Low-Cost Outdoor Air Quality Monitoring and Sensor Calibration: A Survey and Critical Analysis

FRANCESCO CONCAS, JULIEN MINERAUD, EEMIL LAGERSPETZ, SAMU VARJONEN, XIAOLI LIU, KAI PUOLAMÄKI, PETTERI NURMI, and SASU TARKOMA, University of Helsinki, Finland

The significance of air pollution and the problems associated with it are fueling deployments of air quality monitoring stations worldwide. The most common approach for air quality monitoring is to rely on environmental monitoring stations, which unfortunately are very expensive both to acquire and to maintain. Hence, environmental monitoring stations are typically sparsely deployed, resulting in limited spatial resolution for measurements. Recently, low-cost air quality sensors have emerged as an alternative that can improve the granularity of monitoring. The use of low-cost air quality sensors, however, presents several challenges: They suffer from cross-sensitivities between different ambient pollutants; they can be affected by external factors, such as traffic, weather changes, and human behavior; and their accuracy degrades over time. Periodic *re-calibration* can improve the accuracy of low-cost sensors, particularly with machine-learning-based calibration, which has shown great promise due to its capability to calibrate sensors in-field. In this article, we survey the rapidly growing research landscape of low-cost sensor technologies for air quality monitoring and their calibration using machine learning techniques. We also identify open research challenges and present directions for future research.

CCS Concepts: • **Applied computing** → **Environmental sciences**; • **Hardware** → *Sensor applications and deployments*; • **Human-centered computing** → *Ubiquitous and mobile computing systems and tools*;

Additional Key Words and Phrases: Air quality sensors, calibration, low-cost, machine learning, review, survey

ACM Reference format:

Francesco Concas, Julien Mineraud, Eemil Lagerspetz, Samu Varjonen, Xiaoli Liu, Kai Puolamäki, Petteri Nurmi, and Sasu Tarkoma. 2021. Low-Cost Outdoor Air Quality Monitoring and Sensor Calibration: A Survey and Critical Analysis. *ACM Trans. Sen. Netw.* 17, 2, Article 20 (May 2021), 44 pages.
<https://doi.org/10.1145/3446005>

1 INTRODUCTION

Air pollution is one of the most significant environmental challenges of our time. According to the **World Health Organization (WHO)**, in 2016, air pollution was linked to over 4.2 million deaths per year (11.6% of all deaths), with mortality in low and middle-income countries particularly

Authors' addresses: F. Concas, J. Mineraud, E. Lagerspetz, S. Varjonen, X. Liu, K. Puolamäki, P. Nurmi, and S. Tarkoma, University of Helsinki, P. O. 68 (Pietari Kalmin Katu 5), Helsinki, Finland, FI-00014; emails: {francesco, julien.mineraud, eemil.lagerspetz, samu.varjonen, xiaoli.liu, kai.puolamaki, ptnurmi, sasu.tarkoma}@helsinki.fi.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1550-4859/2021/05-ART20 \$15.00

<https://doi.org/10.1145/3446005>

heavily affected by air pollution [141]. Besides having a direct effect on mortality, air pollution is strongly associated with a broad spectrum of acute and chronic diseases, including cardiovascular diseases [5, 18, 50], lung diseases [3, 4, 49], several types of cancer [50, 105, 112], and even conditions affecting cognitive capabilities and the central nervous system [17, 132, 150]. Air pollution is also a significant economic burden worldwide, with estimates suggesting that the world spends 2–5% of overall GDP on treating diseases linked with air pollution [99]. The severity of air pollution is exacerbated by ever-increasing urbanization, with estimates suggesting that 96% of the world's population lives in areas where air pollution exceeds safe limits [72].

Understanding the characteristics of pollutants in urban environments is essential for counter-acting problems linked to poor air quality and for assessing the effectiveness of initiatives designed for tackling it. This need for detailed air quality information is driving deployments of air quality monitoring technology worldwide, particularly in metropolitan regions.^{1,2,3} Traditionally, air pollutant concentrations are monitored using professional air quality monitoring stations that meet strict accuracy criteria [14, 67, 130]. Such stations are highly accurate but also very expensive, with the cost of a single station often reaching hundreds of thousands or even million dollars [69]. Operating such stations is also costly, requiring periodic maintenance from specially trained engineers. Due to their high deployment and operating costs, professional stations are deployed sparsely, with most metropolitan regions only having a single measurement station. While in line with official recommendations, such density is not sufficient, as even a single city block can witness significant variations in pollutant concentrations. For example, congested traffic corridors, such as intersections or bus stops, tend to have significantly higher pollution concentrations than areas around them [94, 111]. To accurately assess the health and environmental risks of pollutants, it is also necessary to understand the chemical composition of pollutants, which varies depending on the season and characteristics of industry and traffic in the region [13, 135, 136]. For these reasons, accurate monitoring of air pollution inside metropolitan regions would require deploying hundreds or even thousands of air quality monitoring stations. In contrast, the WHO recommends deploying one air quality monitoring station per square kilometre,⁴ whereas the EU clean air directive suggests one station per approximately⁵ 200,000 inhabitants [31].

Low-cost air quality sensors costing less than \$10,000 have recently emerged as a way to reach higher density deployments and achieve a higher spatial resolution in air quality monitoring [54, 121, 123]. Low-cost sensors are typically small in size, making it possible to deploy them densely as part of the urban infrastructure. For example, low-cost sensors have been deployed onto light poles and public transport vehicles [6, 74]. The main drawback of low-cost air quality sensors, however, is that their accuracy tends to be poor compared to professional monitoring stations [11, 32, 87, 95]. Indeed, measurements provided by low-cost sensors can vary significantly and have poor correspondence with professional-grade monitoring stations [15], with their performance best suited for specialized tasks where exact measurements are not required, such as detecting pollution hotspots [69].

The accuracy of low-cost sensors can be improved through periodic re-calibration, with a single calibration cycle improving accuracy for up to a fortnight, before drift [63] and other errors [95] start to decrease accuracy. Periodic calibration alone, however, is insufficient, since

¹<https://www.hindustantimes.com/delhi-news/delhi-gets-18-more-monitoring-stations-to-keep-tab-on-air-quality/story-kBtKpMeuPyz0KgOeDB1z9M.html>.

²<http://www.baaqmd.gov/about-air-quality/air-quality-measurement/ambient-air-monitoring-network>.

³http://www.chinadaily.com.cn/china/2016-02/22/content_23595631.htm.

⁴https://m.economictimes.com/small-biz/startups/newsbuzz/making-sense-of-air-quality-using-sensors/amp_articleshow/69262232.cms.

⁵The recommended density varies across pollutants, with 200 000 being the average density across all pollutants.

sensors are vulnerable to cross-sensitivities between different pollutants [32] and variations in atmospheric conditions, with temperature, humidity, and wind direction being examples of factors that influence sensor performance [87]. The calibration process is also highly time-consuming and laborious, which makes it unfeasible for large-scale deployments [107].

Machine-learning-based calibration has recently emerged as a potential solution for improving the generality of calibration techniques and reducing work effort in the calibration process. The general idea of these approaches is to co-locate low-cost sensors in proximity of a professional station, which is used as a reference, and to train a model that can correct the error of the low-cost sensors using weather measurements and other sources of information. While several solutions for machine-learning-based calibration have been proposed [37, 82, 154], the overall research landscape around machine-learning-based calibration is not sufficiently understood. Indeed, there is limited information on which methods are best suited for tackling the research challenges posed by low-cost air quality monitoring stations, how to best evaluate low-cost sensor calibration techniques, and which are the other major research challenges in the area.

In this article, we contribute by surveying and critically analyzing the current research landscape on *low-cost outdoor air quality monitoring stations* and their *calibration using machine learning techniques*. We focus specifically on low-cost technology aimed at *improving the resolution* of monitoring and *increasing the density* of deployments. Previous surveys on air quality monitoring (see Section 2.1) either focus on covering specific sensor technologies or on dealing with a specific measurement challenge without comprehensively reviewing the entire research landscape. Besides reviewing existing work, we perform a rigorous analysis of the field to highlight significant open research challenges to establish a path forward.

2 SCOPE OF THE SURVEY

Research on low-cost air quality sensing has been recently gaining momentum, as sensor technology has matured to a point where increasingly large-scale deployments are possible. For example, Cheng et al. [24] consider a testbed consisting of 1,000 low-cost sensors deployed in Beijing, and Motlagh et al. [96] present a testbed with 100+ sensors located in three different districts in Helsinki, Finland. This gain in momentum also reflects in the number of published research papers on the calibration of low-cost sensors, with a query for low-cost air quality calibration returning over 180,000 results on Google Scholar. Despite this increase, the research challenges in the field are still poorly understood, and there is a lack of critical surveys assessing the state of the art. Indeed, existing surveys mostly focus on individual pollutants and specific parts of the processing pipeline, without covering issues surrounding machine-learning-based calibration⁶ of low-cost air quality sensors in depth. Our survey addresses this gap, providing a thorough review of the research field and performing a rigorous gap analysis to identify significant open research challenges in the field.

2.1 Related Surveys

Table 1 summarizes previous surveys having partially overlapping scopes with our work. In terms of sensor technology, several existing surveys focus on specific types of sensors or technology.

---

⁶In this article, we use *calibration* exclusively to refer to the use of machine learning to improve the quality and accuracy of sensor measurements. The alternative to machine-learning-based calibration is metrological calibration, where the response of the device taking measurements is adjusted to match a reference signal. Metrological calibration is typically conducted in carefully controlled laboratory conditions, whereas machine-learning-based calibration operates using measurements collected from a real-world deployment.

Table 1. Existing Related Surveys

Scope	Survey
Sensors	Gas sensor technologies [9, 74, 95, 134]
	MOS sensors [102]
	NDIR sensors [38]
	Portable sensors [119, 126]
	Wearable sensors [85]
	Commercial sensors [2, 123]
	Low-cost sensors quality [15, 16]
Deployment	Usability of low-cost AQSs for atmospheric measurements [73]
	Cities and projects [95]
Calibration	Adaptation to drift [47]
	Optical PM sensors [79, 151]
	Error sources in calibration [84]
Integration	Air quality sensor networks [145]
	Land-use regression models [58, 64]
	Satellite-based estimations [40, 122]

Table 2. Advantages and Disadvantages of Gas Sensor Technologies with Respect to Their Use within Low-cost Sensor Arrays

Type	Cost	Size	Lifespan	Sensitivity	Drift	Accuracy	Energy	Calibration	Response time	Known issues
MOS	Low	Small	Long	High	Yes	Low	High	Frequent	Fast	Cross-sensitivity to humidity and other gases. Sensitivity reduced in high temperature.
EC	Low	Small	Short	High	2–15% per year	Good	Low	Reasonable	≈ 120s	Sensitive to temperature. Low humidity and high temperatures can cause the electrolyte of the sensor to dry out.
NDIR	High	Small	Long	High	(0.4 ± 0.4)%	High	Low	Frequent	≈ 20s	Spectral interference and high detection limit. Cross-sensitivity at least to water vapour.
PID	High	Small	Long	High	20% in weeks	High	low	Frequent	Fast ≈ 1s	High sensitivity to high humidity levels or water vapour.

However, these surveys have not addressed the suitability of different technologies for large-scale air quality monitoring or how technology affects the processing pipeline.

In terms of operations performed on sensor devices, only a small number of previous surveys exist. These predominantly focus on a specific research challenge or specific parts of the pipeline. For example, Morawska et al. [95] provide an overview of deployments of low-cost air quality sensors but do not cover other parts of the processing pipeline. Gama et al. [47] provide a general review of concept drift without focusing specifically on outdoor calibration. As part of their studies

on sensor calibration, Liu et al. [79] and Zheng et al. [151] survey studies on the calibration of optical **particulate matter (PM)** sensors. Our survey supplements these surveys by providing an overview of approaches and research challenges in the calibration process of outdoor air quality sensors. Closest to our work, Maag et al. [84] provide an overview of low-cost sensor calibration, focusing on different sources of error. While the scope of our survey is similar, we address the field from a different angle, focusing on the operations needed to implement the calibration pipeline and the effect that different sensing technologies have on this pipeline. Our survey also has a broader coverage of the calibration pipeline, thus supplementing the survey by Maag et al. [84]. Having an in-depth understanding of the different sensor technologies and algorithms, as well as their advantages and disadvantages, is essential for developing best practices, identifying the most important research challenges, and advancing research in the field of low-cost air quality sensors outdoor calibration.

Finally, there have been surveys about application areas that require fusing air quality information from several sensors, such as how to use air quality information to generate spatiotemporal pollution maps [58, 64]. These surveys focus on what to do with air quality information produced by calibration models and thus supplement our survey, which addresses the operations needed to calibrate air quality sensors and to provide better quality data.

2.2 Selection of Articles

We determined which articles to include (and exclude from) the survey through a three-stage process. In the first stage, an iterative search strategy was used to determine potentially relevant articles to be included. We first identified an initial set of articles using searches with a small set of keywords on Google Scholar, IEEE Xplore, and ACM Digital Library. The following keywords were used: *air quality*, *sensors*, *low cost*, and *machine learning*. We complemented the results with follow-up works from most prominent researchers. Once we formed the initial set, we searched for articles citing them or that were cited by the articles in this set. We continued this process until we did not find more articles. In line with the interdisciplinary nature of the topic, we carried out searches separately within computer science and atmospheric sciences publications.

In the second stage, during a preliminary screening phase, we pruned the articles found by the search engines by labeling them as potentially relevant or irrelevant by one of the researchers contributing to the survey. We preserved any article relating to the scope of the survey.

In the third stage, once we identified the articles, we filtered them by asking one of the researchers that contributed to the survey to read the articles and to present the main findings for the other authors. We then decided through a majority decision whether the article was within the scope of the survey. In this survey, we aimed to focus on applying machine learning for the re-calibration of low-cost sensors. Hence, we selected only papers that used machine learning in their re-calibration process. However, we also evaluated several papers specific on low-cost sensors implementations to provide insights into the different sensing technologies in air pollution monitoring.

3 LOW-COST AIR QUALITY SENSING PIPELINE

Low-cost air quality sensing follows a typical machine learning pipeline for sensor data, illustrated in Figure 1. Within the pipeline, we can separate two types of operations: *per-device* operations, which need to be performed separately for each sensor, and *integration* operations, which combine data from multiple sensing units. Note that this distinction does not mean that the devices cannot collaborate, but instead, whether a specific operation needs to be applied to each device or the aggregated data. Indeed, calibration functions usually use measurements from multiple different sensors to construct the underlying calibration model, but, at runtime, the

Fig. 1. A reference calibration pipeline for low-cost air quality sensors. Different services may use variants of the pipeline.

model is applied separately to each device, making this a per-device operation. The specifics of the pipeline can vary for different services, even if the general structure remains similar. For example, local pollution prediction requires a different time scale and update interval than daily city-level pollution estimates [26, 39, 78]. In this survey, our focus is on per-device operations, as several existing surveys are focused on application domains that are related to integration operations [40, 58, 64, 122, 145]. Below we briefly give an overview of the different steps in the sensing pipeline.

3.1 Per-Device Operations

The low-cost air-quality sensing pipeline includes six operations that address issues related to individual devices: (1) the design of low-cost sensing units, (2) deciding where the devices should be deployed, (3) the actual data collection, (4) pre-processing, (5) **machine learning- (ML)** based calibration and (6) the evaluation of that calibration. Applications, such as prediction and real-time maps, can then be built on the results of per-device data.

Low-cost Sensing Design. A sensing unit can be seen as a low-cost monitoring station integrating one to several sensors, each measuring a specific pollutant or environmental variable. In current air quality sensing research, the design of the sensing units is largely overlooked. Indeed, most of the research relies on off-the-shelf sensing units and focuses on other parts of the pipeline. However, the choice of sensor models and the overall design of the sensing unit can have a significant impact on how these sensors produce data. For example, as we discuss in Section 4, some sensors need to be heated, while others are sensitive to temperature changes. When sensors for several pollutants are integrated into the same sensing unit, heating some sensors can result in significant inaccuracies by sensors sensitive to temperature changes. Another concern is related

to variations in sampling rates that can make it difficult to synchronize measurements from different sensors or to relate the measurements with real-world events. We discuss the properties of low-cost gas and particle matter sensors in Section 4.

Sensor Placement. The sensors can be mobile, mounted onto vehicles, carried as personal sensors (i.e., wearables), or deployed to specific areas of cities. The benefit of mobile sensors is that their measurements can cover a large area. However, this can result in high sparsity, as many areas will only have a small set of measurements. Another challenge with mobile sensors is how to manage them and how to ensure they are operational. Instead of relying on mobile sensors, most real-world deployments rely on sensors placed in fixed positions for an extended time. Examples include the Barcelona Lighting Masterplan⁷ and the Chicago Array of Things.⁸ In current research, the deployment of sensors is largely driven by practical constraints, such as availability of power and availability of locations where the airflow to the sensor remains unobstructed. We briefly discuss sensor placement in Section 4 and refer to the survey by Morawska et al. [95] for a more thorough overview of existing sensor deployments.

Data Collection. This step consists in collecting air pollution data, i.e., the concentration of selected air pollutants and environmental parameters. Pollutants require different collection densities, therefore, not every pollutant needs to be collected by all devices [125]. For example, temperature and humidity have similar patterns within one region, whereas pollutants resulting from vehicles can fluctuate significantly, even within a small region. Data collection is discussed in Section 6.1.

Data Pre-processing. Examples of pre-processing operations include synchronization of different measurements and removal of erroneous measurements from periods where the device is compromised. For instance, the devices could have been operating in extreme conditions that are not supported by the sensors (e.g., high temperatures), the measurement units may be clogged, or power spikes can disrupt the functionality of a sensor. Other common pre-processing techniques include interpolation to achieve a consistent sampling rate and aggregation to reduce sampling rate, e.g., when low-cost sensors produce data at a higher rate than professional-grade reference stations. State-of-the-art low-cost air quality research typically considers measured pollutant concentrations and environmental variables (such as temperature and relative humidity) as features. Pre-processing is discussed in detail in Section 6.1.

Calibration Based on Machine Learning. This step consists in the application of machine learning techniques to calibrate the measurements of the low-cost sensors, and it is the main focus of our survey. We critically compare existing machine learning solutions for the calibration of low-cost air quality sensors and identify the main advantages and disadvantages of the methods. Machine-learning-based calibration is discussed in detail in Section 6.1.

Sensor Calibration Evaluation. The final step related to a single sensing unit is the performance evaluation of their calibration. We discuss this step in detail in Section 7 with an emphasis on the selection of performance measures and test data length.

Integration Operations. Air quality data from a single sensor is limited in usefulness as it captures pollutant concentrations only within a small region. In practice, to create high-quality spatiotemporal air quality information with fine granularity, we need to combine data from numerous sensing units. This operation is conducted in the seventh step, which includes spatiotemporal modeling and fusion with additional sources of information, including but not limited

⁷<http://ajuntament.barcelona.cat/ecologiaurbana/en>.

⁸<https://arrayofthings.github.io>.

to land-use, weather, and traffic data. The eighth step includes the performance evaluation of the models produced by the previous step. Finally, the ninth and last step is the production of air quality services built upon high-quality air pollution models. Such services include more advanced air quality index [101] models and green path routing [53] to enhance the quality of life of citizens. We note that while such services are not strictly part of the calibration pipeline, air quality services can have contrasting accuracy, resolution, and other requirements. These requirements, in turn, affect the calibration pipeline, and thus the services are intrinsically linked with it.

4 LOW-COST AIR QUALITY SENSORS

Low-cost solutions for air pollution monitoring typically consist of sensing units that package together multiple low-cost air quality sensing components. Besides the components responsible for monitoring pollutants, sensing units can incorporate other components, such as power sources, processing units, local data storage, networking interfaces, and atmospheric sensors (e.g., temperature, relative humidity, and wind direction) [41]. Individual components usually cost between \$20 and \$100. However, a complete low-cost sensing unit typically costs upwards of \$500 [20]. In this article, we use the term *sensor* to refer to a single sensing component responsible for measuring a (single) pollutant or environmental variable, whereas we use the term *sensing unit* to refer to the overall device used to collect measurements for the individual studies in question.

In this section, we survey the most commonly used sensor technologies employed by low-cost air quality sensing units and discuss how their characteristics affect in-field calibration. Characteristics of components and the overall design of a sensing unit can have a significant impact on resulting measurements, and hence need to be carefully chosen to avoid negatively influencing other parts of the air quality monitoring pipeline. Indeed, as we discuss in Section 7.3, the best model for a pollutant rarely is the same that works well for another pollutant. For this reason, it is important to understand the characteristics of the underlying sensor technology, as well as the key advantages and disadvantages of the sensor technology. Before delving into the details of the sensor technologies, we briefly discuss the two different types of pollutants considered in air quality research.

4.1 Types of Pollutants

Air quality sensing research typically categorizes pollutants as gaseous or particulate matter, according to the composition of the pollutant [109]. Commonly considered gaseous pollutants include **carbon monoxide** (CO), **carbon dioxide** (CO_2), **ozone** (O_3), **nitrogen oxides** (NO_x), and sulfur oxides. Carbon dioxide monitoring is restricted to indoor environments. In outdoor monitoring, the most common gases are the ones belonging to prominent air quality indexes, such as EPA in the USA or the Air Quality Index of China. Specifically, these indexes include the following gases: **sulfur dioxide** (SO_2), O_3 , CO and **nitrogen dioxide** (NO_2). Particulates (or aerosols), however, refer to tiny particles of solid or liquid compounds suspended in gases. The source of particulates can be natural (e.g., dust or sea spray) or caused by human activity (e.g., burning of fossil fuels or wood, dust from roads and tires, and power plants).

Gas sensors can be typically tailored to support different gases by changing (parts of) the sensing materials or operating parameters of the sensing unit. Particulate matter sensors, in contrast, only monitor the extent of particulates in the air without being able to identify their exact source or composition. However, particulate matter sensors can be categorized based on the size of the particles they can monitor, with fine ($PM_{2.5}$) and coarse-grained (PM_{10}) being the most common categories in low-cost air quality research and belonging to all major air quality indexes. With some sensor technologies, it is also possible to detect so-called ultra-fine particulates ($PM_{0.1}$). However,

these mostly require expensive professional-grade instruments and thus are rarely considered in low-cost sensing research.

4.2 Sensing Gaseous Pollutants

Within low-cost sensing units for gaseous pollutants we can identify four main types of sensing technologies: **metal oxide semiconductor (MOS)** sensors, **electrochemical (EC)** sensors, **non-dispersive infrared (NDIR)** sensors, and **photo-ionisation detector (PID)** sensors. In this section, we describe the key properties of these technologies. We discuss MOS and EC sensors in more detail as they are the least expensive and most widely used technologies in low-cost sensing units. NDIR and PID sensors have higher costs than MOS and EC sensors (\$500–\$5,000); hence, they are rarely used in low-cost sensing units. We note that there are also other technologies for monitoring concentrations of gaseous air pollutants, but these have even higher costs, making them unsuitable for low-cost sensing units. For example, gas chromatography sensors cost between \$15,000 and \$100,000.

Solid-State Metal Oxide Sensors

MOS sensors are a popular sensor technology for monitoring gas concentrations of several pollutants, such as **non-methane hydrocarbon (NMHC)**, CO, CO₂, NO₂, NO_x, and O₃ [36, 37, 55, 85, 100, 103, 114, 120, 121]. MOS sensors consist of a heating element and a semiconducting metal oxide sensing element. The heater warms the surface of the sensing element up to 300°–500°C, which is then able to detect gases through a chemical reaction occurring on its surface. This reaction causes a change in the electrical conductivity of the sensing element, which can be monitored using an external circuit to measure the detected gas level [102].

Advantages: MOS sensors are very low-cost and compact. Furthermore, these sensors have a high sensitivity and can even reach sub-**parts per billion (ppb)** sensitivity for some gases. MOS sensors also have a short response time, i.e., they can produce data at a high sampling frequency. Other advantages of MOS sensors include their long lifespan and resilience against extreme weather conditions. Indeed, MOS sensors can operate in high temperature and humidity environments, making them well suited for (very) long-term deployments.

Disadvantages: MOS sensors have several disadvantages that can affect other parts of the processing pipeline. First, while they are resilient, their sensitivity is affected by atmospheric conditions and other gases. In terms of calibration, this implies that inputs from temperature and relative humidity sensors are needed while calibrating MOS sensors. Another disadvantage of MOS sensors is that they tend to have low accuracy and are subject to drift [134], which requires them to be re-calibrated often to maintain good-quality data outputs [86, 127]. Drift manifests as a reduced measurement accuracy, which results from a decrease in the conductance of the sensing element over time [110]. The response of MOS sensors also depends on humidity, with higher humidity resulting in higher sensor error [117]. Finally, MOS sensors require access to a sufficiently large power source due to their need to power an electric heater.

MOS sensors are well suited for low-cost sensing units due to their low cost and high resilience against environmental conditions. However, their high power requirement and sensitivity to environmental conditions are a concern. High power requirements make MOS sensors better suited for deployments where a fixed power supply is available than for deployments requiring battery power. As an example of the use of MOS sensors for air quality monitoring, Hasenfratz et al. [54] used a MiCS-OZ-47 sensor as part of a participatory air quality monitoring campaign, where they

connected the sensor to a separate battery pack and used a smartphone to store and transmit measurements. The authors evaluated the impact of the sensor on the battery, which resulted in an estimated operation time of 50 hours when using a separate battery. Burgues et al. [19] state that shutting down the heating elements for some time and then heating them in cycles rather than continuously can reduce the power consumption of gas sensors up to 90%. However, the energy savings come at the cost of reduced measurement accuracy, suggesting that the duty cycle of the sensors needs to be taken into consideration when building calibration models. In particular, two MOS sensors with similar specifications but different duty cycles are likely to have different accuracy characteristics that need to be taken into account when training and sharing measurements or calibration models across devices. Another alternative is to design the duty cycle to maximize measurement quality. For example, if multiple measurements are collected in the same cycle and the first few measurements are dropped, then the resulting measurements are more accurate than when only a single measurement is collected in each cycle [19].

Recently, there have been advances in the synthesis of MOS gas sensing materials that can ensure sensitivity remains high even when air humidity increases [131]. These materials are not yet available in mass quantities, hence they are not used by current low-cost sensors.

Electrochemical Sensors

Electrochemical sensors are the other popular low-cost technology for monitoring gas concentrations. EC sensors have been used to monitor CO, **nitrogen oxide (NO)**, NO₂, O₃, and SO₂ [43, 61, 98, 100, 114, 120, 121, 154]. These sensors detect gases by oxidation-reduction reactions, employing electrodes separated by an electrolyte substance, such as mineral acid. The working electrode is in contact with both the electrolyte and the ambient air, which is monitored via a porous membrane. The reaction produces an electrical current between the electrodes, which can be measured from the outer pins of the sensor.

Advantages: Similarly to MOS sensors, EC sensors are inexpensive, have high sensitivity (in ppb levels for some gases), and good specificity. Compared to MOS sensors, the main benefit of EC sensors is their lower power draw as they do not require powering an electric heater. Another advantage of EC sensors is that their sensitivity is less affected by temperature and humidity than with MOS sensors.

Disadvantages: The main drawback of EC sensors is that the speed of the chemical reaction depends on the operating conditions. The operating range of the sensors is also dependent on the characteristics of the chemicals used in the sensor and tends to be narrower than in MOS sensors. EC sensors also have a shorter lifespan than MOS sensors, with the overall duration depending on the amount of pollution they are exposed to. As an example, the Alphasense NO₂ sensor used by Castell et al. [20] has a lifetime of 2–5 years, whereas MOS sensors can last 10 years or even longer. For example, Romain and Nicolas [110] used MOS sensors to collect data over a 7-year period. EC sensors are also less resilient to weather conditions than MOS sensors. A combination of low humidity and high temperature is particularly problematic to EC sensors, as it can dry out the sensor's electrolyte and break the sensor. Another drawback with EC sensors is that other gases may interfere with the measurements, even though they are less sensitive than MOS sensors [126].

Electrochemical sensors are well suited for low-cost deployments, as the sensors are inexpensive and their performance is less affected by temperature and humidity variations than MOS sensors. The accuracy of EC sensors tends to be good, as long as the weather conditions fall within their operational range. Examples of research relying on EC sensors include Nikzad et al. [98] and Wei et al. [140]. Ultimately the choice between MOS and EC sensors depends on the goals

of the deployment. In areas with fairly stable temperature and weather conditions, EC sensors are well suited due to their better specificity and accuracy. Also, the scale and intended duration of the deployment affect the choice of technology, with MOS sensors being capable of supporting deployments lasting 10 years or more, whereas EC sensors tend to be limited to deployments lasting few years at most. In terms of calibration, the two technologies are reasonably similar, with the main difference being how often the calibration functions need to be re-trained and validated.

Non-dispersive Infrared

NDIR sensors consist of an infrared light source, an atmospheric sampling chamber, an optical filter, and a detector. When a gas passes through the chamber, the light emitted by the infrared source travels through it, and some frequencies get absorbed depending on the gas. The rest of the light hits the optical filter and the detector, which outputs the frequencies through an electrical current. NDIR sensors have been mostly used to detect CO₂ concentrations [103, 120, 121, 154]. However, they can be used to detect also other gases through changes in the wavelength of the light.

Advantages: They are simple and require little power, and small units also are available. They have a long lifespan (not degraded by exposure to gases) and require only little maintenance.

Disadvantages: They have high detection limits, i.e., cannot measure small pollutant concentrations, and they are susceptible to spectral interference from different gases as well as water [126]. NDIR sensors are also subject to drift [38] and they cost considerably more than MOS or EC sensors (up to a 10-fold increase in price).

Even considering the disadvantages, the long lifespan of NDIR sensors makes them a good choice for long-term deployments in dry areas. The high cost and high detection limit mean that they are better suited for sparser deployments than what most low-cost sensing aims at accomplishing.

Photo-ionisation Detectors

PID sensors operate by illuminating compounds using high energy UV photons. The process results in compounds becoming ionized as they absorb the UV photons and results in an electrical current that can be captured by a detector inside the sensor. The greater the concentration of the measured component in the compound, the more ions are produced, and the greater the current.

Advantages: photo-ionisation detectors (PIDs) are very sensitive and have a short response time. They are small in size and weight, and they have low energy requirements [1, 2].

Disadvantages: PIDs affect all molecules whose ionization potential is lower than the UV light affecting them, which means PIDs are not specific to a particular pollutant [2]. PIDs are sensitive to high humidity levels or water vapor. They are not very low-cost (between \$500 and \$5,000), even if affordable compared to high-end monitoring stations [119]. PIDs are also subject to drift and need to be re-calibrated often (once per month).

The ability of PID sensors to analyze samples of low concentration in ambient temperatures and pressures makes them suitable for deployment in a wide range of environments, which has made PID sensors particularly well suited for the analysis of small particles and gases in controlled small-scale experiments. However, the sensitivity to water and the relatively high cost make PID sensors poorly suited for dense long-term deployments.

Spinelle et al. [120, 121] studied the performance of MOS and EC sensors for detecting O_3 , NO , NO_2 , and CO . The authors concluded that no significant differences in sensor outputs can be observed between the two sensor technologies. In total, they evaluated 15 sensor models for the five gases. However, their study included only five months of data, which is within the range of the lifespan of EC sensors, therefore EC sensor degradation effects were not considered.

Sensor mobility has been discussed in numerous studies that we review in this survey [26, 48, 54, 80, 82, 85]. For example, Hasenfratz et al. [54] simulated mobility by carrying experiments in a room with constant O_3 concentration and artificial wind, generated using a table fan. The authors found that, when O_3 concentration is low, the wind does not affect the measurements much. However, when O_3 concentration is high, wind effects produce a measurement offset. Therefore, they recommend shielding the sensor from the wind when moving relatively fast, for example, when riding a bicycle.

Generally, the choice between sensor technologies depends on the characteristics of the deployment. MOS and EC sensors are cheapest and thus best suited for large-scale deployments, whereas PID and NDIR sensors are better suited to sparser deployments. In terms of accuracy, MOS and EC sensors have comparable performance, with EC sensors being more energy-efficient, but MOS sensors being more resilient. In particular, EC sensors do not require an electric heater, but they can break in high temperatures and low humidity conditions. MOS sensors also are more durable than EC sensors. In terms of calibration, both sensors are somewhat sensitive to weather conditions and concentrations of other gases, suggesting that these variables need to be incorporated into the calibration model. In terms of sensor design, MOS sensors can cause cross-interference to other sensors by heating the air inside the sensing unit. Thus, the sensing unit and the sensor cycles need to be carefully designed to minimize the risk of such effects.

4.3 Particulate Matter Sensors

Particulate matter sensors monitor particle density by assessing changes in the properties of air passing through the sensor unit. Unlike gas sensors that can be tailored to detect different pollutants, particulate matter sensors cannot identify the exact composition of pollutants. However, they can be adapted to recognize particles of different sizes.

Existing low-cost particulate matter sensing technology is predominantly based on **Diffusion Size Classifiers (DiSCs)** or **light-scattering particle sensors (LSPs)**. DiSC operate by charging air passing through the sensor and estimating particle density from the total electricity charge after applying different filtering operations on the charged air. LSPs operate by using light scattering to estimate the density of the particles. Traditional laboratory-grade instruments for particulate matter sensing are based on similar principles, but use additional components to improve detection accuracy. For example, **optical particle counters (OPCs)** are high-quality variants of LSP, whereas **condensation particle counters (CPCs)** use alcohol or water vapor to change the physical properties of particulates before passing them through an LSP sensor [22, 52, 115, 118]. However, these sensors are typically bulky and more expensive than basic DiSC and LSP and hence rarely used in low-cost air quality sensing.

At the top end of the scale, **Tapered Element Oscillating Microbalance (TEOM)** sensors use changes in the oscillation frequency of a vibrating glass tube, and **beta attenuation monitor (BAM)** sensors use absorption of beta radiation for estimating particle density. TEOM and BAM sensors are quite expensive, costing over \$20,000 [115]. Another high-end option is scanning mobility particle sizer, which estimates the size and concentration of particles using electrical mobility sizing to monodisperse the output, which is then monitored with a CPC sensor. In the following subsections, we describe DiSC and LSPs in more detail, since these technologies are the most affordable and most widely used in low-cost air quality sensing research.

Diffusion Size Classifiers

DiSC detect particles by applying electrical signals to identify physical changes on the sensing surface. They contain an air inlet, a corona charger, an induction stage, a diffusion stage, and a backup filter. When air enters the sensor, particles go through a diffusion charger that produces ions using a corona wire. A small fraction of these ions attach to particles in the air, and an ion trap placed between the diffusion charger and the induction stage captures the remaining ions. The particles then pass through the induction stage, where they produce a small electrical current that is proportional to their concentration. After leaving the induction stage, the particles arrive at the diffusion stage that precipitates the particles to produce a small electrical current, which is proportional to their concentration. Since the particles also have an induction effect in the diffusion stage, the current measured in the induction stage is subtracted from the current measured in the diffusion stage to compensate for the induction effect. Larger particles not precipitated by the diffusion stage eventually reach the backup filter, which produces an electrical current proportional to their concentration [44, 89]. This way DiSC sensors can separate between different particle sizes.

Advantages: The sensitivity of DiSC sensors is extremely high, and they have low power consumption. Manufacturing costs are low when the sensors are manufactured in large bulks.

Disadvantages: Manufacturing costs of DiSC have a high upfront setup cost, as they require clean rooms and other special facilities. Therefore, the production and assembly of low quantities of sensors have a high unit cost ($\approx \$10,000$). Also, testing equipment for assessing the quality and performance of DiSC sensors can be expensive. Another problem with DiSC sensors is that the sensing area can become unclean, making it necessary to clean it frequently. Research on how to automatize the cleaning process, e.g., using oscillation and electrostatics, is actively pursued [66, 89].

Light-Scattering Particle Sensors

LSPs are small, low-cost sensing units widely used to detect particulate matter [26, 28, 68, 80, 81]. They are composed of an air inlet, a light sensor, and a light source, usually infrared or laser. When air enters the sensor, the light source is focused on a sensing point. An infrared LED is positioned at a forward angle relative to a photodiode. Particles passing through the light beam scatter light, which generates a measurable signal in the sensor circuitry. The scattered light is focused onto the photodiode by a lens. The sensors may contain a light scattering focusing lens and a focusing lens also for the infrared light source. The resolution at which different particle sizes can be detected depends on the configuration of these lenses. Finally, the sensor produces a signal that can be measured to estimate the number of particles in the air [29].

Advantages: LSPs are small and their cost is very low. The sensors mostly require very low power [28].

Disadvantages: Light-scattering based instruments fail to detect very small particles [44]. The sensor readings are also impacted by temperature and relative humidity, which means both temperature and humidity need to be measured. More expensive LSPs can use multi-angular light scattering to reduce the impact of environmental variables [115].

Table 3 summarizes the advantages and disadvantages of the two particulate matter sensing technologies used in low-cost air quality monitoring research. Similarly to gas sensors, the optimal choice of sensing technology depends on the context of the deployment. DiSC sensors have high

Table 3. Advantages and Disadvantages of particulate matter (PM) Sensor Technologies with Respect to Their Use within Low-cost Sensor Arrays

Type	Cost	Size	Lifespan	Sensitivity	Drift	Accuracy	Energy	Calibration	Response time	Known issues
DiSC	High	Small	Good	Good	Yes	Good	High	Yearly (Factory),Hourly (Software)	1 s	The instrument can produce wrong results if the incoming aerosol is highly positively charged. Cannot distinguish between narrow and broad particle size distributions [44].
LSP	Ultra-low	Small	Good	Poor	None	Low	Low	Frequent	≈30 s	Mixes all particle sizes, variation of air influx reduces accuracy.

sensitivity and are mostly unaffected by weather or other environmental conditions. However, they suffer from the need for regular maintenance to ensure the detection surface remains sufficiently clean. LSPs, however, are susceptible to weather changes but require less maintenance, making them better suited for longer-term deployments. While other technologies for particulate matter sensing exist, they are not well suited for low-cost deployments due to the higher cost of sensing equipment and larger size of the sensing units.

Weather affects the concentration of particles in the air, and thus the sensor measurements. Therefore, particulate matter sensors generally require information about weather conditions regardless of whether the underlying sensor technology is sensitive to weather or not. High winds can disperse particulate concentration, whereas high humidity causes particulates to cluster together, increasing its concentration. The effect of temperature, however, is less well understood. Zheng et al. [152] found higher temperatures to lower particulate concentrations when humidity is low, and lower temperatures to decrease particle concentrations when humidity is high. Wang et al. [138] studied the effects of temperature and humidity on low-cost PM sensors and found relative humidity to affect the accuracy of sensor technology. For example, as water in the air absorbs infrared radiation, humidity can result in LSP overestimating particle concentrations as light intensity is reduced. The temperature was not found to directly impact the sensor technology, even if it affects the concentration of particles.

Besides weather, the concentration of particulate matter is affected by the extent of human activity within the area being monitored. The higher the traffic density and the lower the fuel efficiency of the vehicles, the higher the concentration of particulates will be. Note that the density is not solely a result of fuel burning as also tire and road wear produce particulates. In terms of calibration, this implies that the context of the deployment needs to be taken into consideration as locations close to intersections are likely to have an increased particulate matter concentration than other nearby areas [152]. Several research initiatives have explored the possibility of mounting low-cost sensors on vehicles [26, 48, 56]. When the vehicles are moving, the input air flux is constantly changing, which can affect sensor accuracy. In particular, the more air enters into the sensor, the more pollutants will be detected, even if the concentration of particulates would remain constant. Besides vehicles, wind speed can trigger a similar effect. To ensure that an accurate calibration model can be constructed, a possible solution is to include a GPS—or another sensor to estimate the velocity of the sensor unit at the time of measurement—and to use this information to compensate for the density of pollutants detected by the sensor [48].

5 DATA COLLECTION AND PRE-PROCESSING

Low-cost outdoor air quality monitoring commonly relies on sensors deployed as part of the urban infrastructure [95]. The most common approaches are to deploy the sensors as part of fixed infrastructure, such as street lights, or to mount the sensors onto vehicles, such as trams [74], garbage trucks [116], or even Google Street View vehicles [111]. Next, we briefly describe the characteristics of the data collected by low-cost sensor deployments and typical preprocessing operations performed on the measurements.

5.1 Measurements

Low-cost air quality sensors measurements can be interpreted as time-series data consisting of values of different pollutants and environmental variables. The integration of several pollutants is necessary to capture the effects of cross-sensitivities [32], whereas environmental variables are critical for accounting for differences in error as environmental conditions change. As discussed in Section 4.1, the most common pollutants to consider are gaseous pollutants and particulate matter included in prominent air quality indexes. In terms of environmental variables, temperature and relative humidity are the most common variables that need to be considered. Wind speed is another variable that can influence pollutant concentrations. However, wind speed is often difficult to measure with low-cost sensors, as it requires an unobstructed air intake, whereas other environmental variables can be more accurately and reliably measured with low-cost sensors.

Reference Measurements. Calibrating low-cost sensors with machine learning requires access to high-quality reference measurements to consider as ground truth. The most common choice is to deploy low-cost sensors near a high-quality atmospheric station, and use the measurement of the station as ground truth [37, 43, 120, 121]. Another possibility is to use a high-quality mobile measurement laboratory deployed near the low-cost sensors [154]. Generally, the closer the low-cost sensor is to the reference station, the better an appropriate calibration relationship can be established. In cases in which the ground truth is needed in multiple different locations, for example, when calibrating mobile sensors, other approaches must be used [96]. For example, public high-quality pollution data from official authorities can be used [54].

Temporal resolution. The temporal resolution of measurements also influences the calibration process. The resolution is governed by the sampling frequency of the sensors, which varies across different sensor technologies. For example, as MOS sensors require heating, they have a slower sampling rate than sensors that can operate continuously—unless the heating element is run continuously, which would result in prohibitively high power drain. In most of the studies surveyed for this article, the sampling rate is between 5 s and 20 s [37, 82, 114, 154]. However, there are also studies with sampling intervals as high as one hour [16] or as low as 10 ms [120, 121].

5.2 Characteristics of Air Quality Measurements

Low-cost sensor measurements have some characteristics that need to be accounted for while designing calibration solutions. These characteristics all affect the statistical properties of the measurements and can result in natural variations in measurements or from different sources of errors. As they affect the measurements used as input for calibration algorithms, they need to be accounted for when selecting appropriate algorithms. Below we briefly discuss the most important characteristics.

Autocorrelation. Air pollutant concentrations are known to have a strong spatiotemporal correlation with the weather. Furthermore, seasonal patterns also have a significant influence on them [27, 90, 113], which means that the used calibration techniques need to be capable of dealing

with autocorrelation. The data used to test the generality of the model should also be sufficiently long-term to ensure the results are not overfitting to short-term correlations.

Cross-sensitivities. Measurements provided by low-cost sensors suffer from cross-sensitivities between pollutants [32]. Measurements can also be affected by temperature, humidity, and wind direction [87]. In terms of calibration models, this means that the used techniques cannot assume the variables to be independent, but instead, they need to consider the values of environmental variables, and potentially also the values of other pollutants, as input.

Drift. Low-cost air pollutant measurements are vulnerable to *drift* whereby the relationship between environmental variables and pollutants varies over time [47]. For example, an analysis of MOS (see Section 4) has demonstrated that the measurements can differ by over 200% over time [110]. The most common reason for drift is wear. For example, metal oxide sensors are vulnerable to oxidation, which alters the conductivity of the sensing element and results in drift [10], whereas light-scattering particle sensors are vulnerable to deposits forming on the lens of the optical sensor [7]. The traditional way to handle drift is to perform maintenance on the sensors, which requires cleaning and replacing components and re-calibrating the sensors in laboratory conditions. For large-scale deployments, this is not feasible, and an alternative approach is re-training the calibration function to account for changes in the properties of the sensors. The frequency of re-calibration depends on the characteristics of the deployment. For example, light scattering particle sensor re-calibration frequency depends on the extent of pollutants. The higher the amount of pollutants, the more often re-calibration is required.

Concept Drift. Besides being vulnerable to mechanical or chemical drift affecting the sensor hardware, field deployments are vulnerable to *concept drift* where the statistical characteristics related to estimation of the target variables change over time. These changes can be a result of persistent effects, such as clean air policies or changes in human consumption patterns [88], or temporal effects, such as forest fires, volcano eruptions, or other weather phenomena [106]. In terms of calibration function, concept drift can be accounted for by adapting the underlying calibration function or re-training it, depending on the magnitude of changes. De Vito et al. [35] discuss how the risk of concept drift can be recognized by analyzing the statistical difference in distributions of air pollution measurements.

Height differences. The most common placement of low-cost sensors is near the ground level without isolating them from the urban infrastructure. Professional-grade measurement stations, however, are typically at least partially isolated from the urban infrastructure, and they have different air intakes located in different parts of the sensor. For example, in Finland, reference stations are either in dedicated containers that can be near the ground or as part of separate measurement towers that are jointly responsible for weather and pollution measurements [67]. Pollutant concentrations can vary significantly also with elevation. For example, seasonality influences the elevation of the atmospheric mixing layer [124], which in turn affects the extent of pollutants that can be captured [133].

5.3 Preprocessing

The measurements from low-cost sensors typically require preprocessing before they can be used to capture a calibration model. Below we briefly describe the most common operations.

Resampling. Before training a calibration model, measurements need to be resampled to a suitable temporal resolution. A higher resolution implies more samples and a longer model training time, but it can improve the robustness of the model. A common choice is to use a 1-minute resolution [43, 48, 120, 121]. Some studies use a much coarser resolution, such as one hour [16, 37],

which is the standard resolution for deriving air quality index values [91]. Saukh et al. [114] use a resolution in tens of seconds, while Hasenfratz et al. [54] and Maag et al. [82] use 5 s.

Synchronization. To minimize cross-sensitivities, sensor sampling intervals need to be interleaved. For example, as MOS sensors require heating, they can influence measurements for temperature or other pollutants unless the heating period is sufficiently distant from the sampling period of other sensors. To account for differences in sampling times, the measurements need to be synchronized during the modeling phase. This can be accomplished using aggregation, e.g., using the mean value over a given data window, or interpolating values of different sensor units to have consistent timestamps. In most cases, using a simple linear interpolation is sufficient, especially if the synchronization window is short.

Smoothing and Filtering. Pollutant measurements occasionally contain outliers that can significantly decrease the performance of the calibration model, unless the outliers are accounted for. There are several reasons for outliers. For example, sudden wind gusts can result in abnormal measurements, or the air intake of the sensing unit may get temporally obstructed. Common ways to mitigate these issues are to use *smoothing* or *filtering*. In the former case, the measurements are fitted to a model that ensures temporal consistency, whereas in the latter case, values appearing as abnormal or otherwise erroneous are removed. As an example of the former, Cheng et al. [26] smooth the sensor data through a signal reconstruction technique based on a bi-criterion problem with a quadratic smoothing function. As an example of the latter, Hasenfratz et al. [56] used a three-phased filtering process. First, each sensor computes its null-offset and uses it to calibrate the offset of the measured particle concentration. Second, the sensing units check that the sensors are operating correctly and measurements from periods where sensor failures are detected are removed. Third, measurements with poor location accuracy are removed from consideration. Another example of filtering is proposed by Gao et al. [48] who also use GPS data to filter measurements. Smoothing and filtering are popular techniques for preprocessing, and as such, are likely to be used in most studies. However, most studies we surveyed for this article do not indicate which kind of preprocessing has been applied [30, 120, 121, 154].

6 MACHINE LEARNING CALIBRATION OF LOW-COST AIR QUALITY SENSORS

Low-cost sensors increasingly rely on machine-learning-based calibration pipelines to improve the accuracy of sensor measurements. Previous surveys have shown that domain characteristics need to be taken into account when building such calibration pipelines [79, 84, 151]. In the following, we survey existing calibration pipelines for low-cost air quality sensor data and discuss how they address requirements stemming from the specifics of air quality measurements. Generally, these approaches learn the relationship between a specific pollutant, as given by the low-cost sensor, and a ground-truth value obtained from a reference station. We first discuss issues that affect the choice of machine learning algorithms, after which we survey machine learning algorithms used in previous studies. The different models and how they have been applied are summarized in Table 4.

Figure 2 presents a high-level overview of the (continuous) machine learning calibration process. A reference station provides reference data to train calibration models that can correct the measurements of a particular low-cost sensing unit. Note that each sensor (type) can have a different model, as shown in the figure. The performance and usefulness of the models are evaluated against new measurements, similarly collected from low-cost sensor units and a reference station. Once the error of a calibration model is sufficiently small, the corrected measurements can be used by many air quality sensing applications, such as pollution monitoring and prediction, and high-definition pollution maps based on spatiotemporal models [12, 42, 64, 137].

Table 4. Summary of Calibration Studies

ML model	Reference	Training data length	Test data length	Temporal resolution	Mobility	Exploits temporality	Online training
LR	Hasenfratz et al. [54]	NR	≤ 2 mos.	5 s	X		
	Lin et al. [75]	NR	≤ 2 mos.	5 mins			
	Saukh et al. [114]	NR	≤ 6 mos.	10 s, 20 s	X		
MLR	Maag et al. (2016) [82]	4 wks.	1.25 yrs.	5 s	X		
	Maag et al. (2018) [85]	2.1 days	2.7 wks.	NR	X		
	Liu et al. [80]	NR	36 h	1 min	X		
	Cordero et al. [30]	NR	≤ 30 days	NR			
	Zimmerman et al. [154]	NR	1.4–15 wks.	15 mins			
SVM	Cordero et al. [30]	NR	≤ 30 days	NR			
RF	Borrego et al. [15, 16]	12.6 days	1.4 days	1 min/1 h			
	Cordero et al. [30]	NR	≤ 1 mo.	NR			
	Zimmerman et al. [154]	5.6 days	1.4–15 wks.	15 mins			
FFNN	DeVito et al. (2008) [36]	8 mos.	3 mos.	1 h		X	
	DeVito et al. (2009) [37]	2 wks.	7 mos.	1 h			
	Spinelle et al. [120, 121]	1 wk.	4.3 mos.	1 min			
	Esposito et al. [43]	1 wk.	3 wks.	1 min			
	Borrego et al. [15, 16]	12.6 days	1.4 days	1 min/1 h			
	Maag et al. [85]	2.1 days	2.7 wks.	NR	X		
	Cordero et al. [30]	NR	≤ 30 days	NR			
FFNN + GP	Cheng et al. [26]	3.5 mos.	2 mos.	5 mins	X		
	Gao et al. [48]	NR	NR	1 min	X		
NARX	Esposito et al. [43]	1 wk.	3 wks.	1 min		X	
TDNN	Esposito et al. [43]	1 wk.	3 wks.	1 min		X	

Fig. 2. Dataflow diagram of machine-learning-based calibration.

6.1 Issues for Calibration

Generality. Calibration models should be capable of operating under different conditions, including different geographic locations, across different seasons, and potentially also across variations in sensing units. In practice, simultaneously achieving all these goals is not feasible. A model might need to be periodically retrained to adjust for variations. Different models might be needed

for different seasons, hardware, or even locations. In the simplest case, such models correspond to reparametrizations of the same type of model, whereas in some situations it may be necessary to use different classes of models altogether, e.g., a neural network in one setting and a linear regression model in another.

Distribution of Air Quality Data. Many common machine learning techniques, such as common regression models, have been designed for data that are independent and identically distributed (i.i.d.), or at least close to being. As discussed in the previous section, this rarely is the case for air quality data, as measurements contain significant temporal correlations, and the values of different pollutants and environmental variables are dependent on each other. Therefore, calibration models that avoid making strong assumptions about the distribution of measurements are likely to perform and generalize better.

Interpretability. The most powerful machine learning models often are black boxes that hide their internal logic from the user [51, 77]. While a black-box model may produce more accurate estimates, its nature makes it difficult or impossible for the user to understand how the model works or why it produces a particular estimate. If the calibrated measurements will be used to make decisions that have, e.g., economic, legal, or safety consequences, then black-box approaches may be unacceptable. Another concern with black-box models is that errors in the calibration models they capture may be difficult to rectify. Indeed, a complex learning algorithm may, e.g., use unexpected correlations and features in the training data to obtain the estimates, which may lead to unexpected and counter-intuitive calibration behavior in real-life applications. For a thorough discussion on black-box models and issues with them, we refer to Guidotti et al. [51]. The alternative to black boxes is to use white boxes, models of which the performance can be easily understood and explained. In practice, however, the number of parameters influences the interpretability of models, and most common white-box models become gray boxes that can be only partially explained.

Optimization criteria. Most common machine learning algorithms operate by minimizing an objective function that specifies a *loss* between the output of the machine learning model and the desired output. In the case of calibration, the loss function measures differences between the low-cost sensor and the reference station. Generally, the cost function needs to be chosen so that it represents the goal of the application. For example, if we wish to know whether the measured concentration is of the correct order of magnitude, then mean squared error would be a good choice as a cost function, as it penalizes large errors more than small ones. However, if we wish to optimize the median performance of the algorithm, then we could use mean absolute error instead. In some cases, we may be interested in other kinds of objective functions. For example, for detecting drastic short bursts in the concentration of a pollutant, the objective function can be mapped into classification error, where the different classes represent the severity of the pollutant.

Amount of Data. Choosing the right machine learning model for a calibration problem depends on the available amount of air quality measurements. The amount of air quality measurements, in turn, is linked with the generality and complexity of the model. If the model is simple, then it requires relatively few data points for training, but it may not fit the data well. This phenomenon is known as *underfitting*. Conversely, if the model is complex, then it will approximate the function to predict the data better, but it will also need a larger training dataset to avoid *overfitting*, i.e., fitting the training dataset well but having poor performance on unseen data. Generally, the more complex the model, the longer the training dataset should be. Training data should also be sufficiently long to ensure that the model has enough information to learn how the cross-sensitivities between the sensors affect their response. A longer period of training data generally implies a longer training time, which results in fewer data being available for validating and testing the model.

The choice of test data is critical for ensuring the usefulness of the calibration model. Standard machine learning evaluation techniques, such as selecting a subset of all data as test data, or using cross-validation, are not suitable for air quality calibration due to the nature of the measurements. Indeed, these evaluation techniques can result in significant correlations between training and testing data, which would result in the calibration model overfitting on the temporal structure of the measurements. Optimal evaluation of air quality calibration models is currently an open issue as the data should cover a sufficiently long period, different pollutant concentrations, and different environmental conditions. From the studies surveyed for this article, it is difficult to estimate average lengths of training or testing periods, since many studies do not report them exactly. In the studies that report the length, the length of training data spans from 2.1 days to 8 months, whereas the length of the test data spans from 36 hours to 1.25 years.

Computational Time and Complexity. Recent developments with networking (e.g., 5G and edge computing), artificial intelligence, and sensor technology are paving the way to ever-increasing scales for deployments, even for massive-scale deployments that integrate thousands or tens of thousands of sensors [96]. Unlike in traditional sensor network deployments, which rely on tens or hundreds of low-cost sensors and sparse reference station data, in massive-scale deployments, performing the calibration at remote infrastructure becomes difficult, especially if real-time (or near real-time) information is required. For such deployments, the calibration model needs to be sufficiently lightweight in terms of computational time so that a pre-trained model can be used to adjust the measurements directly on the sensor device. This is particularly important for battery-powered monitoring stations, such as those carried by citizens, where any saving from the calibration and network costs offsets the energy cost of the air quality sensors themselves. Note that, even in this case, some of the measurements need to be transmitted to a remote server, edge deployment, or the cloud to support re-training the calibration model. Even if the scale of the deployment is smaller, time requirements of machine learning techniques also have some influence on the choice of machine learning models for calibration, as they influence the resolution of measurements and potentially the energy drain of the sensing units running the calibration models. For fixed deployments where the low-cost sensors have mains power access, computational time and complexity are less of an issue as there is no need to adjust network or CPU power and calibration can be performed in a remote location.

Most machine learning models are slow to train and fast to use while predicting the values of new measurements. The time requirements of the prediction phase should be sufficiently low so that the calibration model can be used to correct any new measurements from the low-cost sensing unit. In practice, any machine learning model can achieve this, since most low-cost air quality sensors require sufficient air intake before they can take measurements. Indeed, the temporal resolution of measurements is usually in minutes (or once per minute) rather than in seconds. For prediction, a more critical concern is memory and storage complexity. Traditional machine learning techniques, such as linear models, random forests, and support vector regression, have a reasonably small model size, but emerging techniques, such as deep learning, often have large a model size, incorporating tens or even hundreds of thousands of parameters. Storing and loading such models on the sensing units may become a bottleneck, especially on sensing units that have been designed to operate for a long time. For model training, time complexity influences how often re-calibration can be performed, as well as the overall system architecture. Simpler models, such as linear models or support vector regression, can be efficiently trained even on low-cost sensing units. However, more complex models, such as artificial neural networks or deep models, can be too computationally heavy for the sensing units. When training cannot be performed on the sensing units, sufficient computing and networking infrastructure need to be available.

6.2 Linear Models

Linear regression (LR) is the simplest machine learning regression model, and it is based on the linear equation. In the case of a single (input) feature, this model is usually referred to as univariate LR or just Linear regression (LR). In the case of more than one feature, this model is usually referred to as **multivariate linear regression (MLR)**. LR models have widely been used as calibration methods for air quality monitoring [54, 75, 114], or as a baseline for comparing the calibration performances of more complex approaches [120, 121, 154]. LR has also been used as a pre-calibration method, with the output of the model fed to other methods [30]. As discussed earlier, low-cost sensors are vulnerable to cross-sensitivities and meteorological conditions, which renders simple LR models insufficient for most environments [120, 121, 154]. MLR has shown improvement in performance [30, 82, 114, 120, 121], because the model can learn cross-sensitivities between different pollutants and meteorological conditions.

When the relationship between the features and the target pollutant is not strictly linear, data transformation can be used to improve the accuracy of LR models. For instance, Liu et al. [80] argued that LR using a logarithmic function reacts better to the cross-sensitivity between $PM_{2.5}$ and wind interference. A variation of MLR called **geometric mean regression (GMR)** has been tested by Saukh et al. [114]. Their comparison with a conventional multivariate linear regression (MLR) model shows that GMR is less vulnerable to concept drift than conventional MLR models.

Conventional linear models assume input data to be independent and identically distributed. As this assumption rarely holds with air pollutants, linear models can only reach modest calibration performance. However, a major benefit of linear models is that they are easily interpretable, since the model corresponds to a (hyper)plane fitted on data. Hence the parameters have an intuitive geometric interpretation and linear models can be considered as white-box models. Due to their simplicity, linear models require less training data than more complex models. However, to ensure good performance, the training data need to be sufficiently representative of the actual distribution of pollutants. For example, the temporal range needs to be sufficiently long to capture potential effects resulting from seasonal variation, differing weather, or other factors. Furthermore, linear models are very quick, both to train and to predict with. This holds also for air quality data as the optimal weights that minimize the cost function can be directly computed, instead of being approximated through iterations, as the number of features is small.

Advantages: Simple to define, and trivial to find the optimum weights.

Disadvantages: The low-cost **air quality sensor (AQS)** calibration problem is too complex for this model [120, 121, 154]. LR cannot automatically find all the cross-sensitivities between the various pollutants, and in some cases, the function that models some relations is not linear. In such cases, the features need to be manually transformed to allow LR to fit them.

6.3 Support Vector Regression

Support vector regression (SVR) is one of the most popular techniques for modeling *non-linear* relationships between input features and the output variable (i.e., calibrated air pollutant). The general idea in SVR is to find the hyperplane that represents the minimum distance between itself and the data points. SVR is well suited for non-linear data, as it uses a so-called kernel function to map non-linear data to a higher-dimensional space, in which the model can then fit a linear hyperplane. In air quality calibration, SVR has mainly been used for comparison against other techniques. For example, Cordero et al. [30] compare SVR, random forest, and **Artificial neural networks (ANNs)** for calibrating NO_2 measurements. The authors report mixed results, with SVR outperforming the other models in some cases but having worse performance in other cases.

SVR models are more complex than linear models, hence, they require more training data. In terms of runtime, they are fast to compute, even if less efficient than linear models. SVR, like linear models, assumes that the data are independent and identically distributed. The interpretability of SVR models depends on the kernel used, with a linear kernel resulting in a white-box model, but non-linear kernels effectively turning the model into a black-box that is difficult to interpret.

Advantages: Model training is defined as a convex optimization problem, for which there are efficient solutions to find the optimal parameters. SVR uses a kernel to transform the input data, enabling capturing nonlinear relationships between input features and calibrated pollutants. SVR incorporates a regularization parameter, which makes it possible to control under and overfitting. There are many efficient, mature, and easy-to-use SVR implementations.

Disadvantages: Sensitive to kernel parameters and assumes data to be independently and identically distributed. Poor interpretability for nonlinear kernels.

6.4 Decision Trees and Random Forests

Decision trees (DTs) are another common off-the-shelf machine learning technique. In a DT, every node of the tree has a conditional check, and each branch corresponds to an outcome of the check. While determining the value of a new measurement, we progress through the tree, starting from the root node and following the branches, until we reach a leaf node. The value of the leaf node is then used as the outcome of the calibration. Each check in the DT corresponds to a rule that can be used to subdivide the measurements into an increasingly smaller range of values. Thus, unlike linear models and SVR, DT do not fit a model or function to the data, hence they can support both linear and nonlinear relationships. To the best of our knowledge, DT have found limited use in air quality calibration. They have only been used for calibrating CO , and even then only as a baseline for other methods [60].

When applied to complex problems, DT can become highly complex and overfit on the training data. These issues can be mitigated using **random forests (RFs)**, which combine multiple simple DT into a single powerful model. Random forest is an example of the *bagging* technique, which operates by creating different subsets of the training data, learning separate models for each subset, and aggregating the outputs of all these models together while predicting unseen data. The main disadvantage of RF models is that they can be difficult to interpret. Indeed, while the outputs of individual DT can be easily interpreted, understanding the joint effect of tens or even hundreds of DT is much more complex. Similarly, the number of DT to integrate into the model influences overall runtime and performance. Having a small number of DT as part of the RF model is efficient to learn; however, it can result in the model underfitting the data, whereas integrating a large number of trees increases the model size and training data requirement. Random forest models have been widely applied to the calibration of sensors for many pollutants, such as CO [16, 154], CO_2 [154], NO_2 and O_3 [16, 30, 154], and NO , SO_2 , $PM_{2.5}$, and PM_{10} [16]. Lin et al. [76] propose a hybrid model that combines LR and RF to simultaneously learn linear and non-linear relationships.

Advantages: Reduces overfitting by training different models on different artificial datasets generated from the original dataset. The training of different models can be parallelized. Decision trees and RF do not require choosing the function for non-linear problems, which is an advantage in our scenario.

Disadvantages: Decision trees, when used for regression, can have a very high depth unless properly regularized. The number of models, and therefore of generated datasets adds additional complexity to training the parameters correctly. The computational needs of

the RF are higher than a single DT. Random forest can be considered gray-box or even black-box algorithms, depending on the number of models that are integrated.

6.5 Boosting Algorithms

Similarly to bagging, *boosting* algorithms operate by creating subsets of the training data and learning a different model on each subset. The difference to bagging is that boosting trains all the weak models sequentially, aggregating them into a single strong model, instead of running each model separately and aggregating their outputs. Boosting, when training a new weak model, also takes into account the success of the previously trained weak model, and weights the training data accordingly. Example of boosting algorithms are adaptive boosting [45], gradient boosting [46], and extreme gradient boosting [23]. In the context of air quality monitoring, boosting algorithms have been used for predicting $PM_{2.5}$ [65].

Advantages: Reduces overfitting by training different models on different artificial datasets generated from the original dataset. No need to choose the function for non-linear relationships. When training the new weak learners to combine into the strong learner, the model takes into account the success of the previous weak learner and weighs data accordingly.

Disadvantages: Interpretability and complexity depend on the number of learners that are aggregated together, with a higher number resulting in better performance, but higher complexity and reduced interpretability. Risk of overfitting when the number of learners to aggregate grows large.

6.6 Artificial Neural Networks

ANNs are a popular machine learning technique for modeling time series data. ANNs consist of a set of nodes called *neurons* grouped into *layers*; their structure has been initially motivated by the structure of the human brain [62]. The first layer of the model is called the input layer, and the last is referred to as the output layer. Intermediate layers between the input and output layers are referred to as hidden layers. Depending on the number of intermediate layers, the network is called either *shallow* or *deep*. For the calibration of low-cost air quality models, the most common approach is to use a shallow **feedforward neural network (FFNN)**, with most models containing one or two hidden layers. This type of ANN have been applied to the calibration of sensors for measuring NMHCs [36], $PM_{2.5}$ [26, 48], CO [37], CO_2 [85], NO [30, 37, 43], NO_x [43], NO_2 [30, 37, 43], and O_3 [30, 85]. Spinelle et al. [120, 121] use an ANN model based on FFNNs applied to the calibration of CO, CO_2 , NO, NO_2 , and O_3 . Lee et al. [71] design a combination calibration process selecting the prevailing calibration models between LR and ANN based on those two models' distribution of residuals.

In an ANN, each neuron contains a function, named *activation function*, which is applied to the data that it receives as input, to produce an output value. Common activation functions include the rectified linear unit function, the sigmoid function, and the hyperbolic tangent function. Similarly to LR, the activation functions consist of some weights that need to be set so that the final output of the model is as close as possible to the ground truth.

Just a few of the studies that we survey discuss the used activation function of the neurons. Some report using the hyperbolic tangent function, or a variation of it [36, 37]. Spinelle et al. [120, 121] use different activation functions in different layers in their ANN model. They also report trying the radial basis function but discarding it, because it did not yield good results.

In terms of performance, there have been studies that compare ANNs against regression models [30, 85, 120, 121]. The conclusions from these studies are mixed. Some studies report a better

performance from ANN [85, 120, 121], whereas some studies report the opposite [30]. This might be explained by the fact that the relationship between the pollutants and how they vary in time is a very complex function. Therefore, the distribution of pollutants and characteristics of environmental variables can influence whether LR and MLR models are sufficient for approximating the underlying relationships in data.

The complexity of the relationships that can be captured with ANNs depends on the structure of the network. Standard FFNN structures mostly have been designed for capturing a non-linear function between input and output. However, more complex structures that can incorporate additional considerations as part of the network, such as temporal structure or feature learning, have been proposed. As an example, **recurrent neural networks (RNNs)** are a class of ANN designed for incorporating temporal structure and thus well suited for modeling time series. In an RNN, input values or neuron outputs of the previous time step can influence the state of the ANN in the current time step. For the calibration of low-cost air quality measurements, Esposito et al. [43] tested two more complex RNN architectures: **time delay neural network (TDNN)** and **nonlinear autoregressive exogenous model (NARX)**. They report that NARX have a better performance than FFNNs, and TDNN achieve the best performance, which might be explained by the fact that RNN take into consideration previous time steps, therefore encoding input value changes over time.

It is also possible to build hybrid models by using layers of different architectures. For example, we can use an RNN layer to learn relationships between different points in the input data, an FFNN layer to learn another level of abstraction, and a final FFNN layer to constrain the size of the output. These hybrid models are called Deep learning models. The idea behind this is to create a model that can learn multiple levels of abstraction of the input data [70].

Advantages: Very flexible learner that can approximate any function, given enough layers and neurons. It can automatically learn the relationships between the features and learn multiple response values simultaneously.

Disadvantages: Its complexity makes it extremely expensive in computing resources to train, requiring dedicated hardware for doing so quickly. It also needs a huge amount of data for avoiding overfitting. ANN are typically black-box algorithms.

6.7 Gaussian Processes

A **Gaussian process (GP)** is a model that combines multiple Gaussian random variables into a joint distribution to estimate the function that models the data. It is a non-parametric approach, which means that there is no need to specify the number of parameters. However, similarly to SVR, GPs need a kernel function to be specified. Detailed information about GPs and how to train them can be found in Rasmussen et al. [108].

Gaussian processes make weak assumptions on the distribution of input data, and hence are well suited for calibrating low-cost air quality sensors. Another benefit of GPs is that it is possible to plot the probability distributions used to model data. Gaussian processes may require a training dataset that is larger than linear models and **support vector machines (SVMs)** would require, but generally smaller than more complex models such as ANN. Gaussian processes are lazy learners, meaning they do not need to be trained, but instead, they approximate the function of the training dataset while predicting. However, this means that the whole training dataset needs to be kept in memory, and every time that it needs to predict the target variable, it needs to compare it to the probability distribution of the features. This means that GPs have high memory complexity, which might be impractical if the deployed AQS units do not communicate with a central infrastructure and need to run calibration in the field. Gaussian processes have been used in the context of low-cost AQS calibration by Cheng et al. [26] and Gao et al. [48] on top of FFNN to improve

its performance. In both studies, GPs were able to improve calibration performance compared to a stand-alone FFNN.

Advantages: Non-parametric model, no need to specify the number of parameters except for the kernel. Makes only weak assumptions on the distribution of data. Gives probability distribution for the predictions.

Disadvantages: High memory complexity, requires a kernel function and can be sensitive to parameters of the kernel function.

6.8 Other Machine Learning Paradigms

Most of the early work on air quality sensor calibration focused on validating and testing traditional machine learning algorithms, as discussed in the previous sections. Recently the focus has shifted towards more advanced paradigms that address specific challenges in the development of calibration systems. Below we briefly discuss some of the most important ones.

Transfer learning. It may happen that a model trained in one domain (e.g., at a specific time interval, at a specific location, or for specific sensor hardware) is not accurate in another domain and the machine learning model has to be trained separately for each domain. The idea in *transfer learning* is to mitigate this problem and to make it possible to use commonalities between the different domains to train better and more accurate models, with a smaller amount of data needed per domain [149, 153]. Transfer learning can transform a pre-trained model so that it can be applied to a different domain. While parameters of the model need to be tuned for the new environment, the technique allows adapting previous solutions to new problems using existing models and data.

Speeding up Training. Several machine learning paradigms are closely related to transfer learning, but instead of attempting to generalize to other domains, focus on making the learning process easier. First, *meta learning*, also known as *learning to learn*, focuses on systematically observing how machine learning algorithms learn on different domains and using this knowledge to speed up the learning process on a new domain [129]. Another related paradigm is *few-shot learning*, where a collection of techniques is used to generalize to new tasks with only a small amount of labeled training samples [139]. *Semi-supervised learning* focuses on using a limited amount of labeled ground-truth measurements to support training [128]. Semi-supervised learning algorithms utilize the structures found in the measurements with and without the ground truth to make better models than would be possible with the ground-truth measurements alone.

Incremental Learning. Calibration models suffer accuracy loss in long-term field deployments due to many negative effects, such as sensor drift, changes in the probability distribution of priors due to seasonal changes, and so on. Hence, it is important to have a more accurate model by updating the original model, taking into consideration new measurements and changes. Incremental learning is a continuous learning approach where the learning process takes place whenever new samples emerge, and it adjusts what has been learned according to the new samples. Continuous learning, online learning, and adaptive learning are quite similar machine learning paradigms to incremental learning. De Vito et al. [34] show that adaptive and incremental strategies, such as performing periodic recalibration, can improve the performance of initially trained calibration models for low-cost sensors. Through experiments carried out on measurements collected from 18-months electrochemical sensors deployments, monitoring CO, NO₂, and O₃ concentrations, the authors demonstrate that such strategies improve the overall performance of calibration models and make them less sensitive to seasonal changes or other variations.

Dimensionality Reduction. When a large number of variables are used for learning the calibration function, the complexity of the measurements and the resulting models can be reduced using dimensionality reduction techniques. Dimensionality reduction is a special case of *unsupervised learning*, and can be implemented using methods such as cluster analysis [57].

6.9 Network Calibration Strategies

In practice, individual sensors rarely have continuous access to reference measurements or even to the same source of reference measurements [96]. *Network calibration* refers to the process of post-deployment calibration, which can either focus on establishing a calibration function for a new sensor with the help of existing sensors or re-training the model when concept drift occurs. Below we briefly cover the main techniques for network calibration, and we refer to the survey by Maag et al. [84] for a more detailed discussion on network (re-)calibration strategies.

Blind Calibration. This technique attempts to adapt calibration models so that there is a high similarity between measurements across the entire measurement network. Blind calibration typically assumes that the deployment is dense and neighboring sensors have closely matching or at least highly correlated measurements. Balzano et al. [8] alleviate the deployment density requirement but require measurements to be correlated across the network. Unfortunately, this is not the situation in the air quality monitoring system, as we have mentioned earlier that the measurements of sensors cannot be guaranteed to be identical, and even a single city block can witness significant variations in pollutant concentrations. Tsujita et al. [127] developed a gas sensor system that incorporates an auto-calibration method. They developed a case study of NO_2 distribution, performed by continuously installing a low-cost NO_2 sensor in different locations where there is no professional-grade reference station in the vicinity. Sensors were then calibrated to the reference station when the NO_2 concentrations were low and expected to have almost identical values in deployed areas. Such approach is only applicable in areas where the pollution level regularly reaches a low level. Blind calibration usually operates under the assumption that the distribution of the concentration of pollutants is uniform in a certain region [93, 97], and is only useful for offset and gain calibration [24]. Miskell et al. [92] designed a hierarchical network framework incorporating a proxy model, a measurement model, and a semi-blind calibration model. They adjust the gain and offset of a low-cost sensor's measurements to the first two moments of the probability distribution of the nearest reference station, and use the adjusted data as a proxy for most sites. The developed algorithms achieve great success in detecting and correcting sensor drift, and the designed framework can deliver reliable high temporal-resolution ozone data at neighborhood scales.

Opportunistic and Collaborative Calibration. When the sensors are mobile, opportunistic encounters between devices can be used to create virtual reference points that can be utilized for calibration [142]. Saukh et al. [114] propose a PM calibration approach that uses opportunistic encounters to obtain reference sensor measurements. Their approach uses these measurements in a multi-sensor data fusion approach to learn the calibration function. Opportunistic calibration can either rely on measurements from a single device or use a *collaborative* approach, where devices share reference measurements and use them to learn the calibration function [96]. A further extension is *multi-hop calibration*, which leverages encounter-based opportunistic calibration to propagate the calibration to other sensors. Specifically, newly calibrated sensors are considered as references for devices that cannot directly access reference devices, and thus calibration parameters and measurements are recursively propagated throughout the measurement network. As an example, Maag et al. [83] propose a multi-hop calibration technique, **sensor array network calibration (SCAN)**, for dependent low-cost sensors. SCAN minimizes error accumulation over sensor arrays and has been theoretically proven to be free from regression dilution.

Calibration Transfer. This technique refers to the use of transfer learning to adapt calibration functions across domains. Calibration transfer has been mainly performed in-lab for pre-deployment calibrating electronic noses to reduce calibration overhead in mass production where each device needs to be calibrated due to inter-device difference [143, 144]. During the calibration process, a pre-trained calibration model with fine-tuning parameters is produced and later adopted on the target sensor. Cheng et al. [24] propose an in-field calibration transfer with a large-scale real-world PM monitoring deployment. The authors assume the target location holds a similar distribution of the ground truth compared to the source location, and that the required transformation is approximately linear. Cheng et al. [25] recently designed an air quality map generation scheme named MapTransfer. Their main idea is to enlarge the current sensor measurements in a downscaled sparse deployment with suitable historical data from a shorter duration dense deployment. A learning-based data selection scheme is adopted to select the best matching data, and a multi-output Gaussian process model is used for fusing the best-selected data with the current measurement. In the experiment, they consider data spanning the whole year 2018 from 200 sensors as the dense deployment, and data spanning half of the year 2019 from 50 randomly selected sensors out of 220 as the sparse deployment. These calibration techniques are very powerful for dealing with post-deployment scenarios, where irregular or even no access to reference measurements can be a common issue.

7 MEASURING PERFORMANCE AND COMPARING MODELS

Machine-learning-based calibration models are only useful if they can consistently improve the accuracy of the sensor measurements produced by a low-cost sensing unit. To ensure this indeed is the case, calibration models need to be validated against measurements collected from high-cost reference stations. Next, we discuss validation methods and how they can be applied to low-cost calibration, and compare existing low-cost air quality calibration studies by selecting the most commonly used performance measures in them.

7.1 Performance Measures

The performance of calibration models is typically expressed through one or more performance measures, which are functions that characterize the dissimilarity between the output of the calibration model and the ground-truth values obtained from a reference station. Existing studies have used differing performance measures, which makes it difficult to compare performance across studies. In the following, we briefly summarize the main performance measures.

Absolute Error Measures. Mean square error (MSE) is the standard error measure for assessing the performance of regression models. MSE is defined as:

$$\text{MSE} = \frac{1}{m} \sum_{i=1}^m (y_i - \hat{y}_i)^2, \quad (1)$$

where m is the number of samples, $\hat{y}_i$ is the predicted value and y_i is the actual value of a sample. MSE is useful to evaluate the performance during training and to define a cost function because of its simplicity. A variation of MSE is **root-mean-square error (RMSE)**, defined as:

$$\text{RMSE} = \sqrt{\frac{1}{m} \sum_{i=1}^m (y_i - \hat{y}_i)^2}. \quad (2)$$

Another absolute error measure is **mean absolute error (MAE)**, which is defined as:

$$\text{MAE} = \frac{1}{m} \sum_{i=1}^m |y_i - \hat{y}_i|. \quad (3)$$

Finally, **mean bias error (MBE)** is defined as:

$$\text{MBE} = \frac{1}{m} \sum_{i=1}^m (y_i - \hat{y}_i). \quad (4)$$

MSE and RMSE are very similar. MSE can be interpreted geometrically as the average fit of points to a regression model, whereas RMSE is the average distance of points from the regression model. RMSE and MSE weigh errors proportionally to their magnitude, whereas MAE weighs all errors equally. This makes RMSE and MSE more sensitive to outliers [21], suggesting that MAE is a better measure for measuring the average performance, whereas (R)MSE is useful for measuring a model's sensitivity to outliers. In practice, it is recommended to use both measures together as this provides complementary information on the model's performance [21]. MBE, however, measures whether the average error is positive or negative and can be used to determine whether the model underestimates or overestimates the pollutant values. Spinelle et al. [120, 121] divide RMSE and MBE by the standard deviation of the reference measurements and combine the resulting values into a target diagram. Target diagrams are useful for visualizing these two performance metrics and for quickly comparing different models.

Relative Error Measures. The alternative to absolute measures is to rely upon a relative error measure that expresses the error proportionally to the true measurements. The most popular relative error measure is **mean relative error (MRE)**, which is defined as:

$$\text{MRE} = \frac{1}{m} \sum_{i=1}^m \left| \frac{y_i - \hat{y}_i}{y_i} \right|. \quad (5)$$

A related measure is **mean absolute percentage error (MAPE)**, which expresses MRE as a percentage:

$$\text{MAPE} = \frac{1}{m} \sum_{i=1}^m \left| \frac{y_i - \hat{y}_i}{y_i} \right| \cdot 100\%. \quad (6)$$

MRE is useful for expressing how far estimated values are from the reference values, whereas MAPE is useful for characterizing performance when the same model is applied for multiple pollutants.

Coefficient of Determination. The coefficient of determination, or R^2 , measures how much a variable influences another variable. For a calibration model, R^2 measures the percentage of variance that the model explains. To compute R^2 we need to compute two variability measures, namely the total sum of squares:

$$\text{SS}_{\text{tot}} = \frac{1}{m} \sum_{i=1}^m (y_i - \bar{y})^2 \quad (7)$$

and the sum of squares of residuals:

$$\text{SS}_{\text{res}} = \frac{1}{m} \sum_{i=1}^m (y_i - \hat{y}_i)^2 = \text{MSE}. \quad (8)$$

Here $\bar{y}$ is the mean of the target data. The R^2 is now given as

$$R^2 = 1 - \frac{SS_{\text{res}}}{SS_{\text{tot}}}. \quad (9)$$

R^2 can be useful in the low-cost AQS calibration scenario to assess how closely the distribution of the predicted values matches the distribution of the ground-truth measurements.

Uncertainty Measures. Air quality standards typically associate bounds on the uncertainty that the measurements can contain. From an algorithmic point-of-view, uncertainty can be considered as a measure of *robustness* as it provides insights into the operating bounds of the calibration framework. As an example, the Clean Air directive of the European Union assigns maximum deviations for the 95% uncertainty of measurements, with the precise deviation depending on the pollutant [31]. Uncertainty is typically measured using the standard deviation of measurements, which specifies the *standard uncertainty* of the measurements. Alternatively, the *relative standard uncertainty* specifies measurement uncertainty relative to the magnitude of the measurements. The standard deviation of measurements can be interpreted as a confidence interval of 68%, which leaves a broad margin for the measurements to deviate. A tighter bound can be obtained using *expanded uncertainty*, which is defined as the standard deviation multiplied by a *coverage factor* k , which determines the bounds of the uncertainty region. A coverage factor of $k = 2$ roughly translates into a confidence interval of 95% whereas a coverage factor of $k = 3$ corresponds to a confidence interval of 99%. Similarly to the standard uncertainty, *expanded relative uncertainty* is a measurement of uncertainty defined relative to the magnitude of the measurements. In practice, the robustness of the evaluation setup of air quality models is measured using expanded relative uncertainty, as that allows to compare uncertainty across areas with differing pollutant concentrations.

Best Practices. MSE, RMSE, and R^2 are closely related. This means that if we rank some models according to one or the other measure separately, then the ranking positions for the models in both rankings will be the same. The same holds for MRE and MAPE. However, measures that are not directly related, such as RMSE, MAE, and MRE, do not necessarily result in the same ranking. Hence, in practice, the recommended approach is to use multiple performance measures and take into account how they are affected by the properties of the data. Visual aids, such as target diagrams, should also be used so that different performance measures can be visually compared. A summary of the performance measures can be found in Table 5.

7.2 Evaluation Criteria for Low-Cost Deployments

To evaluate the suitability of air quality sensor calibration methods considered in this survey for long-term deployments of low-cost sensors, we classify them based on four evaluation criteria.

- The robustness of the evaluation setup of the method, based on the length of the test dataset
- The resolution, determined by the length of the smallest temporal step modeled
- The accuracy of the method as reported by the authors
- The scalability of the sensor technology, emphasizing technologies are low-cost and that can operate independently for long periods

Each method is classified as *Low*, *Medium*, or *High* with respect to each of the four criteria. Methods that could not be classified based on the information available on it are marked with N/A. The rest of this section explains how the criteria are determined for a given method.

Robustness. This criterion is based on the variability of the measurements, and it is indirectly linked to the length of the test dataset. Longer test datasets have a higher probability to include

Table 5. Comparison of Common Performance Measures

Method	Metric	Formula	Advantages	Disadvantages
MSE	No	$\frac{1}{m} \sum_{i=1}^m (y_i - \hat{y}_i)^2$	Simple measure, can be used as a cost function. Useful for measuring the model's sensitivity to outliers.	Tends to exaggerate errors, especially with noisy data. For very clean data, it might overestimate the model performance.
RMSE	Yes	$\sqrt{\frac{1}{m} \sum_{i=1}^m (y_i - \hat{y}_i)^2}$	Same as MSE, but in the same dimension as the target values.	Same as MSE.
MAE	Yes	$\frac{1}{m} \sum_{i=1}^m y_i - \hat{y}_i $	Useful for measuring the "average" performance of a model.	Underestimates the outliers.
MBE	No	$\frac{1}{m} \sum_{i=1}^m (y_i - \hat{y}_i)$	Useful for measuring the bias, and to see to which value the average error tends.	It takes into account the bias only. It can't be used to evaluate the actual performance of a model.
MRE	Yes	$\frac{1}{m} \sum_{i=1}^m \left \frac{y_i - \hat{y}_i}{y_i} \right $	Useful for expressing the average error in proportion to the target values.	Tends to exaggerate the error for small values, and underestimate the error for big values.
MAPE	Yes	$\frac{1}{m} \sum_{i=1}^m \left \frac{y_i - \hat{y}_i}{y_i} \right \cdot 100\%$	Same as above.	Same as above.
R^2	No	$1 - \frac{SS_{res}}{SS_{tot}}$	Useful for measuring how much the variance is accounted for by the model.	Same as MSE and RMSE.

seasonal and weekly variations, which result in a wider spectrum of pollution and environmental values in the data. This in turn ensures the evaluation setup considers the robustness of the model in the presence of such variations—and consequently can provide better estimates of uncertainty. Methods with one-month long datasets or shorter are classified *Low* in terms of Robustness, and one-year-long datasets or longer are classified as having *High* robustness. Note that optimal test data would include multiple years, so the effects of variations between years in the model can be mitigated. In practice, collecting such long-term data with an identical measurement setup is often very difficult due to practical constraints, sensor failures, and other sources of errors, such as sensitivity drift of the measurement hardware.

Resolution. This score is based on the temporal resolution of the data used in a model. Methods using data with a resolution of one hour or coarser are classified *Low* and those using data with a resolution of 1 minute or finer are classified *High*. We note that the main purpose of this classification is to consider the suitability of the technologies for near real-time services, such as estimation of current air pollution levels and calculation of the relevant air quality index value. In some applications, such as long-term assessment of pollution levels or prediction of hazardous areas, a one-hour resolution is sufficient. Hence, classification is not intended as an assessment of quality, but rather as a statement on the level of granularity that is offered.

Accuracy. This score is based on the accuracy of the model. This is the most complex to estimate, since different studies use different evaluation measures. We define it as follows. If a study uses RMSE, or as an alternative, MAE, then it is used as the base value of the performance of the models. Since Spinelle et al. [120, 121] do not provide an exact value, we use a conservative estimate of MAE obtained by analyzing the residual plots that they provide. Models of studies that do not use comparable similarity measures are marked *N/A*. We group the models by the type of pollutant they predict. In each group, the model with the best performance is classified *High*, and the model with the worst performance is classified *Low*. Other models are classified as having *Medium* Accuracy.

Scalability. This criterion is based on the sensor technologies used to produce the input data of a model. The score is determined based on the suitability of such sensor technology for large-scale deployments, taking into account the typical cost, robustness, and lifespan of such technology. DiSC sensors are classified *Low* due to their higher cost, non-dispersive infrareds (NDIRs), PIDs, and OPCs are classified *Medium*, while LSP, MOS, and EC sensors are classified *High* (higher value for scalability is indicative of the technology being a better fit for large-scale deployments). The Scalability classification of each method is computed as the mean of the classifications of the sensors used in the method. For example, a method with one MOS, one EC and one NDIR sensor is classified *High*, while a method with one MOS and one NDIR sensor is classified *Medium*.

7.3 Comparing Studies

The different studies surveyed for this article are compared and summarized in Table 6. Below we separately discuss the studies according to each of the four criteria.

Robustness. The study by Maag et al. [82] is the only one scoring high for robustness. This is because they use a dataset that spans more than a year to test the model, which is by far the longest among all studies surveyed for this article. The next highest is the study by De Vito et al. (2009) [37], which uses a test dataset spanning about 7 months. From the comparison, we can observe that most studies use relatively short test datasets, which are unlikely to capture the full extent of seasonal variations. Most datasets used in the literature are proprietary and specific to a single deployment. Taken together, these factors mean that an objective comparison of calibration methods is currently highly difficult due to short measurement periods and a lack of openly available reference datasets.

Resolution. Many studies score high in Resolution, because they construct models using data with a temporal resolution lower or equal than 1 minute. These studies are Maag et al. (2016) [82], Spinelle et al. (2015, 2017) [120, 121], Cheng et al. [26], Esposito et al. [43], Liu et al. [79], Saukh et al. [114], Hasenfratz et al. (2012) [54], and Gao et al. [48]. Other studies with a good resolution score are Lin et al. [75] and Zimmerman et al. [154]. Every other study uses data with a 1-hour resolution or does not report the used resolution. These result in Resolution scores of Low and N/A, respectively. This suggests that most studies use data with a good temporal resolution. Note that we only compared temporal resolution, and the picture would be completely different if we considered also the spatial resolution of the studies. While air quality deployments are increasingly commonplace, practically all studies only consider measurements from a single geographic location. Together with the lack of openly available datasets, this means that it is difficult to currently assess how well the methods perform across spatial variations—let alone considering the combined effect of spatial and temporal variations.

Accuracy. The RF model by Zimmerman et al. [154] has the top ranking in accuracy for most pollutants, namely CO , CO_2 , NO_2 , and O_3 . Notable mentions with good performance are the MLR model by Maag et al. [82] for CO , the FFNN model by Spinelle et al. [121] for CO_2 , the TDNN and NARX models by Esposito et al. [43] for NO_2 , and the RF model by Borrego et al. [16] for NO_2 and O_3 . Only a few models have been developed for the rest of the pollutants, and a few of the studies that present them do not present meaningful performance measures. Because of these reasons, we will only mention the best model for each. For NO , the best model is the FFNN model by Spinelle et al. (2017) [121], for NO_x , the best is the TDNN model by Esposito et al. [43], for SO_2 , the best model is the RF model by Borrego et al. (2018) [16], and for $PM_{2.5}$, the best is the FFNN model with GP by Cheng et al. [26]. It is noteworthy that the best performing model for almost all pollutants is different (with the sole exception being the FFNN model that is best for

Table 6. Comparison of the Different Machine Learning Approaches for Low-cost air quality sensors (AQSS) Calibration, Including Performance Measures and Our Evaluation Criteria

Target	Method	Reference	Features										Atmospheric features				Performance measures				Evaluation Criteria			
			CO	carbon dioxide (CO ₂)	NO	a combination of nitrogen oxide (NO) and nitrogen dioxide (NO ₂) (NO _x)	NO ₂	ozone (O ₃)	sulphur dioxide (SO ₂)	VOCs	PM _{2.5}	PM ₁₀	T	RH	AH	WS	RMSE	MAE	MAPE	R ²	Robustness	Resolution	Accuracy	Scalability
CO	MLR	[82]	x				x	x					x	x			0.048 ppm	NR	NR	NR	High	High	Medium	Medium
	RF	[154]	x	x			x	x					x	x			NR	7.9 ppb	NR	0.99	Medium	Medium	High	Medium
	FFNN	[121]	x	x			x						x		x		NR	0.05-0.12	NR	0.367	Medium	High	Medium	High
	MLR	[154]	x										x				NR	39 ppb	NR	0.94	Medium	Medium	Medium	High
	FFNN	[37]	x							x							NR	0.31 ppm	27.00 %	NR	Medium	Low	Low	High
	RF	[16]	x	x	x		x	x	x				x	x			NR	0.07 ppm	280.00 %	0.88	Low	Low	Medium	Medium
	FFNN	[16]	x	x	x		x	x	x				x	x			NR	0.09 ppm	34.00 %	0.51	Low	Low	Medium	Medium
	LR	[114]	x														0.23 ppm	NR	NR	NR	Low	High	Medium	High
	FFNN	[121]	x	x			x										NR	2-18	NR	0.787	Medium	High	Medium	Medium
	RF	[154]	x	x			x	x					x	x			NR	1.7 ppb	NR	0.99	Medium	Medium	High	Medium
CO ₂	MLR	[154]		x									x	x			NR	19 ppb	NR	0.49	Medium	Medium	Medium	Medium
	FFNN	[85]						x		x			x				NR	44 ppb	NR	0.94	Medium	N/A	Medium	High
	MLR	[85]						x		x			x				NR	135 ppb	NR	0.49	Medium	N/A	Low	High
	FFNN	[121]			x		x						x	x			NR	0.1-0.5	NR	0.208	Medium	High	High	Medium
	FFNN	[16]	x	x	x		x	x	x				x	x			NR	2.39 ppb	84.00 %	0.39	Low	Low	Medium	Medium
NO	RF	[16]	x	x	x		x	x	x				x	x			NR	4.76 ppb	127.00 %	0.74	Low	Low	Low	Medium
	MLR	[82]	x				x	x					x	x			5.13 ppb	NR	NR	NR	High	High	Medium	Medium
	RF	[154]	x	x			x	x					x	x			NR	0.5 ppb	NR	0.99	Medium	Medium	High	Medium
	FFNN	[120]	x				x	x							x		NR	2-4.5 ppb	NR	0.596	Medium	High	Medium	Medium
	FFNN	[37]	x				x	x		x							NR	10.1 ppb	22.00 %	NR	Medium	Low	Low	High
	MLR	[154]					x			x			x	x			NR	4.6 ppb	NR	0.59	Medium	Medium	Medium	High
	TDNN	[43]			x		x	x					x	x			NR	1.27 ppb	22.00 %	NR	Medium	High	Medium	High
	NARX	[43]			x		x	x					x	x			NR	1.39 ppb	21.00 %	NR	Low	High	Medium	High
	FFNN	[43]			x		x	x					x	x			NR	1.59 ppb	25.00 %	NR	Medium	High	Medium	High
	RF	[16]	x	x	x		x	x	x				x	x			NR	1.97 ppb	17.00 %	0.89	Low	Low	Medium	Medium

(Continued)

Table 6. Continued

Target	Method	Reference	Features										Atmospheric features				Performance measures				Evaluation Criteria			
			CO	carbon dioxide (CO ₂)	NO	a combination of nitrogen oxide (NO) and nitrogen dioxide (NO ₂) (NO _x)	NO ₂	ozone (O ₃)	sulphur dioxide (SO ₂)	VOCs	PM _{2.5}	PM ₁₀	T	RH	AH	WS	RMSE	MAE	MAPE	R ²	Robustness	Resolution	Accuracy	Scalability
	FFNN	[16]	x	x	x		x	x	x		x	x	x	x			NR	2.00 ppb	25.00 %	0.81	Low	Low	Medium	Medium
	MLR	[30]			x		x	x				x	x				2.92–3.88 ppb	2.34–3.14 ppb	12.3–35.4%	0.81–0.93	N/A	Low	Medium	High
	RF	[30]			x		x	x				x	x				2.29–4.73 ppb	1.86–3.35 ppb	16.0–32.9%	0.85–0.95	N/A	Low	Medium	High
	SVM	[30]			x		x	x				x	x				2.87–4.94 ppb	2.07–4.36 ppb	19.8–34.3%	0.79–0.95	N/A	Low	Medium	High
	FFNN	[30]			x		x	x				x	x				3.24–7.44 ppb	2.76–6.22 ppb	20.2–93.4%	0.62–0.88	N/A	Low	Medium	High
	LR	[75]				x											NR	NR	NR	0.88	Low	Medium	N/A	High
	TDNN	[43]			x		x	x				x	x				NR	1.37 ppb	20.00 %	NR	Medium	High	High	High
	FFNN	[43]			x		x	x				x	x				NR	1.95 ppb	29.00 %	NR	Medium	High	Low	High
	MLR	[82]	x				x	x				x	x				2.8 ppb	NR	NR	NR	High	High	Medium	Medium
	FFNN	[120]	x				x	x									NR	1–4.5 ppb	NR	0.915	Medium	High	Medium	Medium
O ₃	RF	[154]	x	x			x	x				x	x				NR	0.7 ppb	NR	0.99	Medium	Medium	High	Medium
	MLR	[154]					x	x				x	x				NR	5.1 ppb	NR	0.81	Medium	Medium	Medium	High
	LR	[54]					x	x									NR	1.46 ppb	NR	NR	Low	High	Medium	High
	FFNN	[85]					x	x		x		x					NR	3.5 ppb	NR	0.91	Medium	N/A	Medium	High
	TDNN	[43]			x		x	x				x	x				NR	7.45 ppb	42.00 %	NR	Medium	High	Medium	High
	RF	[16]	x	x	x		x	x		x		x	x				NR	1.32	20.00 %	0.97	Low	Low	Medium	Medium
	FFNN	[43]			x		x	x				x	x				NR	7.90 ppb	70.00 %	NR	Medium	High	Medium	High
	LR	[114]					x	x									12.9 ppb	NR	NR	NR	Low	High	Low	High
	FFNN	[16]	x	x	x		x	x			x		x	x			NR	2.60 ppb	18.00 %	0.86	Low	Low	Medium	Medium
	LR	[75]					x	x									NR	NR	NR	0.92	Low	Medium	N/A	High
PM _{2.5}	MLR	[85]					x			x			x				NR	10.7 ppb	NR	0.16	Medium	N/A	Medium	High
	FFNN + GP	[26]							x				x	x			96.69 ug/m ³	NR	NR	NR	Medium	High	High	High
	MLR	[79]							x						x		NR	NR	NR	0.9959	Medium	High	N/A	High
SO ₂	FFNN + GP	[48]							x				x	x			NR	NR	5.40 %	NR	N/A	High	N/A	High
	RF	[16]	x	x	x		x	x		x		x	x	x			NR	0.09 ppb	5.00 %	0.95	Low	Low	High	Medium
	FFNN	[16]	x	x	x		x	x		x		x	x	x			NR	0.16 ppb	10.00 %	0.86	Low	Low	Low	Medium

For each pollutant, rows are ranked using our evaluation criteria-based final ranking in descending order.

two pollutants). Whether this is due to a lack of sufficiently generalizable models or differences in pollutant characteristics is currently an open issue that warrants further investigation. Indeed, the target and input variables of the models tend to vary across studies, which makes it difficult to draw any generalizable conclusions about the types of models that are needed to support different pollutants. The comparison is further exacerbated by the fact that the sensor technology used in each study tends to differ from the other studies.

Scalability. The Scalability score evaluates the sensor technologies used in the studies. The studies with the highest score are those by De Vito et al. [37], Hasenfratz et al. [54], Lin et al. [75], Saukh et al. [114], Maag et al. [85] and Zimmerman et al. [154]. All of these studies have in common the fact that they use MOS sensors, which, along with EC sensors, seem to be the easiest and cheapest to integrate into large-scale deployments. Studies with EC sensors were also ranked *High*. Other studies with a good score are Cheng et al. [26], Gao et al. [48], and Liu et al. [79], which all use LSP sensors. The rest of the studies have a score that ranges from middle to low, since they use a combination of sensors including also high-cost sensors, such as Zimmerman et al. [154], which uses NDIR. Studies with a higher than average score are Spinelle et al. [120, 121], and Zimmerman et al. [154]. Scalability is an essential metric for assessing the suitability of technologies for large-scale deployments and a useful tool for fostering the adoption of denser sensor deployments. Currently comparing scalability is difficult due to a lack of suitable points of reference. The cost of low-cost sensors will always be several orders of magnitude lower than that of professional reference stations, but a combination of both is required to ensure the low-cost sensors can be calibrated. To measure scalability, suitable economic and geographic references are needed. As an example of this, Motlagh et al. [96] estimated a cost for a dense deployment of sensors in central Beijing. Having such geographic estimates that can relate resolution of monitoring, cost of deployment and improvements in accuracy, would improve comparison of sensor units and their suitability for large-scale deployments. Naturally, these reference points should be sufficiently varied, as population density is another factor that guides the requirements for the density of deployments.

8 DISCUSSION AND ROADMAP

In this survey, we have critically compared common technologies and methods for machine-learning-based calibration of low-cost sensing units, including the sensing units themselves, machine learning algorithms used for constructing the calibration models, and evaluation measures for assessing the usefulness of the models. We next reflect on the current state of the field, highlighting open issues that need addressing, and briefly presenting some directions for future research. As a basis for identifying open issues, we use the comparison of previous studies shown in Table 6.

Combination of Sensors. Considering that sensors have cross-sensitivities between pollutants, an important research problem is to find the best combinations of sensors that capture as many pollutants as possible without suffering in calibration performance. As also shown in Table 6, most studies currently focus either on calibrating gas sensors or particular matter counters, instead of attempting to calibrate both types of pollutants simultaneously. Capturing actionable pollution information requires accurate measurements for all pollutants included in relevant air quality indexes. A related topic is finding the best combinations of sensors so that they can complement each other, e.g., combining infrared and laser-based LSP sensors in the same sensing unit, to optimize energy consumption and accuracy while offsetting each others' disadvantages. Indeed, most sensing units currently encapsulate a single sensor per pollutant instead of combining different types of sensors for the same pollutant.

Lifecycle Management. Massive-scale AQSS deployments are often built out of a heterogeneous base of sensors that are unattended and installed in hard-to-reach locations. Routine tasks such as cleanup or software updates become hard to manage, which can lead to high maintenance costs. Device lifecycle management with minimal need for manual intervention is critical for continuous long-term operation of these deployments. Achieving this with massive deployments is an open issue. Another open problem related to lifecycle management is detecting or predicting when a sensor has failed or is about to fail. This can be potentially accomplished using ML techniques, but these techniques have not yet been investigated in the context of low-cost AQSSs.

Mobility Effects. Mobility can significantly affect the accuracy of sensors. When a sensor is in movement, the quantity of air that enters the sensor increases proportionally to its speed of travel, which in turn can increase the concentration of the pollutant detected by the sensor. As we have already discussed, there are some ways to measure movement speed, so that it can be taken into account on the accuracy.

Universal Models. Most studies that attempt to calibrate multiple pollutants report mixed performance, with the best model differing for different pollutants. The most evident example can be found in Cordero et al. [30], where no single model is the best performing for all gases and PM classes. Developing models that perform well for multiple pollutants is currently an open issue.

Deep Learning. Little work has been done on applying deep learning for calibrating low-cost AQSSs, with the only works we are aware of being those of Yu et al. [146, 147]. While these have shown encouraging results, the lack of openly accessible standard test datasets makes it difficult to fully assess the benefits of deep learning in the context of air quality calibration. In particular, the many complexities of air quality measurements, including drift, autocorrelation, cross-sensitivities, and other sources of uncertainty, mean that the risk of overfitting is high. Indeed, deep learning typically requires a large number of measurements before the model converges, and in heterogeneous environments can easily overfit without this being easy to identify [104]. Another challenge with deep learning is to determine how to optimally represent and model the data. The works of Yu et al. [146, 147] rely on so-called sequence-to-point modeling where the idea is to use a sliding window over input measurements to derive an aggregate point estimate for the window. These models can improve overall performance, but do so at the cost of temporal resolution and are best suited for deployments having frequent sampling rates (e.g., minute or even less). Understanding these kinds of trade-offs in modeling is another important research direction.

Dataset Length. The studies by Maag et al. and De Vito et al. [34, 82, 85], to the best of our knowledge, are the only ones that use a test dataset longer than a year. This is important for capturing seasonal phenomena and the effects of different weather patterns. In the future, it would be important to see more studies that test the models for periods longer than a year, so that they are tested on different conditions.

Concept Drift and Re-Calibration. Concept drift has been widely reported as an issue for low-cost sensor technologies. However, most studies have not been able to assess its effect due to using only short measurement periods. Concept drift can be mitigated using periodic re-calibration or online training, where new training data are continuously used to improve model performance. Developing methods for detecting drift and triggering these mitigation techniques are currently open issues.

Other Machine Learning Paradigms. Machine learning paradigms, such as transfer learning or semi-supervised learning, could be readily utilized in sensor calibration. Transfer learning would

be able to, e.g., speed up learning on new sensor hardware by utilizing commonalities with other domains, and semi-supervised learning could utilize the potentially large number of measurements for which calibrated ground truth does not exist.

Use of Multiple Performance Measures. Most studies use only one or two performance measures for evaluating the performance of calibration models. Further work is needed to assess calibration accuracy using multiple, complementary measurements. Also, the practical results of calibration have not been thoroughly assessed, with the study by Cheng et al. [26] being the only one to consider calibration performance in practical applications. Specifically, they consider how the results of the calibration model would affect the values of an air quality index.

Virtual Sensor. The dominant approach for air quality monitoring relies on low-cost sensors, a reference station, and a calibration model. Another possibility is to combine calibration with *virtual sensors* where so-called proxy variables—variables that are correlated with the target variable—are used to construct a model that can be used to estimate the value of the target pollutant. For example, we can measure BC concentration using low-cost PM sensors by building a model that can learn the relationship between PM and BC to estimate BC concentration from PM concentration [148]. Using virtual sensors and proxy variables can reduce the required number of sensors, as fewer specialized sensors are needed, allowing for denser deployments or lower costs, or both.

5G and Edge Computing. High-resolution spatial-temporal air quality monitoring requires a dense deployment of low-cost sensors in an urban area, which brings many challenges for providing system support, such as ubiquitous high-speed connectivity and real-time analysis. Emerging 5G networks and edge computing deployments provide many of the necessary mechanisms, e.g., by offering faster data rates, energy-efficient networking, and having support for computing and data storage closer to the sensors. As an example of the potential benefits, Das et al. [33] report a 10-fold reduction in the energy costs of network transmissions when using 5G. These improvements offer opportunities for new types of air quality monitoring solutions. For example, hyperspectral remote sensing has recently been used to estimate air quality [59]. The improvements in networking technology offered by 5G and edge computing make it possible to deploy and use low-cost hyperspectral cameras for monitoring urban areas, e.g., by mounting the cameras on buildings or other infrastructure. Exploring novel modalities for air quality monitoring enabled by these improvements is another important research direction.

9 SUMMARY AND CONCLUSION

Low-cost air quality monitoring technology is emerging as a complementary technology to professional-grade air quality stations. The high cost of professional-grade stations limits the granularity at which they can capture pollutant concentrations, whereas low-cost sensors can be deployed densely to increase the spatial granularity of collected information. Unfortunately, the accuracy of low-cost sensors tends to be poor as the sensors are vulnerable to several sources of noise. In this article, we have critically surveyed machine-learning-based calibration of low-cost air quality sensors, the main technique for improving the usefulness of measurements provided by low-cost air quality sensors. Our focus has been on individual sensing units, each of which typically integrates several different sensors (e.g., environmental sensors, particulate matter sensors, and sensors for gaseous pollutants). In this survey, we have covered the sensor technology itself, the processing pipeline required for calibration, the machine learning techniques that are used in calibration, and different ways to evaluate the performance of calibration models. Based on our survey, we have highlighted open research issues in the field, with the inconsistency of studies, lack of sufficiently long datasets, and lack of models that perform well across several pollutants being among the most critical research problems.

REFERENCES

- [1] S. O. Agbroko and J. Covington. 2018. A novel, low-cost, portable PID sensor for the detection of volatile organic compounds. *Sens. Actuat. B: Chem.* 275 (2018), 10–15. DOI: <https://doi.org/10.1016/j.snb.2018.07.173>
- [2] M. Aleixandre and M. Gerboles. 2012. Review of small commercial sensors for indicative monitoring of ambient gas. *Chem. Eng. Trans.* 30 (Sep. 2012), 169–174. DOI: <https://doi.org/10.3303/CET1230029>
- [3] Z. J. Andersen, K. Bønnelykke, M. Hvidberg, S. S. Jensen, M. Ketzel, S. Loft, M. Sørensen, A. Tjønneland, K. Overvad, and O. Raaschou-Nielsen. 2012. Long-term exposure to air pollution and asthma hospitalisations in older adults: A cohort study. *Thorax* 67, 1 (2012), 6–11. DOI: <https://doi.org/10.1136/thoraxjnl-2011-200711>
- [4] Z. J. Andersen, M. Hvidberg, S. S. Jensen, M. Ketzel, S. Loft, M. Sørensen, A. Tjønneland, K. Overvad, and O. Raaschou-Nielsen. 2011. Chronic obstructive pulmonary disease and long-term exposure to traffic-related air pollution. *Am. J. Respir. Crit. Care Med.* 183, 4 (2011), 455–461. DOI: <https://doi.org/10.1164/rccm.201006-0937OC>
- [5] Z. J. Andersen, L. C. Kristiansen, K. K. Andersen, T. S. Olsen, M. Hvidberg, S. S. Jensen, M. Ketzel, S. Loft, M. Sørensen, A. Tjønneland, K. Overvad, and O. Raaschou-Nielsen. 2012. Stroke and long-term exposure to outdoor air pollution from nitrogen dioxide. *Stroke* 43, 2 (2012), 320–325. DOI: <https://doi.org/10.1161/STROKEAHA.111.629246>
- [6] Array of Things. 2018. Array of Things Expands with Partner Projects Around U.S. Retrieved November 12, 2020 from <https://medium.com/array-of-things/array-of-things-expands-with-partner-projects-around-u-s-c0ef9171f314>.
- [7] E. Austin, I. Novosselov, E. Seto, and M. G. Yost. 2015. Laboratory evaluation of the Shinyei PPD42NS low-cost particulate matter sensor. *PLoS One* 10, 9 (Sep. 2015), 1–17. DOI: <https://doi.org/10.1371/journal.pone.0137789>
- [8] L. Balzano and R. Nowak. 2007. Blind calibration of sensor networks. In *Proceedings of the 6th International Conference on Information Processing in Sensor Networks (IPSN'07)*. Association for Computing Machinery, New York, NY, 79–88. DOI: <https://doi.org/10.1145/1236360.1236372>
- [9] R. Baron and J. Saffell. 2017. Amperometric gas sensors as a low cost emerging technology platform for air quality monitoring applications: A review. *ACS Sens.* 2, 11 (2017), 1553–1566. DOI: <https://doi.org/10.1021/acssensors.7b00620>
- [10] N. Barsan, D. Koziej, and U. Weimar. 2007. Metal oxide-based gas sensor research: How to? *Sens. Actuat. B: Chem.* 121, 1 (2007), 18–35. DOI: <https://doi.org/10.1016/j.snb.2006.09.047>
- [11] M. Bart, D. E. Williams, B. Ainslie, I. McKendry, J. Salmond, S. K. Grange, M. Alavi-Shoshtari, D. Steyn, and G. S. Henshaw. 2014. High density ozone monitoring using gas sensitive semi-conductor sensors in the lower Fraser valley, British Columbia. *Environ. Sci. Technol.* 48, 7 (2014), 3970–3977. DOI: <https://doi.org/10.1021/es404610t>
- [12] R. Beelen, G. Hoek, D. Vienneau, M. Eeftens, K. Dimakopoulou, X. Pedeli, M. Tsai, N. Künzli, T. Schikowski, A. Marcon, K. T. Eriksen, O. Raaschou-Nielsen, E. Stephanou, E. Patelarou, T. Lanki, T. Yli-Tuomi, C. Declercq, G. Falq, M. Stempfelet, M. Birk, J. Cyrys, S. von Klot, G. Nádor, M. J. Varró, A. Dédélé, R. Gražulevičienė, A. Mölter, S. Lindley, C. Madsen, G. Cesaroni, A. Ranzi, C. Badaloni, B. Hoffmann, M. Nonnemacher, U. Krämer, T. Kuhlbusch, M. Cirach, A. de Nazelle, M. Nieuwenhuijsen, T. Bellander, M. Korek, D. Olsson, M. Strömberg, E. Dons, M. Jerrett, P. Fischer, M. Wang, B. Brunekreef, and K. de Hoogh. 2013. Development of NO₂ and NO_x land use regression models for estimating air pollution exposure in 36 study areas in Europe—The ESCAPE project. *Atmos. Environ.* 72 (2013), 10–23. DOI: <https://doi.org/10.1016/j.atmosenv.2013.02.037>
- [13] M. L. Bell, K. Ebisu, R. D. Peng, J. M. Samet, and F. Dominici. 2009. Hospital admissions and chemical composition of fine particle air pollution. *Am. J. Respir. Crit. Care Med.* 179, 12 (2009), 1115–1120. DOI: <https://doi.org/10.1164/rccm.200808-1240OC>
- [14] R. Berkowicz, F. Palmgren, O. Hertel, and E. Vignati. 1996. Using measurements of air pollution in streets for evaluation of urban air quality—meteorological analysis and model calculations. *Sci. Total Environ.* 189-190 (1996), 259–265. DOI: [https://doi.org/10.1016/0048-9697\(96\)05217-5](https://doi.org/10.1016/0048-9697(96)05217-5)
- [15] C. Borrego, A. M. Costa, J. Ginja, M. Amorim, M. Coutinho, K. Karatzas, Th. Sioumis, N. Katsifarakis, K. Konstantinidis, S. De Vito, E. Esposito, P. Smith, N. André, P. Gérard, L.A. Francis, N. Castell, P. Schneider, M. Viana, M. C. Minguiñón, W. Reimringer, R. P. Otjes, O. von Sicard, R. Pohle, B. Elen, D. Suriano, V. Pfister, M. Prato, S. Dipinto, and M. Penza. 2016. Assessment of air quality microsensors versus reference methods: The EuNetAir joint exercise. *Atmos. Environ.* 147 (2016), 246–263. DOI: <https://doi.org/10.1016/j.atmosenv.2016.09.050>
- [16] C. Borrego, J. Ginja, M. Coutinho, C. Ribeiro, K. Karatzas, Th. Sioumis, N. Katsifarakis, K. Konstantinidis, S. De Vito, E. Esposito, M. Salvato, P. Smith, N. André, P. Gérard, L.A. Francis, N. Castell, P. Schneider, M. Viana, M. C. Minguiñón, W. Reimringer, R. P. Otjes, O. von Sicard, R. Pohle, B. Elen, D. Suriano, V. Pfister, M. Prato, S. Dipinto, and M. Penza. 2018. Assessment of air quality microsensors versus reference methods: The EuNetAir Joint Exercise—Part II. *Atmos. Environ.* 193 (2018), 127–142. DOI: <https://doi.org/10.1016/j.atmosenv.2018.08.028>
- [17] I. Bos, Patrick De Boever, Luc Int Panis, and Romain Meeusen. 2014. Physical activity, air pollution and the brain. *Sports Med.* 44, 11 (01 Nov 2014), 1505–1518. DOI: <https://doi.org/10.1007/s40279-014-0222-6>
- [18] R. D. Brook, S. Rajagopalan, C. A. Pope, J. R. Brook, A. Bhatnagar, A. V. Diez-Roux, F. Holguin, Y. Hong, R. V. Luepker, M. A. Mittleman, A. Peters, D. Siscovick, S. C. Smith, L. Whitsel, and J. D. Kaufman. 2010. Particulate

- matter air pollution and cardiovascular disease. *Circulation* 121, 21 (2010), 2331–2378. DOI : <https://doi.org/10.1161/CIR.0b013e3181dbee1>
- [19] J. Burgués and S. Marco. 2018. Low power operation of temperature-modulated metal oxide semiconductor gas sensors. *Sensors* 18, 2 (2018), 15. DOI : <https://doi.org/10.3390/s18020339>
- [20] N. Castell, F. R. Dauge, P. Schneider, M. Vogt, U. Lerner, B. Fishbain, D. Broday, and A. Bartonova. 2017. Can commercial low-cost sensor platforms contribute to air quality monitoring and exposure estimates? *Environ. Int.* 99 (2017), 293–302. DOI : <https://doi.org/10.1016/j.envint.2016.12.007>
- [21] T. Chai and R. R. Draxler. 2014. Root mean square error (RMSE) or mean absolute error (MAE)? - Arguments against avoiding RMSE in the literature. *Geosci. Model Dev.* 7, 3 (2014), 1247–1250. DOI : <https://doi.org/10.5194/gmd-7-1247-2014>
- [22] B. Chen, D. Schwegler-Berry, A. Cumpston, J. Cumpston, S. Friend, S. Stone, and M. Keane. 2016. Performance of a scanning mobility particle sizer in measuring diverse types of airborne nanoparticles: Multi-walled carbon nanotubes, welding fumes, and titanium dioxide spray. *J. Occupat. Environ. Hygiene* 13 (Feb. 2016). DOI : <https://doi.org/10.1080/15459624.2016.1148267>
- [23] T. Chen and C. Guestrin. 2016. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'16)*. Association for Computing Machinery, New York, NY, 785–794. DOI : <https://doi.org/10.1145/2939672.2939785>
- [24] Y. Cheng, X. He, Z. Zhou, and L. Thiele. 2019. ICT: In-field calibration transfer for air quality sensor deployments. *Proc. ACM Interact. Mob. Wearable Ubiqu. Technol.* 3, 1, Article 6 (Mar. 2019), 19 pages. DOI : <https://doi.org/10.1145/3314393>
- [25] Y. Cheng, X. He, Z. Zhou, and L. Thiele. 2020. MapTransfer: Urban air quality map generation for downscaled sensor deployments. In *Proceedings of the 2020 IEEE/ACM 5th International Conference on Internet-of-Things Design and Implementation (IoTDI'20)*. IEEE, 14–26. DOI : <https://doi.org/10.1109/IoTDI49375.2020.00010>
- [26] Y. Cheng, X. Li, Z. Li, S. Jiang, Y. Li, J. Jia, and X. Jiang. 2014. AirCloud: A cloud-based air-quality monitoring system for everyone. In *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems (SenSys'14)*. Association for Computing Machinery, New York, NY, 251–265. DOI : <https://doi.org/10.1145/2668332.2668346>
- [27] D. P. Chock, T. R. Terrell, and S. B. Levitt. 1975. Time-series analysis of Riverside, California air quality data. *Atmos. Environ.* 9, 11 (1975), 978–989. DOI : [https://doi.org/10.1016/0004-6981\(75\)90017-7](https://doi.org/10.1016/0004-6981(75)90017-7)
- [28] M. R. Chowdhury, S. De, N. K. Shukla, and R. N. Biswas. 2018. Energy-efficient air pollution monitoring with optimum duty-cycling on a sensor hub. In *Proceedings of the 2018 24th National Conference on Communications (NCC)*. IEEE, 1–6. DOI : <https://doi.org/10.1109/NCC.2018.8600133>
- [29] Z. Chowdhury, R. D. Edwards, M. Johnson, K. Naumoff Shields, T. Allen, E. Canuz, and K. R. Smith. 2007. An inexpensive light-scattering particle monitor: Field validation. *J. Environ. Monit.* 9, 10 (2007), 1099–1106. DOI : <https://doi.org/10.1039/B709329M>
- [30] J. M. Cordero, R. Borge, and A. Narros. 2018. Using statistical methods to carry out in field calibrations of low cost air quality sensors. *Sens. Actuat. B: Chem.* 267 (2018), 245–254. DOI : <https://doi.org/10.1016/j.snb.2018.04.021>
- [31] Council of the European Union. 2008. Directive 2008/50/EC of the European parliament and of the council of 21 May 2008 on ambient air quality and cleaner air for Europe. *Off. J. Eur. Union* L152 (2008), 1–44.
- [32] E. S. Cross, L. R. Williams, D. K. Lewis, G. R. Magoon, T. B. Onasch, M. L. Kaminsky, D. R. Worsnop, and J. T. Jayne. 2017. Use of electrochemical sensors for measurement of air pollution: Correcting interference response and validating measurements. *Atmos. Meas. Techn.* 10, 9 (2017), 3575–3588. DOI : <https://doi.org/10.5194/amt-10-3575-2017>
- [33] P. Das, S. Ghosh, S. Chatterjee, and S. De. 2020. Energy harvesting-enabled 5G advanced air pollution monitoring device. In *Proceedings of the 2020 IEEE 3rd 5G World Forum (5GWF'20)*. IEEE, 218–223. DOI : <https://doi.org/10.1109/5GWF49715.2020.9221330>
- [34] S. De Vito, G. Di Francia, E. Esposito, S. Ferlito, F. Formisano, and E. Massera. 2020. Adaptive machine learning strategies for network calibration of IoT smart air quality monitoring devices. *Pattern Recogn. Lett.* 136 (2020), 264–271. DOI : <https://doi.org/10.1016/j.patrec.2020.04.032>
- [35] S. De Vito, E. Esposito, N. Castell, P. Schneider, and A. Bartonova. 2020. On the robustness of field calibration for smart air quality monitors. *Sens. Actuat. B: Chem.* 310 (2020), 127869. DOI : <https://doi.org/10.1016/j.snb.2020.127869>
- [36] S. De Vito, E. Massera, M. Piga, L. Martinotto, and G. Di Francia. 2008. On field calibration of an electronic nose for benzene estimation in an urban pollution monitoring scenario. *Sens. Actuat. B: Chem.* 129, 2 (2008), 750–757. DOI : <https://doi.org/10.1016/j.snb.2007.09.060>
- [37] S. De Vito, M. Piga, L. Martinotto, and G. Di Francia. 2009. CO, NO₂ and NO_x urban pollution monitoring with on-field calibrated electronic nose by automatic bayesian regularization. *Sens. Actuat. B: Chem.* 143, 1 (2009), 182–191. DOI : <https://doi.org/10.1016/j.snb.2009.08.041>

- [38] T. Dinh, I. Choi, Y. Son, and J. Kim. 2016. A review on non-dispersive infrared gas sensors: Improvement of sensor detection limit and interference correction. *Sens. Actuat. B: Chem.* 231 (2016), 529–538. DOI: <https://doi.org/10.1016/j.snb.2016.03.040>
- [39] A. Donnelly, B. Misstear, and B. Broderick. 2015. Real time air quality forecasting using integrated parametric and non-parametric regression techniques. *Atmos. Environ.* 103 (2015), 53–65. DOI: <https://doi.org/10.1016/j.atmosenv.2014.12.011>
- [40] B. N. Duncan, A. I. Prados, L. N. Lamsal, Y. Liu, D. G. Streets, P. Gupta, E. Hilsenrath, R. A. Kahn, J. E. Nielsen, A. J. Beyersdorf, S. P. Burton, A. M. Fiore, J. Fishman, D. K. Henze, C. A. Hostetler, N. A. Krotkov, P. Lee, M. Lin, S. Pawson, G. Pfister, K. E. Pickering, R. B. Pierce, Y. Yoshida, and L. D. Ziemba. 2014. Satellite data of atmospheric pollution for U.S. air quality applications: Examples of applications, summary of data end-user resources, answers to FAQs, and common mistakes to avoid. *Atmos. Environ.* 94 (2014), 647–662. DOI: <https://doi.org/10.1016/j.atmosenv.2014.05.061>
- [41] P. Dutta, P. M. Aoki, N. Kumar, A. Mainwaring, C. Myers, W. Willett, and A. Woodruff. 2009. Common sense: Participatory urban sensing using a network of handheld air quality monitors. In *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems (SenSys'09)*. ACM, New York, NY, 349–350. DOI: <https://doi.org/10.1145/1644038.1644095>
- [42] M. Eeftens, R. Beelen, K. de Hoogh, T. Bellander, G. Cesaroni, M. Cirach, C. Declercq, A. Dèdelè, E. Dons, A. de Nazelle, K. Dimakopoulou, K. Eriksen, G. Falq, P. Fischer, C. Galassi, R. Gražulevičienė, J. Heinrich, B. Hoffmann, M. Jerrett, D. Keidel, M. Korek, T. Lanki, S. Lindley, C. Madsen, A. Mölter, G. Nádor, M. Nieuwenhuijsen, M. Nonnemacher, X. Pedeli, O. Raaschou-Nielsen, E. Patelarou, U. Quass, A. Ranzi, C. Schindler, M. Stempfelet, E. Stephanou, D. Sugiri, M. Tsai, T. Yli-Tuomi, M. J. Varró, D. Vienneau, S. von Klot, K. Wolf, B. Brunekreef, and G. Hoek. 2012. Development of land use regression models for PM_{2.5}, PM_{2.5} absorbance, PM₁₀ and PM_{coarse} in 20 European study areas; results of the ESCAPE project. *Environ. Sci. Technol.* 46, 20 (2012), 11195–11205. DOI: <https://doi.org/10.1021/es301948k>
- [43] E. Esposito, S. De Vito, M. Salvato, V. Bright, R.L. Jones, and O. Popoola. 2016. Dynamic neural network architectures for on field stochastic calibration of indicative low cost air quality sensing systems. *Sens. Actuat. B: Chem.* 231 (2016), 701–713. DOI: <https://doi.org/10.1016/j.snb.2016.03.038>
- [44] M. Fierz, C. Houle, P. Steigmeier, and H. Burtcher. 2011. Design, calibration, and field performance of a miniature diffusion size classifier. *Aerosol Sci. Technol.* 45, 1 (2011), 1–10. DOI: <https://doi.org/10.1080/02786826.2010.516283>
- [45] Y. Freund and R. E. Schapire. 1996. Experiments with a new boosting algorithm. In *Proceedings of the 13th International Conference on International Conference on Machine Learning (ICML'96)*. Morgan Kaufmann, San Francisco, CA, 148–156. DOI: <https://doi.org/10.5555/3091696.3091715>
- [46] J. H. Friedman. 2001. Greedy function approximation: A gradient boosting machine. *Ann. Stat.* 29, 5 (2001), 1189–1232.
- [47] J. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy, and A. Bouchachia. 2014. A survey on concept drift adaptation. *ACM Comput. Surv.* 46, 4, Article 44 (Mar. 2014), 37 pages. DOI: <https://doi.org/10.1145/2523813>
- [48] Y. Gao, W. Dong, K. Guo, X. Liu, Y. Chen, X. Liu, J. Bu, and C. Chen. 2016. Mosaic: A low-cost mobile sensing system for urban air quality monitoring. In *Proceedings of the 35th Annual IEEE International Conference on Computer Communications (IEEE INFOCOM'16)*. IEEE, 1–9. DOI: <https://doi.org/10.1109/INFOCOM.2016.7524478>
- [49] U. Gehring, A. H. Wijga, M. Brauer, P. Fischer, J. C. de Jongste, M. Kerkhof, M. Oldenwening, H. A. Smit, and B. Brunekreef. 2010. Traffic-related air pollution and the development of asthma and allergies during the first 8 years of life. *Am. J. Respir. Crit. Care Med.* 181, 6 (2010), 596–603. DOI: <https://doi.org/10.1164/rccm.200906-0858OC>
- [50] M. Goldberg. 2011. A systematic review of the relation between long-term exposure to ambient air pollution and chronic diseases. *Rev. Environ. Health* 23, 4 (2011), 243–298. DOI: <https://doi.org/10.1515/REVEH.2008.23.4.243>
- [51] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi. 2018. A survey of methods for explaining black box models. *ACM Comput. Surv.* 51, 5, Article 93 (Aug. 2018), 42 pages. DOI: <https://doi.org/10.1145/3236009>
- [52] L. H. Schmoll, T. M. Peters, and P. O'Shaughnessy. 2010. Use of a condensation particle counter and an optical particle counter to assess the number concentration of engineered nanoparticles. *J. Occupat. Environ. Hygiene* 7 (Sep. 2010), 535–545. DOI: <https://doi.org/10.1080/15459624.2010.496072>
- [53] D. Hasenfratz, T. Arn, I. de Concini, O. Saukh, and L. Thiele. 2015. Health-optimal routing in urban areas. In *Proceedings of the 14th International Conference on Information Processing in Sensor Networks (IPSN'15)*. ACM, New York, NY, 398–399. DOI: <https://doi.org/10.1145/2737095.2737135>
- [54] D. Hasenfratz, O. Saukh, S. Sturzenegger, and L. Thiele. 2012. Participatory air pollution monitoring using smart-phones. In *Proceedings of the 2nd International Workshop on Mobile Sensing*. Academic Press, 1–5.
- [55] D. Hasenfratz, O. Saukh, and L. Thiele. 2012. On-the-fly calibration of low-cost gas sensors. In *Wireless Sensor Networks*, G. Picco and W. Heinzelman (Eds.). Springer, 228–244. DOI: https://doi.org/10.1007/978-3-642-28169-3_15
- [56] D. Hasenfratz, O. Saukh, C. Walser, C. Hueglin, M. Fierz, and L. Thiele. 2014. Pushing the spatio-temporal resolution limit of urban air pollution maps. In *Proceedings of the 2014 IEEE International Conference on Pervasive Computing and Communications (PerCom'14)*. IEEE, 69–77. DOI: <https://doi.org/10.1109/PerCom.2014.6813946>

- [57] T. Hastie, R. Tibshirani, and J. Friedman. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, New York, NY.
- [58] G. Hoek, R. Beelen, K. de Hoogh, D. Vienneau, J. Gulliver, P. Fischer, and D. Briggs. 2008. A review of land-use regression models to assess spatial variation of outdoor air pollution. *Atmos. Environ.* 42, 33 (2008), 7561–7578. DOI: <https://doi.org/10.1016/j.atmosenv.2008.05.057>
- [59] W. Hou, J. Wang, X. Xu, J. S. Reid, S. J. Janz, and J. W. Leitch. 2020. An algorithm for hyperspectral remote sensing of aerosols: 3. Application to the GEO-TASO data in KORUS-AQ field campaign. *J. Quant. Spectrosc. Radiat. Transf.* 253 (2020), 107161. DOI: <https://doi.org/10.1016/j.jqsrt.2020.107161>
- [60] K. Hu, A. Rahman, H. Bhugubanda, and V. Sivaraman. 2017. HazeEst: Machine learning based metropolitan air pollution estimation from fixed and mobile sensors. *IEEE Sens. J.* 17, 11 (Jun. 2017), 3517–3525. DOI: <https://doi.org/10.1109/JSEN.2017.2690975>
- [61] K. Hu, V. Sivaraman, H. Bhugubanda, S. Kang, and A. Rahman. 2016. SVR based dense air pollution estimation model using static and wireless sensor network. In *Proceedings of the 2016 IEEE Conference on Sensors (SENSORS'16)*. IEEE, 1–3. DOI: <https://doi.org/10.1109/ICSENS.2016.7808827>
- [62] A. K. Jain, Jianchang Mao, and K. M. Mohiuddin. 1996. Artificial neural networks: A tutorial. *Computer* 29, 3 (1996), 31–44. DOI: <https://doi.org/10.1109/2.485891>
- [63] W. Jiao, G. Hagler, R. Williams, R. Sharpe, R. Brown, D. Garver, R. Judge, M. Caudill, J. Rickard, M. Davis, L. Weinstein, S. Zimmer-Dauphinee, and K. Buckley. 2016. Community Air Sensor Network (CAIRSENSE) project: Evaluation of low-cost sensor performance in a suburban environment in the southeastern United States. *Atmos. Meas. Techn.* 9, 11 (2016), 5281–5292. DOI: <https://doi.org/10.5194/amt-9-5281-2016>
- [64] M. Johnson, V. Isakov, J. S. Tuma, S. Mukerjee, and H. Özkaynak. 2010. Evaluation of land-use regression models used to predict air quality concentrations in an urban area. *Atmos. Environ.* 44, 30 (2010), 3660–3668. DOI: <https://doi.org/10.1016/j.atmosenv.2010.06.041>
- [65] N. E. Johnson, B. Bonczak, and C. E. Kontokosta. 2018. Using a gradient boosting model to improve the performance of low-cost aerosol monitors in a dense, heterogeneous urban environment. *Atmos. Environ.* 184 (2018), 9–16. DOI: <https://doi.org/10.1016/j.atmosenv.2018.04.019>
- [66] M. Kim, J. Choi, and W. Park. 2018. MEMS PZT oscillating platform for fine dust particle removal at resonance. *Int. J. Precis. Eng. Manufact.* 19, 12 (01 Dec. 2018), 1851–1859. DOI: <https://doi.org/10.1007/s12541-018-0214-9>
- [67] M. Kulmala. 2018. Build a Global Earth Observatory. Retrieved from <https://www.nature.com/articles/d41586-017-08967-y>.
- [68] J. Kuula, T. Mäkelä, R. Hillamo, and H. Timonen. 2017. Response characterization of an inexpensive aerosol sensor. *Sensors (Switzerland)* 17 (Nov. 2017), 14. DOI: <https://doi.org/10.3390/s17122915>
- [69] E. Lagerspetz, N. H. Motlagh, M. Arbayani Zaidan, P. L. Fung, J. Mineraud, S. Varjonen, M. Siekkinen, P. Nurmi, Y. Matsumi, S. Tarkoma, and T. Hussein. 2019. MegaSense: Feasibility of low-cost sensors for pollution hot-spot detection. In *Proceedings of the 2019 IEEE 17th International Conference on Industrial Informatics (INDIN'19)*, Vol. 1. IEEE, 1083–1090. DOI: <https://doi.org/10.1109/INDIN41052.2019.8971963>
- [70] Y. LeCun, Y. Bengio, and G. Hinton. 2015. Deep learning. *Nature* 521 (May 2015), 436–444. <https://doi.org/10.1038/nature14539>
- [71] H. Lee, J. Kang, S. Kim, Y. Kim, S. Yoo, and D. Lee. 2020. Long-term evaluation and calibration of low-cost particulate matter (PM) sensor. *Sensors* 20, 13 (2020), 3617. DOI: <https://doi.org/10.3390/s20133617>
- [72] A. Lewis and P. Edwards. 2016. Validate personal air-pollution sensors. *Nature* 535 (2016), 29–31. DOI: <https://doi.org/10.1038/535029a>
- [73] A. C. Lewis, E. von Schneidmesser, and R. E. Peltier. 2018. *Low-cost Sensors for the Measurement of Atmospheric Composition: Overview of Topic and Future Applications*. Technical Report 1215. World Meteorological Organization, Geneva, Switzerland.
- [74] J. J. Li, B. Faltings, O. Saukh, D. Hasenfratz, and J. Beutel. 2012. Sensing the air we breathe: The opensense Zurich dataset. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI'12)*. AAAI Press, 323–325.
- [75] C. Lin, J. Gillespie, M. D. Schuder, W. Duberstein, I. J. Beverland, and M. R. Heal. 2015. Evaluation and calibration of Aeroqual series 500 portable gas sensors for accurate measurement of ambient ozone and nitrogen dioxide. *Atmos. Environ.* 100 (2015), 111–116. DOI: <https://doi.org/10.1016/j.atmosenv.2014.11.002>
- [76] Y. Lin, W. Dong, and Y. Chen. 2018. Calibrating low-cost sensors by a two-phase learning approach for urban air quality measurement. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 1, Article 18 (Mar. 2018), 18 pages. DOI: <https://doi.org/10.1145/3191750>
- [77] Z. C. Lipton. 2018. The mythos of model interpretability. *Commun. ACM* 61, 10 (Sep. 2018), 36–43. DOI: <https://doi.org/10.1145/3233231>
- [78] B. Liu, A. Binaykia, P. Chang, M. K. Tiwari, and C. Tsao. 2017. Urban air quality forecasting based on multi-dimensional collaborative Support Vector Regression (SVR): A case study of Beijing-Tianjin-Shijiazhuang. *PLoS One* 12, 7 (07 2017), 1–17. DOI: <https://doi.org/10.1371/journal.pone.0179763>

- [79] D. Liu, Q. Zhang, J. Jiang, and D. Chen. 2017. Performance calibration of low-cost and portable particulate matter (PM) sensors. *J. Aerosol Sci.* 112 (2017), 1–10. DOI : <https://doi.org/10.1016/j.jaerosci.2017.05.011>
- [80] H. Liu, H. Wu, H. Lee, Y. Ho, and L. Chen. 2017. A system calibration model for mobile PM_{2.5} sensing using low-cost sensors. In *Proceedings of the 2017 IEEE International Conference on Internet of Things (iThings'17) and IEEE Green Computing and Communications (GreenCom'17) and IEEE Cyber, Physical and Social Computing (CPSCom'17) and IEEE Smart Data (SmartData'17)*. IEEE, 611–618. DOI : <https://doi.org/10.1109/iThings-GreenCom-CPSCom-SmartData.2017.97>
- [81] L. Liu, W. Liu, Y. Zheng, H. Ma, and C. Zhang. 2018. Third-eye: A mobilephone-enabled crowdsensing system for air quality monitoring. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 1, Article 20 (Mar. 2018), 26 pages. DOI : <https://doi.org/10.1145/3191752>
- [82] B. Maag, O. Saukh, D. Hasenfratz, and L. Thiele. 2016. Pre-deployment testing, augmentation and calibration of cross-sensitive sensors. In *Proceedings of the 2016 International Conference on Embedded Wireless Systems and Networks (EWSN'16)*. Junction Publishing, 169–180. <http://dl.acm.org/citation.cfm?id=2893711.2893735>
- [83] B. Maag, Z. Zhou, O. Saukh, and L. Thiele. 2017. SCAN: Multi-hop calibration for mobile sensor arrays. *Proc. ACM Interact. Mobile Wear. Ubiqu. Technol.* 1, 2 (2017), Article 19, 21 pages. DOI : <https://doi.org/10.1145/3090084>
- [84] B. Maag, Z. Zhou, and L. Thiele. 2018. A survey on sensor calibration in air pollution monitoring deployments. *IEEE IoT J.* 5, 6 (Dec. 2018), 4857–4870. DOI : <https://doi.org/10.1109/JIOT.2018.2853660>
- [85] B. Maag, Z. Zhou, and L. Thiele. 2018. W-air: Enabling personal air pollution monitoring on wearables. *Proc. ACM Interact. Mob. Wearable Ubiqu. Technol.* 2, 1, Article 24 (Mar. 2018), 25 pages. DOI : <https://doi.org/10.1145/3191756>
- [86] N. Masson, R. Piedrahita, and M. Hannigan. 2015. Approach for quantification of metal oxide type semiconductor gas sensors used for ambient air quality monitoring. *Sens. Actuat. B: Chem.* 208 (2015), 339–345. DOI : <https://doi.org/10.1016/j.snb.2014.11.032>
- [87] N. Masson, R. Piedrahita, and M. Hannigan. 2015. Quantification method for electrolytic sensors in long-term monitoring of ambient air quality. *Sensors* 15, 10 (2015), 27283–27302. DOI : <https://doi.org/10.3390/s151027283>
- [88] B. C. McDonald, J. A. de Gouw, J. B. Gilman, S. H. Jathar, A. Akherati, C. D. Cappa, J. L. Jimenez, J. Lee-Taylor, P. L. Hayes, S. A. McKeen, Y. Y. Cui, S. Kim, D. R. Gentner, G. Isaacman-VanWertz, A. H. Goldstein, R. A. Harley, G. J. Frost, J. M. Roberts, T. B. Ryerson, and M. Trainer. 2018. Volatile chemical products emerging as largest petrochemical source of urban organic emissions. *Science* 359, 6377 (2018), 760–764. DOI : <https://doi.org/10.1126/science.aag0524>
- [89] R. Meier, K. Clark, and M. Riediker. 2012. Comparative testing of a miniature diffusion size classifier to assess airborne ultrafine particles under field conditions. *Aerosol Sci. Technol.* 47 (Jan. 2012). DOI : <https://doi.org/10.1080/02786826.2012.720397>
- [90] P. H. Merz, L. J. Painter, and P. R. Ryason. 1972. Aerometric data analysis—Time series analysis and forecast and an atmospheric smog diagram. *Atmos. Environ.* 6, 5 (1972), 319–342. DOI : [https://doi.org/10.1016/0004-6981\(72\)90199-0](https://doi.org/10.1016/0004-6981(72)90199-0)
- [91] D. Mintz. 2018. *Technical Assistance Document for the Reporting of Daily Air Quality—The Air Quality Index (AQI)*. U.S. Environmental Protection Agency, Research Triangle Park, North Carolina 27711. Retrieved from <https://www.airnow.gov/sites/default/files/2020-05/aqi-technical-assistance-document-sept2018.pdf>.
- [92] G. Miskell, K. Alberti, B. Feenstra, G. S. Henshaw, V. Papapostolou, H. Patel, A. Polidori, J. A. Salmond, L. Weissert, and D. E. Williams. 2019. Reliable data from low cost ozone sensors in a hierarchical network. *Atmos. Environ.* 214 (2019), 116870. DOI : <https://doi.org/10.1016/j.atmosenv.2019.116870>
- [93] S. Moltchanov, I. Levy, Y. Etzion, U. Lerner, D. M. Broday, and B. Fishbain. 2015. On the feasibility of measuring urban air pollution by wireless distributed sensor networks. *Sci. Total Environ.* 502 (2015), 537–547. DOI : <https://doi.org/10.1016/j.scitotenv.2014.09.059>
- [94] A. Moore, M. Figliozzi, and C. M. Monsere. 2012. Air quality at bus stops: Empirical analysis of exposure to particulate matter at bus stop shelters. *Transport. Res. Rec.* 2270, 1 (2012), 76–86. DOI : <https://doi.org/10.3141/2270-10>
- [95] L. Morawska, P. K. Thai, X. Liu, A. Asumadu-Sakyi, G. Ayoko, A. Bartonova, A. Bedini, F. Chai, B. Christensen, M. Dunbabin, J. Gao, G. S.W. Hagler, R. Jayaratne, P. Kumar, A. K. H. Lau, P. K. K. Louie, M. Mazaheri, Z. Ning, N. Motta, B. Mullins, M. Rahman, Z. Ristovski, M. Shafiei, D. Tjondronegoro, D. Westerdaahl, and R. Williams. 2018. Applications of low-cost sensing technologies for air quality monitoring and exposure assessment: How far have they gone? *Environ. Int.* 116 (2018), 286–299. DOI : <https://doi.org/10.1016/j.envint.2018.04.018>
- [96] N. H. Motlagh, E. Lagerspetz, P. Nurmi, X. Li, S. Varjonen, J. Mineraud, M. Siekkinen, A. Rebeiro-Hargrave, T. Hussein, T. Petaja, M. Kulmala, and S. Tarkoma. 2020. Toward massive scale air quality monitoring. *IEEE Commun. Mag.* 58, 2 (2020), 54–59. DOI : <https://doi.org/10.1109/MCOM.001.1900515>
- [97] M. Mueller, J. Meyer, and C. Hueglin. 2017. Design of an ozone and nitrogen dioxide sensor unit and its long-term operation within a sensor network in the city of Zurich. *Atmos. Meas. Techn.* 10, 10 (2017), 3783–3799. DOI : <https://doi.org/10.5194/amt-10-3783-2017>
- [98] N. Nikzad, N. Verma, C. Zifti, E. Bales, N. Quick, P. Zappi, K. Patrick, S. Dasgupta, I. Krueger, T. Šimunić Rosing, and W. G. Griswold. 2012. CitiSense: Improving geospatial environmental assessment of air quality using a wireless

- personal exposure monitoring system. In *Proceedings of the Conference on Wireless Health (WH'12)*. ACM, New York, NY, Article 11, 8 pages. DOI : <https://doi.org/10.1145/2448096.2448107>
- [99] OECD. 2016. *The Economic Consequence of Outdoor Air Pollution*. OECD Publishing, Paris. DOI : <https://doi.org/10.1787/9789264257474-en>
- [100] D. Oletic and V. Bilas. 2015. Design of sensor node for air quality crowdsensing. In *Proceedings of the IEEE Sensors Applications Symposium (SAS'15)*. IEEE, 1–5. DOI : <https://doi.org/10.1109/SAS.2015.7133628>
- [101] H. Olstrup, C. Johansson, B. Forsberg, A. Tornevi, A. Ekeboom, and K. Meister. 2019. A Multi-Pollutant Air Quality Health Index (AQHI) based on short-term respiratory effects in Stockholm, Sweden. *Int. J. Environ. Res. Publ. Health* 16, 1 (Jan. 2019), 105. DOI : <https://doi.org/10.3390/ijerph16010105>
- [102] Ö. Özgür. 2012. An overview of metal oxide semiconducting sensors in electronic nose applications. In *Proceedings of the 3rd International Symposium on Sustainable Development*. 506–515. <https://core.ac.uk/download/pdf/153447895.pdf>
- [103] R. Piedrahita, Y. Xiang, N. Masson, J. Ortega, A. Collier, Y. Jiang, K. Li, R. P. Dick, Q. Lv, M. Hannigan, and L. Shang. 2014. The next generation of low-cost personal air quality sensors for quantitative exposure monitoring. *Atmos. Meas. Techn.* 7, 10 (2014), 3325–3336. DOI : <https://doi.org/10.5194/amt-7-3325-2014>
- [104] T. Pulkkinen, J. Nurminen, and P. Nurmi. 2020. Understanding WiFi cross-technology interference detection in the real world. In *Proceedings of the 40th International Conference on Distributed Computing Systems (ICDCS'20)*. IEEE, 11.
- [105] O. Raaschou-Nielsen, Z. J. Andersen, M. Hvidberg, S. S. Jensen, M. Ketzel, M. Sørensen, J. Hansen, S. Loft, K. Overvad, and A. Tjønneland. 2011. Air pollution from traffic and cancer incidence: A Danish cohort study. *Environ. Health* 10, 1 (19 Jul. 2011), 67. DOI : <https://doi.org/10.1186/1476-069X-10-67>
- [106] L. F. Radke, J. L. Stith, D. A. Hegg, and P. V. Hobbs. 1978. Airborne studies of particles and gases from forest fires. *J. Air Pollut. Contr. Assoc.* 28, 1 (1978), 30–34. DOI : <https://doi.org/10.1080/00022470.1978.10470566>
- [107] N. Ramanathan, L. Balzano, M. C. Burt, D. Estrin, T. Harmon, C. K. Harvey, J. Jay, E. Kohler, S. E. Rothenberg, and M. Srivastava. 2006. *Rapid Deployment with Confidence: Calibration and Fault Detection in Environmental Sensor Networks*. Technical Report. University of California. Retrieved from <https://escholarship.org/uc/item/8v26b5qh>.
- [108] C. E. Rasmussen. 2004. *Gaussian Processes in Machine Learning*. Springer, Berlin, 63–71. DOI : https://doi.org/10.1007/978-3-540-28650-9_4
- [109] S. T. Rinne and J. D. Kaufman. 2012. Air pollution. In *Clinical Respiratory Medicine* (3rd ed.), Stephen G. Spiro, Gerard A. Silvestri, and Alvar Agustí (Eds.). W.B. Saunders, Philadelphia, 937–945. DOI : <https://doi.org/10.1016/B978-1-4557-0792-8.00080-5>
- [110] A. C. Romain and J. Nicolas. 2010. Long term stability of metal oxide-based gas sensors for e-nose environmental applications: An overview. *Sens. Actuat. B: Chem.* 146, 2 (2010), 502–506. DOI : <https://doi.org/10.1016/j.snb.2009.12.027>
- [111] J. S. Apte, K. Messier, S. Gani, M. Brauer, T. Kirchstetter, M. Lunden, J. D. Marshall, C. J. Portier, R. Vermeulen, and S. P. Hamburg. 2017. High-resolution air pollution mapping with Google street view cars: Exploiting big data. *Environ. Sci. Technol.* 51 (Jun. 2017), 6999–7008. DOI : <https://doi.org/10.1021/acs.est.7b00891>
- [112] E. M. Saber and G. Heydari. 2012. Flow patterns and deposition fraction of particles in the range of 0.1–10 μm at trachea and the first third generations under different breathing conditions. *Comput. Biol. Med.* 42, 5 (2012), 631–638. DOI : <https://doi.org/10.1016/j.combiomed.2012.03.002>
- [113] R. L. R. Salcedo, M. C. M. Alvim Ferraz, C. A. Alves, and F. G. Martins. 1999. Time-series analysis of air pollution data. *Atmos. Environ.* 33, 15 (1999), 2361–2372. DOI : [https://doi.org/10.1016/S1352-2310\(99\)80001-6](https://doi.org/10.1016/S1352-2310(99)80001-6)
- [114] O. Saukh, D. Hasenfratz, and L. Thiele. 2015. Reducing multi-hop calibration errors in large-scale mobile sensor networks. In *Proceedings of the 14th International Conference on Information Processing in Sensor Networks (IPSN'15)*. ACM, New York, NY, 274–285. DOI : <https://doi.org/10.1145/2737095.2737113>
- [115] W. Shao, H. Zhang, and H. Zhou. 2017. Fine particle sensor based on multi-angle light scattering and data fusion. *Sensors (Switzerland)* 17 (May 2017), 15. DOI : <https://doi.org/10.3390/s17051033>
- [116] Y. Shirai, Y. Kishino, F. Naya, and Y. Yanagisawa. 2016. Toward on-demand urban air quality monitoring using public vehicles. In *Proceedings of the 2nd International Workshop on Smart (SmartCities'16)*. Association for Computing Machinery, New York, NY, Article 1, 6 pages. DOI : <https://doi.org/10.1145/3009912.3009920>
- [117] J. H. Sohn, M. Atzeni, L. Zeller, and G. Pioggia. 2008. Characterisation of humidity dependence of a metal oxide semiconductor sensor array using partial least squares. *Sens. Actuat. B: Chem.* 131, 1 (2008), 230–235. DOI : <https://doi.org/10.1016/j.snb.2007.11.009>
- [118] S. Sousan, K. Koehler, G. Thomas, J. H. Park, M. Hillman, A. Halterman, and T. M. Peters. 2016. Inter-comparison of low-cost sensors for measuring the mass concentration of occupational aerosols. *Aerosol Sci. Technol.* 50, 5 (2016), 462–473. DOI : <https://doi.org/10.1080/02786826.2016.1162901>

- [119] L. Spinelle, M. Gerboles, G. Kok, S. Persijn, and T. Sauerwald. 2017. Review of portable and low-cost sensors for the ambient air monitoring of benzene and other volatile organic compounds. *Sensors* 17, 7 (Jun. 2017), 1520. DOI: <https://doi.org/10.3390/s17071520>
- [120] L. Spinelle, M. Gerboles, M. G. Villani, M. Aleixandre, and F. Bonavitacola. 2015. Field calibration of a cluster of low-cost available sensors for air quality monitoring. Part A: Ozone and nitrogen dioxide. *Sens. Actuat. B: Chem.* 215 (2015), 249–257. DOI: <https://doi.org/10.1016/j.snb.2015.03.031>
- [121] L. Spinelle, M. Gerboles, M. G. Villani, M. Aleixandre, and F. Bonavitacola. 2017. Field calibration of a cluster of low-cost commercially available sensors for air quality monitoring. Part B: NO, CO and CO₂. *Sens. Actuat. B: Chem.* 238 (2017), 706–715. DOI: <https://doi.org/10.1016/j.snb.2016.07.036>
- [122] D. G. Streets, T. Canty, G. R. Carmichael, B. de Foy, R. R. Dickerson, B. N. Duncan, D. P. Edwards, J. A. Haynes, D. K. Henze, M. R. Houyoux, D. J. Jacob, N. A. Krotkov, L. N. Lamsal, Y. Liu, Z. Lu, R. V. Martin, G. G. Pfister, R. W. Pinder, R. J. Salawitch, and K. J. Wecht. 2013. Emissions estimation from satellite retrievals: A review of current capability. *Atmos. Environ.* 77 (2013), 1011–1042. DOI: <https://doi.org/10.1016/j.atmosenv.2013.05.051>
- [123] B. Szulczyński and J. Gębicki. 2017. Currently commercially available chemical sensors employed for detection of volatile organic compounds in outdoor and indoor air. *Environments* 4, 1 (Mar. 2017), 21. DOI: <https://doi.org/10.3390/environments4010021>
- [124] G. Tang, J. Zhang, X. Zhu, T. Song, C. Munkel, B. Hu, K. Schäfer, Z. Liu, J. Zhang, L. Wang, J. Xin, P. Suppan, and Y. Wang. 2016. Mixing layer height and its implications for air pollution over Beijing, China. *Atmos. Chem. Phys.* 16, 4 (2016), 2459–2475. DOI: <https://doi.org/10.5194/acp-16-2459-2016>
- [125] A. C. Targino, M. D. Gibson, P. Krecl, M. V. Costa Rodrigues, M. M. dos Santos, and M. de Paula Corrêa. 2016. Hotspots of black carbon and PM_{2.5} in an urban area and relationships to traffic characteristics. *Environ. Pollut.* 218 (2016), 475–486. DOI: <https://doi.org/10.1016/j.envpol.2016.07.027>
- [126] J. E. Thompson. 2016. Crowd-sourced air quality studies: A review of the literature & portable sensors. *Trends Environ. Anal. Chem.* 11 (2016), 23–34. DOI: <https://doi.org/10.1016/j.teac.2016.06.001>
- [127] W. Tsujita, A. Yoshino, H. Ishida, and T. Moriizumi. 2005. Gas sensor network for air-pollution monitoring. *Sens. Actuat. B: Chem.* 110, 2 (2005), 304–311. DOI: <https://doi.org/10.1016/j.snb.2005.02.008>
- [128] J. E. Van Engelen and H. H. Hoos. 2020. A survey on semi-supervised learning. *Mach. Learn.* 109 (2020), 373–440. DOI: <https://doi.org/10.1007/s10994-019-05855-6>
- [129] J. Vanschoren. 2018. Meta-learning: A survey. arxiv:1810.03548. Retrieved from <http://arxiv.org/abs/1810.03548>.
- [130] S. Vardoulakis, N. Gonzalez-Flesca, B. E. A. Fisher, and K. Pericleous. 2005. Spatial variability of air pollution in the vicinity of a permanent monitoring station in central Paris. *Atmos. Environ.* 39, 15 (2005), 2725–2736. DOI: <https://doi.org/10.1016/j.atmosenv.2004.05.067>
- [131] A. Vasiliev, A. E. Varfolomeev, I. A. Volkov, N. Simonenko, P. V. Arsenov, I. Vlasov, V. V. Ivanov, A. V. Pisyakov, A. Lagutin, I. Jahatspanian, and T. Maeder. 2018. Reducing humidity response of gas sensors for medical applications: Use of spark discharge synthesis of metal oxide nanoparticles. *Sensors* 18 (Aug. 2018), 13. DOI: <https://doi.org/10.3390/s18082600>
- [132] H. Volk, F. Lurmann, B. Penfold, I. Hertz-Picciotto, and R. McConnell. 2013. Traffic-related air pollution, particulate matter, and autism. *JAMA Psychiatr.* 70, 1 (2013), 71–77. DOI: <https://doi.org/10.1001/jamapsychiatry.2013.266>
- [133] P. Wagner and K. Schäfer. 2017. Influence of mixing layer height on air pollutant concentrations in an urban street canyon. *Urban Clim.* 22 (2017), 64–79. DOI: <https://doi.org/10.1016/j.uclim.2015.11.001>
- [134] C. Wang, L. Yin, L. Zhang, D. Xiang, and R. Gao. 2010. Metal oxide gas sensors: Sensitivity and influencing factors. *Sensors* 10, 3 (2010), 2088–2106. DOI: <https://doi.org/10.3390/s100302088>
- [135] J. Wang and S. A. Christopher. 2003. Intercomparison between satellite-derived aerosol optical thickness and PM_{2.5} mass: Implications for air quality studies. *Geophys. Res. Lett.* 30, 21 (2003), 4–1–4–4. DOI: <https://doi.org/10.1029/2003GL018174>
- [136] J. Wang and S. A. Christopher. 2003. Intercomparison between satellite-derived aerosol optical thickness and PM_{2.5} mass: Implications for air quality studies. *Geophys. Res. Lett.* 30, 21 (2003), 3771–3784. DOI: <https://doi.org/10.1029/2003GL018174> arXiv:<https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2003GL018174>
- [137] R. Wang, S. B. Henderson, H. Sbihi, R. W. Allen, and M. Brauer. 2013. Temporal stability of land use regression models for traffic-related air pollution. *Atmos. Environ.* 64 (2013), 312–319. DOI: <https://doi.org/10.1016/j.atmosenv.2012.09.056>
- [138] Y. Wang, J. Li, H. Jing, Q. Zhang, J. Jiang, and P. Biswas. 2015. Laboratory evaluation and calibration of three low-cost particle sensors for particulate matter measurement. *Aerosol Sci. Technol.* 49 (Nov. 2015), 1063–1077. DOI: <https://doi.org/10.1080/02786826.2015.1100710>
- [139] Y. Wang, Q. Yao, J. Kwok, J. T. Kwok, and L. M. Ni. 2020. Generalizing from a few examples: A survey on few-shot learning. *Comput. Surv.* 53, 3 (2020), 63:1–63:34. DOI: <https://doi.org/10.1145/3386252>

- [140] P. Wei, Z. Ning, S. Ye, L. Sun, F. Yang, K. C. Wong, D. Westerdahl, and P. K. K. Louie. 2018. Impact analysis of temperature and humidity conditions on electrochemical sensor response in ambient air quality monitoring. *Sensors* 18 (Jan. 2018), 16. DOI : <https://doi.org/10.3390/s18020059>
- [141] WHO. 2016. *Ambient Air Pollution: A Global Assessment of Exposure and Burden of Disease*. Technical Report. World Health Organization. Retrieved from <https://www.who.int/phe/publications/air-pollution-global-assessment/en/>.
- [142] Y. Xiang, L. S. Bai, R. Piedrahita, R. P. Dick, Q. Lv, M. Hannigan, and L. Shang. 2012. Collaborative calibration and sensor placement for mobile sensor networks. In *Proceedings of the 2012 ACM/IEEE 11th International Conference on Information Processing in Sensor Networks (IPSN'12)*. IEEE, 73–83. DOI : <https://doi.org/10.1109/IPSN.2012.6920952>
- [143] K. Yan, L. Kou, and D. Zhang. 2018. Learning domain-invariant subspace using domain features and independence maximization. *IEEE Trans. Cybernet.* 48, 1 (2018), 288–299. DOI : <https://doi.org/10.1109/TCYB.2016.2633306>
- [144] K. Yan and D. Zhang. 2016. Calibration transfer and drift compensation of e-noses via coupled task learning. *Sens. Actuat. B: Chem.* 225 (2016), 288–297. DOI : <https://doi.org/10.1016/j.snb.2015.11.058>
- [145] W. Y. Yi, K. M. Lo, T. Mak, K. S. Leung, Y. Leung, and M. L. Meng. 2015. A survey of wireless sensor network based air pollution monitoring systems. *Sensors* 15, 12 (2015), 31392–31427. DOI : <https://doi.org/10.3390/s151229859>
- [146] H. Yu, Q. Li, Y. Geng, Y. Zhang, and Z. Wei. 2020. AirNet: A calibration model for low-cost air monitoring sensors using dual sequence encoder networks. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI'20)*, Vol. 34, 1. AAAI Press, Palo Alto, CA, 1129–1136. DOI : <https://doi.org/10.1609/aaai.v34i01.5464>
- [147] H. Yu, Q. Li, R. Wang, Z. Chen, Y. Zhang, Y. A. Geng, L. Zhang, H. Cui, and K. Zhang. 2020. A deep calibration method for low-cost air monitoring sensors with multilevel sequence modeling. *IEEE Trans. Instrument. Meas.* 69, 9 (2020), 7167–7179. DOI : <https://doi.org/10.1109/TIM.2020.2978596>
- [148] M. A. Zaidan, N. Hossein Motlagh, P. L. Fung, D. Lu, H. Timonen, J. Kuula, J. V. Niemi, S. Tarkoma, T. Petäjä, M. Kulmala, and T. Hussein. 2020. Intelligent calibration and virtual sensing for integrated low-cost air quality sensors. *IEEE Sens. J.* 20, 22 (2020), 13638–13652. DOI : <https://doi.org/10.1109/JSEN.2020.3010316>
- [149] J. Zhang, Z. Tang, M. Li, D. Fang, P. Nurmi, and Z. Wang. 2018. CrossSense: Towards cross-site and large-scale WiFi sensing. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking (MobiCom'18)*. Association for Computing Machinery, New York, NY, 305–320. DOI : <https://doi.org/10.1145/3241539.3241570>
- [150] X. Zhang, X. Chen, and X. Zhang. 2018. The impact of exposure to air pollution on cognitive performance. *Proc. Natl. Acad. Sci. U.S.A.* 115, 37 (2018), 9193–9197. DOI : <https://doi.org/10.1073/pnas.1809474115>
- [151] T. Zheng, M. H. Bergin, K. K. Johnson, S. N. Tripathi, S. Shirodkar, M. S. Landis, R. Sutaria, and D. E. Carlson. 2018. Field evaluation of low-cost particulate matter sensors in high- and low-concentration environments. *Atmos. Meas. Techn.* 11, 8 (2018), 4823–4846. DOI : <https://doi.org/10.5194/amt-11-4823-2018>
- [152] Y. Zheng, F. Liu, and H. Hsieh. 2013. U-air: When urban air quality inference meets big data. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'13)*. ACM, New York, NY, 1436–1444. DOI : <https://doi.org/10.1145/2487575.2488188>
- [153] F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, Y. Zhu, H. Xiong, and Q. He. 2021. A Comprehensive survey on transfer learning. *Proceedings of the IEEE* 109, 1 (2021), 43–76. DOI : [10.1109/JPROC.2020.3004555](https://doi.org/10.1109/JPROC.2020.3004555)
- [154] N. Zimmerman, A. A. Presto, S. P. N. Kumar, J. Gu, A. Hauryliuk, E. S. Robinson, A. L. Robinson, and R. Subramanian. 2018. A machine learning calibration model using random forests to improve sensor performance for lower-cost air quality monitoring. *Atmos. Meas. Techn.* 11, 1 (2018), 291–313. DOI : <https://doi.org/10.5194/amt-11-291-2018>

Received August 2020; revised November 2020; accepted December 2020