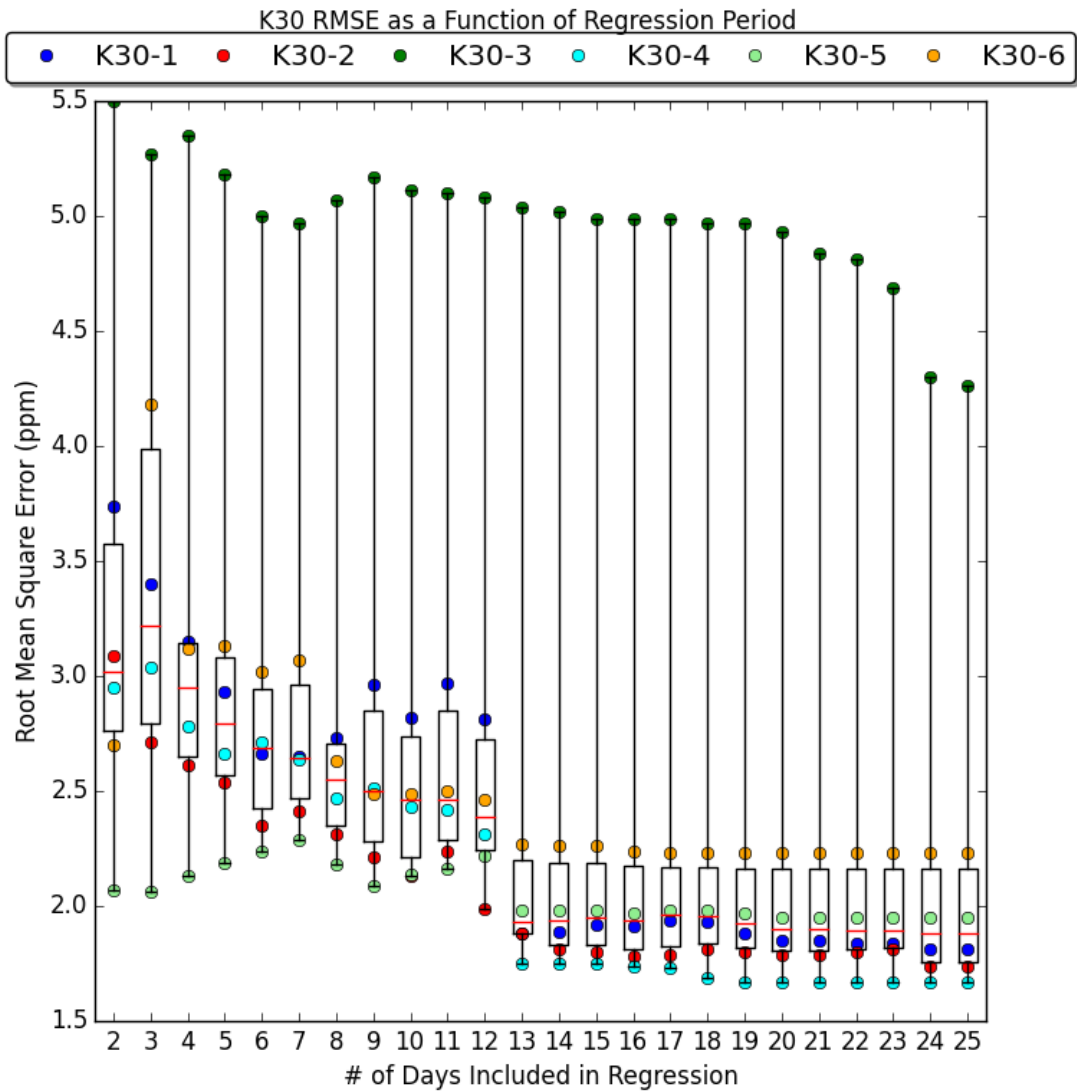


Figure 9. The RMSE of all six K30 NDIR sensors when compared to the LGR over the entire experiment as a function of how many days the regression analysis was performed. The colored dots represent each K30's RMSE, and the box plot shows the median in red, the first and third quartiles within the box, and the min and max values on the whiskers.

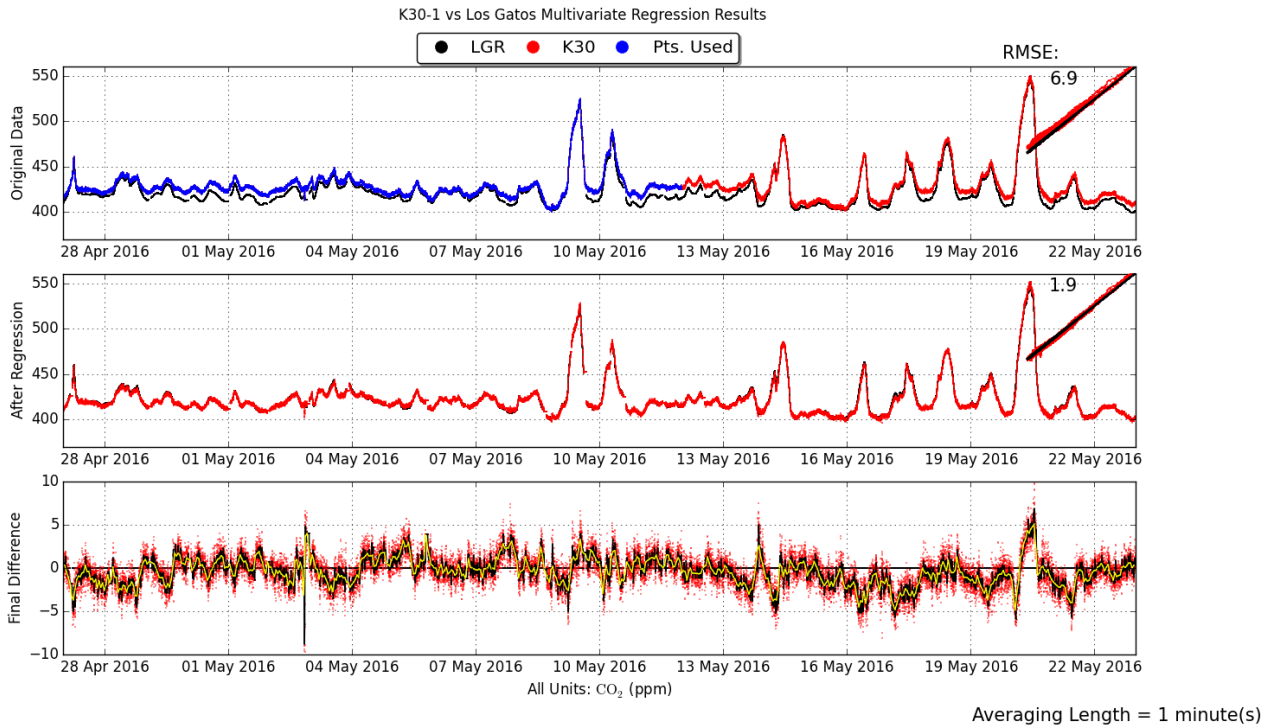


Figure 10. As depicted in Fig. 8, a continuous time series as well as scatter plots for K30 #1 compared to the LGR for the multivariate regression described in Sect. 5.2. Top panel: the original data, middle panel: final time series after correction, and the bottom panel: difference plot between the corrected K30 dataset and the original LGR dataset. However, this regression only includes the first 15 days of data (regression training data in blue, the entire dataset in red) to compute the correction coefficients. The difference plot (bottom) also shows running means for 10 minute (black) and hourly (yellow) averages.