

Reviewer #2

General comments. The manuscript describes an open source, systematic methodology to calibrate low-cost sensors (LCSs). The Authors propose a 7-step statistical method based on: 1) preliminary analysis of raw data; 2) data cleaning; 3) flag data; 4) selection of the model by using both multiple linear regression and random forest and several statistical parameters; 5) model validation; 6) export of the experimental data as concentrations; 7) error predictions. Finally, the Authors tested the proposed model with an example during a field campaign in urban environment.

The manuscript shows a very interesting and systematic methodology to calibrate LCSs, suggesting to employ a univocal and standardized method to let comparable the LCSs measurements, considering the more and more frequently use of this technology. Despite this, the manuscript requires revisions before to be accepted for final publication. Following suggestions and specific comments.

We thank the reviewer for their positive overall evaluation of this manuscript. Their comments have improved the quality of this manuscript. What follows are point-by-point responses to their comments regarding recommended changes to the manuscript.

Specific comments.

For this calibration procedure, reference instruments are needed. Trends due to specific events (i.e., burning etc...) could be not properly described by the sensors, if not calibrated in the same conditions?

As the reviewer points out, this is a general limitation of low-cost sensors, as they do not have the same level of accuracy as reference instruments and are susceptible to various cross-sensitivities. It is likely that the low-cost sensors would respond to specific events (i.e., burning), however if they were not calibrated with reference instruments under conditions in which such events occur, then it would be very challenging to ensure that the applied calibration accurately represented pollutant concentrations under substantially different ambient conditions. Where available, external data from monitoring stations could be used to identify such events and calibrations could be adjusted post-hoc, if need be.

Text has been added to section 3.3, which reads, "Further reference data on other species would have been beneficial to this calibration, as the MOS do exhibit cross-sensitivities to other species, but resources were insufficient, and these data were not collected."

Moreover, did the Authors try to do a calibration procedure by using chemical standard to produce a calibration curve at different concentrations and conditions in laboratory experiment? If yes, could the Authors discuss difference between the two approaches?

We thank the reviewer for this comment. A laboratory calibration was not conducted in this study as we did not have the technical resources to do so. Furthermore, calibration curves for various species are provided in the technical specifications of the individual sensors and served as a basis for understanding their sensitivity. There are many other studies that have produced calibration curves or tested LCS in both laboratory and field conditions, but as this methodology focuses on the field calibration of LCS, we did not include any discussion of laboratory calibrations. Furthermore, as several studies have shown, laboratory calibrations are often not transferrable to calibration of ambient measurements, see e.g., Castell et al., (2017), Rai et al., (2017).

In case of LCSs time drift, did the proposed methodologies take into account of it (or allow to)? Is the proposed frequency (2 weeks every 2-3 months) enough to take into account of seasonal variations and, eventually, time drift of the sensors?

We thank the reviewer for this comment. Sensor drift is indeed a significant challenge for long-term operation of LCS. In our experience, the proposed frequency of co-location is sufficient to account for sensor drift, which generally occurs in significant amounts 4 months after initial calibration (Peterson et al., 2017). As such, the frequency of co-location in this study would account for any significant drift. Between the summer and winter campaigns, drift might have occurred on a significant scale, but co-location data were not used from one campaign to train models in the other, therefore this should not be a problem. This is, however, an important topic that will be explored in future studies with experiments conducted over longer timescales.

Text has been added to Section 3.1, which reads, "These MOS sensors typically experience significant amounts of drift four months after initial calibration, which is why in this study co-locations were conducted at high frequency, before and after each experiment."

Lines 430 and following. LCSs could have a T and RH dependency. Is it appropriate to apply the suggested corrections by the manufacturer for T and RH to the raw data before to apply the models in the proposed methodology? Could the Authors discuss this aspect and if the models relationships with RH and T are in line with what, eventually, suggested by the manufacturer?

We thank the reviewer for this comment. For these sensors (SGX Sensortech MICS-4514), no direct corrections for T or RH are recommended by the manufacturer. Instead, they provide a recommended range of operation for these meteorological parameters and discuss the conditions under which sensor sensitivity tests were conducted. Other studies (i.e., Peterson et al. (2017)) have identified that MOS sensor performance can be affected by T and RH and therefore we include these parameters in the model selection process. In this study, as can be seen in Tables 2-5, the most accurate models for predicting NO₂ and O₃ included both T and RH for the MLR models, and either T or RH for the RF models.

Other species could interfere with measurements: if the concentration of those compounds changes (season, night-day, etc...), this could affect the sensors response. How the Author suggest to deal with this eventuality?

We thank the reviewer for this comment. Such cross-sensitivities do exist and interfere with the performance of low-cost sensors. The methodology presented here accounts for these types of changes by recommending relatively frequent co-locations to capture changes in conditions owing to seasonality, etc. and subsequent effects on ambient conditions, including pollutant changes. We recognize however, that we cannot attribute changes in the model to changes in specific compounds that are interfering. To properly account for all of these, the low-cost sensors would need to be rigorously calibrated alongside reference instruments that measure a much wider range of the species to which the sensors are potentially cross-sensitive. This was not possible for this study and would not be for most studies, and therefore the attribution of impact of these species is unfortunately unknown. It is likely that a portion of the uncertainty in these measurements can be attributed to the impact of these cross-sensitivities on sensor performance.

To account for changes in chemistry due to daytime and nighttime conditions, we included a binary "Time of Day" variable (see section 3.4.4) in the model selection process, though in most cases it was deemed to be unimportant to prediction of NO₂ and O₃. A sentence has been amended in

section 2.1, which now reads “Regular co-location allows for the establishment of datasets that cover not only changes in meteorology, but also sensor functioning and interactions of potentially cross-sensitive species.”

The Authors showed their experiment results only for NO₂ and O₃ MOS sensors. Is this methodology applicable to other compounds (i.e., VOCs) or technologies (i.e., PID, electrochemical) with the same characteristics proposed in the manuscript? This information should be included in the manuscript.

We thank the reviewer for their comment. Reviewer #1 also pointed this out and this has been answered as follows:

A brief physical explanation has been added in the text. While different in their design, both MOS and EC sensors produce a measure of voltage which varies in response to changing concentrations of gas-phase species, and hence can be calibrated using the same methodology. In more recent published work, we have successfully applied this methodology for the calibration of EC sensors, see Schmitz et al., (2021).

The added text reads “Furthermore, while it was applied here to sensor systems containing metal oxide LCS, this methodology is also equally as applicable to electrochemical LCS or photoionization detectors (PID), as these produce a similar measure of voltage that varies in response to changing concentrations of gas-phase species and have similar cross-sensitivities to temperature and relative humidity. It is not directly applicable for optical particle counters (OPC) for the measurement of particulate matter, as the transformation of the raw data into concentrations during calibration functions differently, though some of the principles discussed here are still relevant. For an application of this methodology to EC sensors, please see Schmitz et al., (2021).”

The Authors refer to Ammonia and Reducing gas sensor in the manuscript (see Table 1), but results regarding these sensors are not present. Is this due to lacking of reference instrument?

Yes, this is correct. No reference instrumentation for NH₃ was available for this study and the focus was on NO₂ and O₃.

Could the Author describe in more details how the Step 4 is performed? How the experiment and co-location data have been used in this step? The Authors describe that in Step 5 and 6 the co-locations data were used, but information about data used in Step 4 seem missed.

We thank the reviewer for this comment. The same co-location data that is used in model validation in step 5 and in training the final models in step 6 is used in the model selection process of step 4. In this step, the co-location data are broken into multiple continuous training and test data sets upon which various statistical models using combinations of the various LCS and meteorological variables are tested. Statistical parameters of R², RMSE, MAE, and AIC are used to select models of best fit. Once the best-fitting models are determined, these are then passed on to step 5, the model validation step.

Text has been added to section 2.2.4 to clarify that the data used in the model selection process are those from the co-locations. This now reads, “The co-location data were used in this step to train various models and determine the best fitting MLR and RF models.”

Did the Authors intercompare between them similar sensors, i.e. two Zephyrs, before and after the calibration to check the response of same sensors in same conditions?

We thank the reviewer for this comment. This same issue was raised by reviewer #1. Our response is as follows:

For this study, we use the two low-cost sensors primarily as examples of how to use the seven-step methodology and did not consider their intercomparison as we felt it might distract from the main focus of the work. However, we have added graphs depicting the agreement between standardized raw LCS data of s71 and s72 during the 2 co-locations of the winter measurement campaign into the supplemental information (Figures S4 and S5). The Oxa and O3a MOS sensors of each sensor system are linearly related, but due to differences in sensor sensitivity, have different baselines. In the summer campaign, the relationship between the O3a sensors of s71 and s72 during co-location 2 is non-linear but returns to linear agreement in co-location 3 and in the winter campaign.

A reference to the added graphs in SI and a brief discussion of this point in the text was added [section 3.3] and reads, "To compare sensor performance between s71 and s72, an intercomparison of available co-location raw data was conducted for the oxidizing MOS (Oxa) and ozone MOS (O3a). During all co-locations, the sensors had a linear relationship and an $R^2 > 0.95$ (Figures S4 and S5). In only one instance was this not the case (co-location 2, O3a), where the R^2 was 0.59 and a deviation from linearity was detected. This relationship was linear and normal in all other co-locations."

About the data cleaning, how the Authors correct data for possible bias effect? Line 190: the duration of the moving window chosen to remove the outliers avoid to exclude from the dataset some specific and real events with short duration?

We thank the reviewer for this comment. A similar comment was made by reviewer #1. Our response is as follows:

Indeed, this is a significant challenge in data cleaning and unfortunately requires, in some cases, subjective assessment for an accurate determination to be reached. While it is possible that this data cleaning method removes some non-outlier measurements during peak emissions, it is equally possible that such events are indeed outliers due to technical sensor error. For this reason, we recommend that identified outliers be graphically compared with neighboring points to determine if their removal is justified.

With the optimized moving window and threshold identified in this study, for sensor system s71, a total of 58 outliers were detected from >500,000 data points. In this case individual assessment of each point's 'outlierness' was practical, but there may be cases where this is impractical. In such cases, we recommend a random subset of outliers be graphically assessed to determine the extent to which the data cleaning function is removing actual outliers. Furthermore, while peak emissions may be misidentified as outliers, these short peaks often occur not as individual points, but as small groups that indicate a peak event. As such, the uppermost point might be removed by this data cleaning method, but those surrounding points would not, ensuring that the peak emissions event remains mostly accounted for. This is imperfect, but it is unlikely that a data cleaning method exists which can perfectly separate outliers from peak events. It is with such peak events that other tested methods such as the AutoRegressive Integrated Moving Average performed particularly poorly, identifying most peak events as outliers. If such events are expected due to the deployment environment, particular care in evaluating the outlier removals should address this.

Clarification text has been added to section 2.2.2 regarding this and now reads, "The points flagged as outliers with this method were then graphically assessed against neighboring datapoints to prevent inadvertent removal of peak emission events. In other cases where assessing all outliers is impractical, it is recommended to do so with a random subset of outliers. Furthermore, if substantial short-term events are expected due to the deployment environment, such as during mobile measurements, a more thorough check of the outlier removal should be done."

Lines 218-220. To identify which model better describe the measurements in term of over or under estimation, could the Authors consider to include also a statistical parameter such as the Fractional Bias?

We thank the reviewer for this comment. We are unfamiliar with the use of that statistic and while we could find its definition, we could find no information on how to assess the uncertainty on the central estimate due to sample size. Unless the reviewer can point us to a reference that discusses assessment of uncertainty on fractional bias for a given sample size, we have no way of knowing if a particular value derived from our sample actually represents significant bias. At this stage, we ask the reviewer for greater clarification on how they envision the use of fractional bias in this study.

Lines 365-369. Co-location 3 was at the end of the summer campaign (i.e., October). Anyway data for experiment 2 are not available. It seems from Figure 4 there is a seasonal impact. Did the Author use this co-location for their calibration for Experiment 1? Did the different season affect the calibration procedure? Are the 2 weeks every 2-3 months enough to take into account of it?

We thank the reviewer for this comment. Indeed, there is a seasonal impact on the data, as the co-location in October experienced cooler temperatures than co-location 2, which took place at the end of the summer. The reason co-location 3 was included in the analysis was to see if the data were still useful for training a model that could predict on experiment 1, which took place in August. As can be seen in tables 8 and 9, the most accurate models for prediction on experiment 1 were those that were trained with co-location 2 exclusively. However, models trained using both co-locations were only slightly more inaccurate. We would have expected co-location 3 to be more relevant for experiment 2, due to the seasonal effects that the reviewer points out. Unfortunately, these hypotheses could not be tested due to loss of data for experiment 2 and co-location 1.

Text has been added to section 3.5, in connection with a separate but similar comment. The text reads, "This alludes largely to the fact that seasonal changes caused co-location 3 to experience different meteorological and pollution conditions than were present during the experiment. While results show that this co-location was not useful for accurate prediction, it is likely that it would have been more relevant for prediction on experiment 2, during which the environmental conditions were more comparable. Similarly, co-location 1 would likely have been more valuable for prediction with experiment 1 than with experiment 2. However, due to the loss of data from s71, this could not be assessed more closely in this study."

Lines 390-396. Is the GSM the only way to transfer data to a database? The warm up time was provided by the manufacturer?

There are other methods of data transfer, such as through Wi-Fi or a direct data download via cable, but these were not incorporated into the prototype Zephyrs and therefore GSM was the best method for sending data to the external database. The warm-up time and the impact on sensor

signal was detected during post-processing and was not provided by the manufacturer but confirmed in subsequent discussions with EarthSense. It was known that the sensors needed time to warm up before stabilizing, but it was not clear that the sensor signal would fluctuate as significantly as it did when the sensors briefly turned off. This further highlights the importance of Step 2 in which the raw data are rigorously cleaned.

Lines 421-423. Since the 3rd co-location is in October, could this be indication that closer and more frequent co-location are needed? See also the following Section 3.5 (line 539) and Figures 14-15.

We thank the reviewer for this comment. If data from co-location 1 and experiment 2 were available, this would have been assessed more closely in this study. Tables 8 and 9 reveal that inclusion of co-location 3 training data in final models for prediction did not improve predictive accuracy on experiment 1. This reflects in large part the impact of seasonality as the reviewer points out. However, using training data from co-location 2 alone was suitable for prediction on experiment 2 and therefore we do not feel that more frequent co-location was needed in this case. While co-location 3 was not necessary for accurate prediction of experiment 1, it would likely have been necessary for accurate prediction of experiment 2. Therefore, we feel that the spacing of co-locations was appropriate for this study, but the technical failures that led to loss of data for co-location 1 and experiment 2 prevented proper evaluation.

Text has been added to section 3.5, in connection with a separate but similar comment. The text reads, "This alludes largely to the impact of seasonal changes on co-location 3, which experienced different meteorological and pollution conditions than were present during the experiment. While results show that this co-location was not useful for accurate prediction, it is likely that it would have been more relevant for prediction on experiment 2, during which the environmental conditions were more comparable. Similarly, co-location 1 would likely have been more valuable for prediction with experiment 1 than with experiment 2. However, due to the loss of data from s71, this could not be assessed more closely in this study."

Figures 14-15. Could the Authors add the 1:1 lines and indicate the R² in the plots? How the Authors can explain the constant thresholds in the plots of panels 15e and 15f? Looking at Figure 11, the models using internal T and RH seem to give lower O₃ and NO₂ compare to the ones that use the ambient T and RH. In figure 11 this is less evident: could the Authors explain it and the reasons/meaning of the slopes (typically lower than the unit) and intercepts?

We thank the reviewer for this comment. 1:1 lines as well as R² and RMSE have been added to the plots, which was also requested by reviewer #1.

The constant thresholds seen in panels 15e and 15f reflect the nature of the RF models, which cannot predict outside the minimum and maximum values of the co-location data they are trained with. This is a fundamental flaw of RF models and is referred to in the final paragraph of section 3.5. More text has been added here to explain this more clearly.

The text reads, "This is a fundamental flaw of RF models as they cannot predict outside the bounds of the co-location data they are trained with."

Regarding the interpretation of slopes and intercepts in figures 14 and 15, a slope greater than 1 indicates overprediction, whereas a slope less than 1 indicates underprediction. The intercepts typically indicate the same (just centered around 0), but in this case, as the models tend to overpredict low concentrations and underpredict high concentrations for NO₂, the intercepts are

in some cases offset to be positive. For O₃ this is less often the case, as the slopes and intercepts more often both agree regarding under or overprediction.

To be certain that the models trained with ambient and internal T/RH were giving significantly different concentrations, their predictions were compared to each other using student's t-tests and mann-whitney-wilcoxon U-tests. In all cases the results were statistically significant, confirming the reviewer's suspicions. The likeliest explanation stems from the differences in raw output of the internal and ambient T/RH monitors. These have somewhat non-linear relationships, which in the case of RH becomes especially non-linear under high RH conditions. In fact, the internal T is consistently higher and RH consistently lower than their ambient counterparts. These differences likely lead to the significantly different predictions between ambient and internal MLR and RF models, but the magnitude is not large, as T/RH are less relevant to prediction in all cases than the MOS O_{3a}/O_{3a}. As follows with previous comments regarding the differences between models trained with internal and ambient T/RH, while the results are interesting and warrant further investigation, we feel that this analysis lies outside the scope of this manuscript and should be the subject of future research.

Text that was added to the discussion in line with a previous comment from Reviewer #1 has been modified further, and now reads, "However, given that models using ambient data were more accurate during the validation step and significant differences between predictions of models trained with internal vs ambient T and RH were identified, these results require closer inspection, which should be the subject of future research."

Supplementary. Why for the winter campaign, the Authors use co-location 1 and 2 instead of 4 and 5, which are closer to Experiment 3? Comparing Table 9 and Table S4, the models identified for O₃ are different (and similarly for NO₂): how the Authors could explain this?

We thank the reviewer for this comment and noticing this error. At an earlier stage in the analysis, co-locations 4 and 5 were referred to as 'co-locations 1 and 2' for the winter campaign. The enumeration of the co-locations was not changed in the supplemental information upon submission, but this error has now been corrected.

Technical comments.

Line 88. See "host".

"...a host of..." has been changed to "many".

Line 200. Do the Authors refer to Section 3.5?

Yes, this has been changed.

Line 291. Decent and good agreement should have to be quantitative and not qualitative information.

These qualitative statements have been changed and replaced with the actual R² of the intercomparison between instruments.

Lines 306-307 and 317-318. Information about the date of the campaigns are confused and should be coherent. The information could be furnished only once clearly and I would suggest to add the dates in Figure 2, as well.

The dates for the experiments have been moved to the next paragraph and have been placed alongside the dates for the co-locations. Figure 2 has been updated to include the dates.

Line 328-329. This sentence should be clarified. Zephyr s71 and s72 were located as in Figure 3 or with reference in an office on the 6th floor? In the former case, this information is redundant and could be included in previous paragraphs, when describing the setting (line 311 and following). In the second: how air masses have been sampled?

This sentence has been moved to the previous paragraph when describing the co-location set-up on the 6th floor of the Mathematics building. Text has been modified in the original paragraph of this sentence for coherence.

Line 359. When the Authors refer to “combined” co-location, this means an average of co-location 1 and 2?

Combined co-location refers to the use of both co-locations together, not the average. The text has been updated to be clearer.

Lines 419-423. This section is not well described. Could the Authors explain in more details the criteria to be used to flag the data?

We thank the reviewer for this comment. The details regarding the flagging of data are provided previously in section 2.2.3 and therefore were not repeated here. If they feel that the details in section 2.2.3 are insufficient, then we will gladly be more specific.

Lines 430. Could the Authors specify in this or previous paragraph the units of the input data?

Done. Units have been added to graphs as well where they were missing.

Lines 436-438. The Author report that relationship between O_x and O₃ was determined be inverse; but, since the predictive accuracy for no transformation is similar, they selected the latter. Anyway, in Table 3 there is not inverse relationship and a log dependency between O₃ and O_x was selected (also in Table 5 there is not inverse transformation). Could the Author explain this discrepancy or illustrate better this paragraph?

We thank the reviewer for identifying this discrepancy. The text was not updated alongside new findings during the model selection process, leading to this discrepancy between the text and Table 3. The findings in Table 3 are correct and the text in section 3.4.4 has now been updated accordingly.

Lines 490-498. A comparison with the reference O₃ and NO₂ data should be included here (and in Figure 11).

We thank the reviewer for this comment. An earlier version of this manuscript contained comparisons to NO₂ and O₃ concentrations in this section but was removed in the submitted version. The justification for this decision was to ensure that in section 3.4, the results were

described without reference concentrations as under normal experimental deployment, a comparison of calibrated LCS data to reference instruments would not be possible. Therefore, to keep the example of the use of the methodology more relevant for realistic scenarios, we removed the comparisons to reference concentrations in step 6. The comparison between calibrated data and reference concentrations was then placed exclusively in section 3.5, which describes the atypical extra validation step taken in this study. In this regard, we prefer to not include reference concentrations in step 6, as it is not realistic to expect that the experimental deployment of LCS will occur alongside reference instrumentation in the majority of future studies and applications. If the reviewer feels that the comparisons to reference data made in Tables 8 and 9, as well as in Figures 14 and 15, are insufficient, then we will reconsider including reference data during Step 6.

Line 595. See "this is should be".

This grammatical error has been fixed.

References

Castell, N., et al. (2017). "Can commercial low-cost sensor platforms contribute to air quality monitoring and exposure estimates?" *Environ Int* 99: 293-302. <https://doi.org/10.1016/j.envint.2016.12.007>.

Peterson, P. J. D., et al. (2017). "Practical Use of Metal Oxide Semiconductor Gas Sensors for Measuring Nitrogen Dioxide and Ozone in Urban Environments." *Sensors (Basel)* 17(7). <https://doi.org/10.3390/s17071653>.

Rai, A. C., et al. (2017). "End-user perspective of low-cost sensors for outdoor air pollution monitoring." *Sci Total Environ* 607-608: 691-705. <https://doi.org/10.1016/j.scitotenv.2017.06.266>.

Schmitz, S., et al. (2021). "Do new bike lanes impact air pollution exposure for cyclists?—a case study from Berlin." *Environmental Research Letters* 16(8). <https://doi.org/10.1088/1748-9326/ac1379>.