



This is a repository copy of *Modelling calibration uncertainty in networks of environmental sensors*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/202440/>

Version: Accepted Version

Article:

Smith, M.T., Ross, M., Ssematimba, J. et al. (3 more authors) (2023) Modelling calibration uncertainty in networks of environmental sensors. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 72 (5). pp. 1187-1209. ISSN 0035-9254

<https://doi.org/10.1093/jrssc/qlad075>

This is a pre-copyedited, author-produced version of an article accepted for publication in *Journal of the Royal Statistical Society Series C: Applied Statistics* following peer review. The version of record Michael Thomas Smith, Magnus Ross, Joel Ssematimba, Mauricio A Álvarez, Engineer Bainomugisha, Richard Wilkinson, *Modelling calibration uncertainty in networks of environmental sensors*, *Journal of the Royal Statistical Society Series C: Applied Statistics*, 2023, qlad075 is available online at: <https://doi.org/10.1093/jrssc/qlad075>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Modelling calibration uncertainty in networks of environmental sensors

Michael Thomas Smith¹ | Magnus Ross² | Joel Ssematimba³ | Mauricio A Álvarez² | Engineer Bainomugisha³ | Richard Wilkinson⁴

¹Department of Computer Science, University of Sheffield, UK

²Department of Computer Science, University of Manchester, UK

³Department of Computer Science, Makerere University, Kampala, Uganda

⁴School of Mathematical Sciences, University of Nottingham, UK

Correspondence

Michael Thomas Smith, Department of Computer Science, University of Sheffield, Sheffield, S1 4DP, UK
Email: m.t.smith@sheffield.ac.uk

Funding information

Google.org.
EPSRC GCRF grant, EP/T00343X/1.

Networks of low-cost environmental sensors are becoming ubiquitous, but often suffer from poor accuracies and drift. Regular colocation with reference sensors allows recalibration but is complicated and expensive. Alternatively the calibration can be transferred using low-cost, mobile sensors. However inferring the calibration (with uncertainty) becomes difficult. We propose a variational approach to model the calibration across the network. We demonstrate the approach on synthetic and real air pollution data, and find it can perform better than the state of the art (multi-hop calibration).

In Summary: The method achieves uncertainty-quantified calibration, which has been one of the barriers to low-cost sensor deployment.

KEYWORDS

air pollution, Bayesian modelling, calibration, Gaussian processes, low-cost sensors, variational inference

1 | INTRODUCTION

Large networks of low-cost environmental sensors are becoming increasingly common across a range of domains, including weather (van de Giesen et al., 2014), snow-cover (Pohl et al., 2014), wildlife (Dyo et al., 2010) and air pollution (Khedo et al., 2010; Liu et al., 2020). The aim is to collect far more data at low-cost. Unfortunately, across all these domains is the common issue of how to ensure low-cost sensors remain calibrated, while still achieving substantial

cost-savings. In particular, the problem of maintaining calibration over time, in low cost sensor networks, typically requires regular colocation recalibration (Rai et al., 2017).

Considerable work has already been invested in developing approaches for the reliable use of low-cost sensors for air pollution monitoring (see Section 1.2). Rai et al. (2017) emphasise that the key issue with the use of low-cost sensors is data quality: For the network to have some use for policy makers the sensors need frequent calibration in 'final deployment conditions' such that sensor calibration drift can be accounted for.

This work is motivated by the need to manage the calibration of a network of low-cost PM2.5, air pollution sensors deployed in Kampala, Uganda (PM2.5 is particulate pollution with particle diameter less than $2.5\mu\text{m}$). Kampala is an East African city, associated with a tropical, high-humidity, dusty environment; leading to additional issues with sensor reliability and drift. Calibration in this network is achieved through regular but brief visits to the static (low-cost and reference) sensors, by mobile low-cost sensors mounted on motorbike taxis. The motorbike taxi visits are often opportunistic and usually last less than ninety minutes. This leads to a complex network of in-situ colocation events. Our experience is that drift (due to dust, etc) experienced in Kampala is much greater than similar sensors running in the UK, suggesting that this should be included in any model. We found that most of the research on pollution monitoring were for cities in temperate climates. Very few environmental monitoring networks are on the continent of Africa (Mao et al., 2019). Hopefully the approach described will help guide future deployments across the continent.

1.1 | The task

We have a set of particulate (PM2.5) sensors. Some are mobile and some are static. Some of the static sensors are 'reference' instruments which we assume measure the true pollution. The rest are low-cost sensors. We assume that the true pollution is a parametric function, ϕ (raw measurement, parameters), of the raw observation from the low-cost sensor. The task is to estimate this calibration function (or rather its vector of parameters) for each sensor, for any given time, to allow us to use the low-cost sensors to estimate the true air pollution in locations where there is no reference instrument.

The full problem consists of both estimating the calibration functions and modelling the pollution over space and time across the city. Modelling these two aspects jointly may have advantages, but we focus in this paper on simply solving the calibration problem, where each sensor-pair colocation occurs at an arbitrary, undefined spatial location.

For a given sensor, s at time, t , we wish to find the value of the calibration function's parameter vector, $f_s(t)$. We will estimate this using a list of colocations. Each colocation records: (a) The time of the colocation (b) The id of the two sensors (c) The raw (uncalibrated) measurements of the two sensors. To illustrate with an example, the calibration function could be a simple linear regression with two parameters (an offset and a gradient), such that: true pollution = raw measurement \times gradient + offset. The gradient and offset make up a vector of two parameters, which we wish to estimate. These we believe can change over time, and differ between sensors. We want a method that: (1) Can infer the calibration function's parameters even if the associated sensor has not been directly visited by a reference sensor. (2) Can combine calibration information arriving through multiple pathways of calibration. (3) Uses both static and mobile sensor data. (4) Has a reasonable model of how the measurement is related to the true pollution. (5) Provides principled estimates of the uncertainty in the calibration, taking into account time.

1.2 | Background

1.2.1 | In-situ Calibration

A review by Delaine et al. (2019) divides the calibration literature in several ways, considering reference instrumentation, sensor mobility, types of calibration (time dependency and complexity) and 'grouping strategies'. Calibration can occur prior to deployment or can occur in-situ by finding those times in which pollution is believed to be similar across large distances, for example when few local sources of pollution exist at night (e.g. Miskell et al., 2018). We are interested in those approaches which use physical post-deployment colocation.

Consider first Arfire et al. (2015). They have a model-based approach which assumes a low-cost chemical sensor is calibrated by co-location with a reference sensor regularly. This doesn't then handle the more complex network that we are analysing (in which colocations between low-cost sensors also occurs). They consider temporal drift by either a linear or a hyperbolic term with respect to time. There is no indication that they compute or assess the uncertainty in the model's predictions.

Of more interest to our application are studies looking at networks of sensors such that the low-cost sensors are not necessarily directly calibrated by the reference instrument but potentially by intermediate sensors (e.g. mobile sensors). Tsujita et al. (2004) considers one such model. However they make a series of assumptions (e.g. how the sensor calibration drift is a fixed linear number). The calibration is also done in a non-probabilistic way; with an estimate for the pollution at a location simply the average of all sensors, even though some will, presumably, be more reliable than others (e.g. if they have just visited the reference station). The model also doesn't report its confidence. Setting aside these issues, the paper does suggest that a network approach to calibration might be of benefit. Xiang et al. (2012) use the uncertainty in a given sensor's calibration when combining measurements. The problem this paper considers largely matches ours. One major difference is the error model. In our experience the calibration we are concerned with involves scaling of the true pollution, not just an offset. One potentially could alter their paper by using the log of sensor measurements, which will then make their sum a product. An advantage of their approach though is that the calibration computation and data remains local to the device, while we are computing this using all the data from all sensors. Markert et al. (2018) assign weights to various estimates but, as with almost all the papers, the correlations between the calibration estimates are not considered. Their treatment is partly Bayesian (for example in how different uncertainties are combined), but has a single value for drift (subtracting the calibrated reading from the raw reading). Also by only looking at previous colocations, it fails to propagate all the calibration information available. We suggest therefore that there are likely to be considerable improvements in calibration accuracy and uncertainty quantification by using a process model of calibration over time, which will be able to leverage prior knowledge around temporal structure in calibration functions. A fully Bayesian treatment will also be able to handle more complex measurement functions. These criticisms aside, Markert et al. (2018) provide a useful foundation and demonstration of the benefits of a network of sensors. Another related work (Hasenfratz et al., 2012) looks at a similar problem, this time using a polynomial fit for the sensor data (rather than just adding an offset). They handle multiple sensors by computing weighted sums, with weights dependent on the time difference (between query and calibration) although not in a particularly probabilistic manner. The paper computes a single maximum likelihood (ML) estimate for the calibration.

Most other papers (e.g. Maag et al., 2017; Kizel et al., 2018; Bychkovskiy et al., 2003; Fonollosa et al., 2016) do not consider drift over time. Many, (e.g. Maag et al., 2017; Bychkovskiy et al., 2003) do not handle uncertainty in their calibration estimates. Finally note that most of the papers mentioned in this section assume linear/scaling-only calibration. In summary, there is currently no framework for performing calibration across an arbitrary network of sensor-rendezvous events which probabilistically handles the quantification of uncertainty in the calibration with

temporal drift. The method described in this paper correctly estimates the uncertainty in calibration across a complex network consisting of mobile, static, low-cost and reference instruments.

1.2.2 | Related work in probabilistic modelling

We did originally explore a joint model, in which both the calibration parameters and the pollution were modelled as Gaussian processes. This had potentially more power as non-colocated (but still correlated) measurements could still inform the estimates of the calibration, but it was found to be somewhat intractable (see Section 3.1.1 for a discussion), and hence we focus in this paper on the calibration pair model (in which we treat the colocations independently of their spatial locations) and leave the joint model for future work.

It is however worth considering the necessary tools for inference in the joint model, as it connects with other literature. The joint model involved a calibration function parameterised by a Gaussian process (modelling, e.g. the calibration offset) operating on another Gaussian process (the true pollution over time and space). This is reminiscent, if the calibration function was simply scaling the pollution, of Wilson et al. (2012), in which a weighted sum of Gaussian processes (GPs) is computed, with the weights consisting of GPs themselves. They experimented with both Markov chain Monte Carlo (MCMC) (using elliptical slice sampling) and variational Expectation Maximization (EM) to perform the inference on the latent variables. We experimented with using MCMC for inference in both models but found mixing was challenging as the network grew larger. Wilson et al. (2012) noted that even with their relatively simple model, Gibbs sampling would 'mix poorly because of the tight correlations between the weights and the nodes'. We found that even with whitening, Hamiltonian Monte Carlo (HMC) and other tricks, the strong correlations in the posterior made it very difficult to perform inference. Alvarado and Stowell (2017) more recently computed a similarly structured product of Gaussian processes, and used Gauss-Hermite quadrature. In our experiments we used simple Monte Carlo to compute the expectations, following the approaches in Hensman et al. (2013) and Salimbeni and Deisenroth (2017).

1.3 | Calibration pair model

We will describe the approach and algorithm in detail next, but we first give a brief description of our Bayesian calibration approach. We assume that a sensor's calibration function can be expressed using a set of parameters (e.g. the offset, scaling etc) and then assume that these parameters change slowly over time. To model this we assign Gaussian process priors to these parameters (these allow us to specify, mathematically, these smoothness assumptions). The observations our model is provided with consist of measurements from pairs of sensors, of (what is assumed to be) the same pollution concentration, by visiting the same locations at the same time. Some of the sensors are reference sensors (that provide ground truth observations). To infer the processes that describe the calibration function parameters, we consider each pair of observations: For each sensor measurement we apply the sensor's calibration function, using the values for the parameters for that sensor at that time, sampled from the appropriate process. If the predicted pollution values from the two sensors are similar then the 'likelihood' is high. Bayesian inference effectively balances the likelihood of the observations with how well our calibration function parameters match the prior. To actually compute this we use a stochastic variational inducing point approximation to the Gaussian processes, and then iteratively improve this approximation. This will be explained further below.

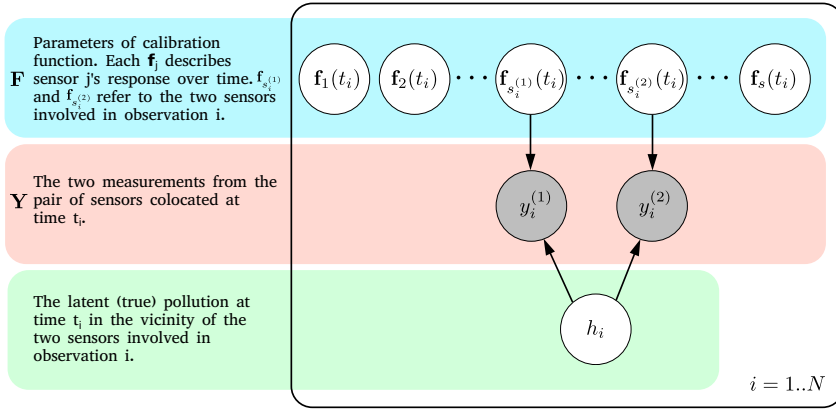


FIGURE 1 Calibration pair model. *A priori* the vectors of GP random variables in \mathbf{F} are independent, but become correlated in the posterior due to collocated observations connecting them together in pairs.

1.3.1 | Model definition

Each of the $i = 1..N$ observations, $\mathbf{Y} \in \mathcal{R}^{N \times 2}$, consist of two measurements, $y_i^{(1)}, y_i^{(2)}$, by two sensors with indices $s_i^{(1)}, s_i^{(2)}$, collocated at time t_i . One could, in principle, have more than two sensors collocated, but for simplicity in representation and likelihood function we have used pair-wise collocation. The model assumes that a deterministic parametric function, ϕ , describes how a measurement is related to the true pollution: true pollution = $\phi(\text{raw measurement, parameters}) + \text{unstructured noise}$.

For each sensor, j , the calibration function's parameter vector consists of C parameters that are each functions of time: $\mathbf{f}_j(t)$. An example might be that, the first parameter is the scaling; the second, the offset; and the third an offset proportional to the humidity. Our model assigns an independent Gaussian process prior (over time) to each of these $c = 1..C$ parameters for each of the $j = 1..S$ sensors. So, let $\mathbf{F}(t)$ be an $S \times C$ matrix of independent Gaussian processes $[f_j]_c(t) \sim GP(0, k_{j,c})$: We are assuming, *a priori*, that the calibration parameters of the sensors are independent between sensors and the parameters are also independent within sensors. In our implementation we have different kernels for sensors of different types and for different parameters, allowing us to model, for example, that one sensor type's offset drifts at a slower rate than its scaling, etc. The aim of the method is to infer the values of $\mathbf{F}(t)$.

We emphasise that, in this model we are discarding location information and just using information about which sensor pairs are collocated, and when.

1.3.2 | The likelihood

Here, by likelihood, we mean the probability of the observations, \mathbf{Y} , given the model, its hyper-parameters and the matrix of latent variables, \mathbf{F} , which describe the calibration function parameters. To compute $\log p(\mathbf{Y}|\mathbf{F})$ we first assume that this can be factorised across observations: $p(\mathbf{Y}|\mathbf{F}) = \prod_{i=1}^N p(\mathbf{y}_i|\mathbf{F})$. Where $p(\mathbf{y}_i|\mathbf{F})$ is the probability of one pair of collocated observations, given our model and latent calibration parameters. We pick from \mathbf{F} the two rows associated with the two sensors $s_i^{(1)}$ and $s_i^{(2)}$ that took the two measurements. So the calibrated predicted pollution from the two observations are, $\phi(y_i^{(1)}, \mathbf{f}_{s_i^{(1)}}(t_i))$ and $\phi(y_i^{(2)}, \mathbf{f}_{s_i^{(2)}}(t_i))$.

Let h_i be the (latent) true pollution associated with the collocated observation pair \mathbf{y}_i . We assume that these

calibration-corrected predictions of the pollution are normally distributed around the true (latent) pollution, i.e.:

$$p\left(\begin{bmatrix} \phi(y_i^{(1)}, \mathbf{f}_{s_i^{(1)}}(t_i)) \\ \phi(y_i^{(2)}, \mathbf{f}_{s_i^{(2)}}(t_i)) \end{bmatrix} \middle| h_i\right) \sim \mathcal{N}\left(\begin{bmatrix} h_i \\ h_i \end{bmatrix}, \begin{bmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{bmatrix}\right). \quad (1)$$

If we place a wide, fairly uninformative, normal prior on h_i (with mean zero, and variance γ^2), we can integrate out h_i in closed form. Allowing us to write the probability of a given (calibrated) observation pair as a 2d joint multivariate Gaussian,

$$p\left(\begin{bmatrix} \phi(y_i^{(1)}, \mathbf{f}_{s_i^{(1)}}(t_i)) \\ \phi(y_i^{(2)}, \mathbf{f}_{s_i^{(2)}}(t_i)) \end{bmatrix}\right) = \int p\left(\begin{bmatrix} \phi(y_i^{(1)}, \mathbf{f}_{s_i^{(1)}}(t_i)) \\ \phi(y_i^{(2)}, \mathbf{f}_{s_i^{(2)}}(t_i)) \end{bmatrix} \middle| h_i\right) \times \mathcal{N}(h_i | 0, \gamma^2) dh_i \quad (2)$$

$$\sim \mathcal{N}\left(0, \begin{bmatrix} \sigma^2 + \gamma^2 & \gamma^2 \\ \gamma^2 & \sigma^2 + \gamma^2 \end{bmatrix}\right). \quad (3)$$

For inference using Hensman et al. (2013) we need an expression for the conditional probability of the observations \mathbf{y} given the latent calibration \mathbf{f} . So we need to perform a change of variable, to covert from the distribution of $p(\phi(\mathbf{y}, \mathbf{f}))$ to $p(\mathbf{y} | \mathbf{f})$. To do this we multiply the probability of a vector of calibrated observations, ϕ , by the determinant of the vector's Jacobian with respect to \mathbf{y} (which is diagonal in this case), evaluated at the two observations.

$$p\left(\begin{bmatrix} y_i^{(1)} \\ y_i^{(2)} \end{bmatrix} \middle| \mathbf{F}\right) = p\left(\begin{bmatrix} \phi(y_i^{(1)}, \mathbf{f}_{s_i^{(1)}}(t_i)) \\ \phi(y_i^{(2)}, \mathbf{f}_{s_i^{(2)}}(t_i)) \end{bmatrix}\right) \times \frac{\partial \phi(z_1, \mathbf{f}_{s_i^{(1)}}(t_i))}{\partial z_1} \bigg|_{z_1=y_i^{(1)}} \times \frac{\partial \phi(z_2, \mathbf{f}_{s_i^{(2)}}(t_i))}{\partial z_2} \bigg|_{z_2=y_i^{(2)}} \quad (4)$$

For implementation, we use the log probabilities, and so the product becomes the sum of the log probability of ϕ and the log of the two partial derivatives.

Reference Instruments We define a binary vector r , which describes which sensors are reference sensors. In those cases, we use the identity function instead of the calibration function, i.e. replacing $\phi(y_i^{(a)}, \mathbf{f}_{s_i^{(a)}}(t_i))$ with $y_i^{(a)}$.

1.3.3 | Variational inference for the calibration pair model

We wish to compute the calibration parameters from the observations, $p(\mathbf{F} | \mathbf{Y})$. A simple application of Bayes rule, $p(\mathbf{F} | \mathbf{Y}) = p(\mathbf{Y} | \mathbf{F})p(\mathbf{F})/p(\mathbf{Y})$, is not soluble or tractable in closed form, as the number of observations may be large and the likelihood function is non-Gaussian. We quickly step through the variational approach, largely based on Hensman et al. (2013). We approximate the posterior, $p(\mathbf{F} | \mathbf{Y})$ with a variational distribution over \mathbf{F} : $q(\mathbf{F})$. We want to make this distribution similar to the true posterior. To this end we aim to minimise the Kullback-Leibler (KL)

divergence between the two:

$$\mathcal{D}_{KL}[q(\mathbf{F})||\rho(\mathbf{F}|\mathbf{Y})] = - \int q(\mathbf{F}) \log \frac{\rho(\mathbf{F}|\mathbf{Y})}{q(\mathbf{F})} d\mathbf{F} \quad (5)$$

$$= - \int q(\mathbf{F}) \log \frac{\rho(\mathbf{F}, \mathbf{Y})}{\rho(\mathbf{Y})q(\mathbf{F})} d\mathbf{F} \quad (6)$$

$$= - \int q(\mathbf{F}) \log \frac{\rho(\mathbf{F}, \mathbf{Y})}{q(\mathbf{F})} d\mathbf{F} + \int q(\mathbf{F}) \log \rho(\mathbf{Y}) d\mathbf{F} \quad (7)$$

$$= - \underbrace{\int q(\mathbf{F}) \log \frac{\rho(\mathbf{F}, \mathbf{Y})}{q(\mathbf{F})} d\mathbf{F}}_{\text{The ELBO}} + \log \rho(\mathbf{Y}) \quad (8)$$

The last term is constant wrt the variational distribution so we can minimise the KL divergence simply by maximising the integral (the evidence lower bound, ELBO). This is still not tractable without the approximation provided by inducing points. We introduce an additional vector \mathbf{u} of values that we assume can describe \mathbf{F} , evaluated at inducing point locations, \mathbf{Z} . Each row of \mathbf{Z} has a time, and a sensor id/parameter id, to index the associated GP in \mathbf{F} . We found that spacing the inducing points evenly over the sensor domains, rather than being optimised works well for the datasets we experimented with. We can compute the posterior over both \mathbf{F} and \mathbf{u} (we assume that \mathbf{Y} is conditionally independent of \mathbf{u} given \mathbf{F}), $\rho(\mathbf{F}, \mathbf{u}|\mathbf{Y}) \propto \rho(\mathbf{Y}|\mathbf{F})\rho(\mathbf{F}|\mathbf{u})\rho(\mathbf{u})$. We don't compute this directly, but instead use our approximation, $q(\mathbf{F}, \mathbf{u}|\mathbf{Y}) = \rho(\mathbf{F}|\mathbf{u})q(\mathbf{u})$. Here we have assumed that \mathbf{u} is a sufficient statistic to determine \mathbf{F} so that $\rho(\mathbf{F}|\mathbf{u}, \mathbf{Y}) = \rho(\mathbf{F}|\mathbf{u})$. We substitute this approximation into the ELBO and augment with the inducing point values,

$$\mathcal{L} = \int \int q(\mathbf{F}, \mathbf{u}) \log \frac{\rho(\mathbf{F}, \mathbf{u}, \mathbf{Y})}{q(\mathbf{F}, \mathbf{u})} d\mathbf{F} d\mathbf{u} \quad (9)$$

$$= \int \int q(\mathbf{F}, \mathbf{u}) \log \frac{\rho(\mathbf{Y}|\mathbf{F})\rho(\mathbf{F}|\mathbf{u})\rho(\mathbf{u})}{\rho(\mathbf{F}|\mathbf{u})q(\mathbf{u})} d\mathbf{F} d\mathbf{u} \quad (10)$$

$$= \int \int q(\mathbf{F}, \mathbf{u}) \log \rho(\mathbf{Y}|\mathbf{F}) d\mathbf{F} d\mathbf{u} + \int \int q(\mathbf{F}, \mathbf{u}) \log \frac{\rho(\mathbf{u})}{q(\mathbf{u})} d\mathbf{u} d\mathbf{F} \quad (11)$$

The second term's integral over \mathbf{F} integrates out, leaving:

$$\mathcal{L} = \int \int q(\mathbf{F}, \mathbf{u}) \log \rho(\mathbf{Y}|\mathbf{F}) d\mathbf{F} d\mathbf{u} + \int q(\mathbf{u}) \log \frac{\rho(\mathbf{u})}{q(\mathbf{u})} d\mathbf{u} \quad (12)$$

$$= \mathbb{E}_{q(\mathbf{F}, \mathbf{u})} \left[\log \rho(\mathbf{Y}|\mathbf{F}) \right] - \mathcal{D}_{KL} \left[q(\mathbf{u}) || \rho(\mathbf{u}) \right] \quad (13)$$

We now maximise \mathcal{L} wrt the parameters of the variational distribution, described by the mean, \mathbf{m} and triangular matrix \mathbf{R} (covariance = $\mathbf{R}\mathbf{R}^\top$), using stochastic gradient descent. Specifically we approximate the expectation by sampling from $q(\mathbf{F})$ and then computing the mean of the log likelihoods for all the samples. $q(\mathbf{u})$ and $\rho(\mathbf{u})$ are Gaussian so the second term can be computed in closed form. We use TensorFlow's automatic differentiation to maximise \mathcal{L} wrt the variational parameters, \mathbf{m} and \mathbf{R} using this stochastic approximation to the true gradient (stochastic gradient descent). This formulation also allows us to use minibatches, if N becomes intractably large. See Algorithm 1 in the supplementary, for the overview of the computation. Finally we can use the estimates of the variational parameters to predict $q(\mathbf{F})$ for a test point at time t_* for all sensors.

1.3.4 | Optimising hyperparameters

To avoid excessive model complexity we used just four hyperparameters: the lengthscale of static sensors; the lengthscale of mobile sensors; likelihood Gaussian noise; and the scale for all kernels. One could treat these as random variables and integrate over them as with other variables. However, a common approach when using Gaussian processes is to find a point estimate either by maximising the marginal likelihood, or in our case, by optimising an error metric in held out data. We used Bayesian optimisation (using the GPyOpt library, GPyOpt authors, 2016) to find such a point estimate, which we used in a model tested on new data. Although we are not now reflecting the uncertainty in the hyperparameters, this point estimate means that the model when deployed in a data processing pipeline, is quicker and more robust. More importantly our lengthscale can reflect our prior belief that a sensor might degrade quickly even if it hasn't yet done so yet in the training data. If such a failure doesn't yet exist in the dataset, the model is liable to select very long lengthscales. By manually selecting lengthscales we can incorporate our domain knowledge in sensor quality and performance.

1.3.5 | Sampling

For the synthetic data example below, we found the algorithm typically optimised the calibration between low-cost sensors first, then would very slowly move these distribution means to match the reference sensors. This behaviour is inevitable due to the highly correlated posterior and the imbalance in samples (with most being between non-reference pairs of instruments). Many elements of the approximating distribution mean must all move together in precisely the right direction, to increase the fit to the reference sensor data. As a side note, this might suggest parametrisations of the approximating distribution's mean vector should be investigated that might allow this manipulation using fewer variables. This problem is compounded by the relatively infrequent reference sensor observations: Many mini-batches contain few or no samples from the reference instrument co-locations. The gradient from the reference sensors only directly applies to those sensors that were co-located with it, while the co-locations of these sensors with the other non-reference collection of sensors will exert a gradient in the opposing direction.

We found for some experiments that importance sampling (Csiba and Richtárik, 2018) largely solved this issue and led to fast, reliable optimisation. We simply oversampled co-locations involving the reference sensors and corrected for the biased sampling by adjusting the noise variance appropriately.

We did not find this a problem in the experiment with real data from Kampala. This might be due to the shallow nature of the network, making it easier for the reference co-locations to influence the variational distribution across the whole network. Part of our preprocessing of the real data involved averaging samples of co-located observations (blocks of 10 observations, typically over 15 minutes, were averaged), such averaging is standard practice in the field (e.g. Stedman et al., 2006). This is necessary due to the considerable noise associated with individual observations (each averaging over only 80 seconds). This may have led to some balancing between the reference-sensors (which make measurements averaging over longer periods) and other sensors, which could explain why importance sampling was unnecessary for the real dataset.

1.3.6 | Other variables

Optical Particle Counter (OPC) PM2.5 sensors are sensitive to environmental factors. Specifically, a high relative humidity can lead to large biases (e.g. Crilley et al., 2018, in particular figure 6). In this framework such variables could be included by simply adding additional columns to \mathbf{Y} . These can then be passed to the calibration function, ϕ , and

involved in the parametric expression that computes the ‘true’ pollution.

1.3.7 | Algorithm

The above steps of the variational calibration pair model algorithm, and how to compute predictions using it, are detailed in the supplementary, in Algorithms 1 and 2 respectively.

To summarise, the inference algorithm involves finding an appropriate mean vector m and covariance matrix RR^T that describe the approximating Gaussian distribution. During each step of optimisation, one iterates over the colocated observations, and for each, one samples values of the calibration parameters for the two sensors involved, from the approximating posterior distribution. One then computes the log likelihood of the observation given the calibration values from each of the samples, and then averages to get the expected log likelihood. Finally the ELBO is computed by summing over these expected log likelihoods from the observations and subtracting the KL divergence between the approximating Gaussian distribution and the prior. An automatic differentiation framework is used to find the gradients of the ELBO with respect to our approximating distribution parameters and these are used to optimise those parameters (using Adam). To make a prediction, one samples the relevant calibration parameter values from the approximating posterior, and for each sample compute the predicted true pollution, given the sensor’s observation and the sample of the calibration. The mean and standard deviation of these predictions of the true pollution could then be computed, for example.

1.4 | Multi-hop instant calibration

For comparison with our calibration pair method we developed and implemented an algorithm loosely based on the [Multi-hop Instant Calibration](#) method from Hasenfratz et al. (2012) and the [rendezvous connection graph](#) from Saukh et al. (2015). The former paper is a little unclear about combining weights, while the latter one handles sensor drift by only using observations within each time window. In Saukh et al. (2015) a shortest path algorithm is applied (personal communication with authors) to convert the network to an acyclic graph.

We augment the original graph in Saukh et al. (2015) by adding edges between the same sensor, in neighbouring time windows. The intuition being that the scaling information provided by a long series of hops within the window might be less accurate than a direct reference calibration ‘carried over’ from a different time window. The edges connecting time windows can have a different weight than those connecting sensors in the same time period. The choice of weight ratio (i.e. time-edge-weight/colocation-edge-weight) will decide whether the model leans towards using longer chains of colocations in the same time window (high ratio) vs relying on colocations that occurred a long time ago (low ratio). Figure 2 illustrates.

The implementation of the graph building and prediction are detailed in the supplementary, in Algorithms 3 and 4 respectively. To summarise, one first builds a graph that connects all sensors that have (at least 5) colocation observations during a given time window. Each edge has a weight of 1 and an associated value of the average log ratio between the two observations during the colocations. This is repeated for all time windows in the dataset, and edges are added between the same sensor in neighbouring time windows. These have a weight that we need to select (see Figure 2) and a value of zero. Once the graph is constructed, Dijkstra algorithm is used to find the shortest path from all sensors (in all time windows) to any reference sensor (in any time window), and the sum of the values along the edges of each path are entered as that sensor’s log-scaling factor for that time window. A prediction is made by simply exponentiating the sensor’s log-scaling factor and multiplying it by the raw observation.




FIGURE 2 Multi-hop: The ratio of edge weights for neighbouring times (green) over weights for colocated sensors (blue) affects the choice of shortest path to a reference instrument. A low ratio will lead to shorter chains of colocation but relying on colocations from different time windows (dashed grey path), a high ratio will lead to a preference for colocations from nearby times, but with longer chains of colocations pairs (solid grey path).




FIGURE 3 Synthetic sensor example: Colocation visits over a representative 350 hours of the total 4320 hours: Visits to the reference (blue) and low-cost (green) static sensors by the mobile ones (red), indicate by vertical black lines.

2 | RESULTS

Because the multi-hop approach was limited to only scaling, we also restrict the calibration-pair model to a simple scaling factor in the experiments in this paper, to make a fair comparison. It is important to remember that the calibration-pair model can easily incorporate more complex calibration functions. We have included an example in the supplementary that demonstrates a ‘scaling plus offset’ calibration function, which shows in more detail how the posterior is able to represent the joint uncertainty in the two parameters. We briefly discuss the choice of calibration function in Section 3.1.2.

The *test* datasets in the Kampala examples either consist of a synthetically scaled reference-station’s measurements (in Section 2.2.3), or its exact values, but using a co-located low-cost sensor to generate the predictions (in Section 2.2.4). Note that the training data’s network of stations must always include a reference station (providing true observations).

2.1 | Synthetic example

2.1.1 | The synthetic data

Before considering real data we explored the method’s response to a representative, but synthetic, dataset. We simulated ten static and four mobile sensors over a 180 day period. Four of the static sensors are assumed to provide

noise-free, unbiased, reference measurements, the rest have a time-varying bias and added Gaussian white-noise with variance of $100 (\mu\text{g}/\text{m}^3)^2$. In the Kampala data discussed later, each mobile sensor is typically localised to one part of the city, with only occasional visits further afield. We simulate this by making the probability of a visit by a mobile sensor to a static sensor, proportional to the inverse distance between a 'home' and the static sensor.

The synthetic data was engineered specifically to provide low-quality, noisy, infrequent and challenging data, in an attempt to reflect the problems that are typical in real data. Figure 3 illustrates how few visits were made to the reference instruments, over a representative sample of the synthetic data.

Simple sinusoidal functions were used to generate both the true pollution and the scaling of the mobile and static sensors. The period of the sinusoids was different for each sensor, but the static sensors had a longer period on average (median = 3070 hours vs 1007 hours). This variation was to (a) emphasise the heterogeneous nature of the sensors and (b) illustrate how the mobile sensors are probably less stable. Figure 4 shows the scalings actually used.

2.1.2 | Variational pair model

We applied the method described in Section 1.3 to the synthetic data. As mentioned in Section 1.3.4, we used Bayesian optimisation (30 iterations, optimising the Normalised Mean Squared Error, NMSE) to select the model hyperparameters for the variational calibration pair method (kernel variance, likelihood variance and two lengthscales). We optimised using the synthetic data with added noise scale of $10\mu\text{g m}^{-3}$ for all the results. The synthetic data was resampled each iteration of the Bayesian optimisation (and for testing): the sensor locations, collocation events and noise were all resampled; but the frequencies of the synthetic calibration sinusoids remained the same.

Even though each sensor's scaling function has a different period, we use just two lengthscales to model the static and mobile sensors. On the first dataset generated, the Bayesian optimisation chose 2201 hours and 581 hours for the lengthscales of the static and mobile sensors, respectively. This likely reflects the underlying generating function, but also trades off the ability to model the short lengthscales of some of the sensors with the utility of long lengthscales which allow more data to be used for any given prediction. It also permits calibration information to be passed over longer periods of time. The likelihood variance and kernel scale (variance) were $6.02 (\mu\text{g}/\text{m}^3)^2$ and $7.64 (\mu\text{g}/\text{m}^3)^2$ respectively. These reflect the variance of the scaling factor, so it is hard to compare directly to the noise added to the data. The choice of likelihood variance seems to overestimate the true value a little. It may be that the effect of selecting a longer lengthscale than some of the static sensors' true lengthscales leads to some of their inaccuracy being 'explained away' by this added noise term.

2.1.3 | Multi-hop calibration and comparison

We also applied the multi-hop calibration method to the synthetic data. We first performed a grid search to optimise the window size and the ratio of edge weights between those edges representing time and those representing a collocation (again, performed on separately generated training data). We unfortunately found this model performed very poorly on this data. The best configurations usually had very long window sizes, leading to little or no ability to model the variation over time. We found that this configuration marginally improved on simply using the raw measurements. See Table 1.

The estimates from the variational method provided considerable improvements over the multi-hop approach. We can also compute the negative log predictive as we predict a distribution. Figure 4 shows these predictions for three of the six non-reference static sensors and two of the four mobile, for both the multi-hop method and the calibration pair model. One can see some errors in the calibration pair model. In particular, the high frequency drift

Noise Scale / $\mu\text{g m}^{-3}$	NMSE			NLPD
	No Method	Multi-hop	Cal. Pair Model	Cal. Pair Model
2	0.77 (0.02)	0.27 (0.01)	0.18 (0.01)	42.6 (2.8)
5	0.78 (0.02)	0.44 (0.03)	0.21 (0.01)	49.5 (2.3)
10	0.97 (0.02)	0.87 (0.07)	0.35 (0.01)	58.2 (2.3)
20	1.46 (0.02)	3.18 (0.69)	0.84 (0.02)	87.7 (1.5)

TABLE 1 NMSE using raw data, the multi-hop calibration and the variational calibration pair model, for synthetic data with four levels of added noise. Bracketed values are the standard error of the mean based on predictions generated from 10 synthetic datasets (for each, the multi-hop was optimised on the test data, while the calibration pair model was optimised once on training data with noise scale = $10\mu\text{g m}^{-3}$). NMSE, Normalised Mean Squared Error; NLPD, Negative Log Predictive Density.

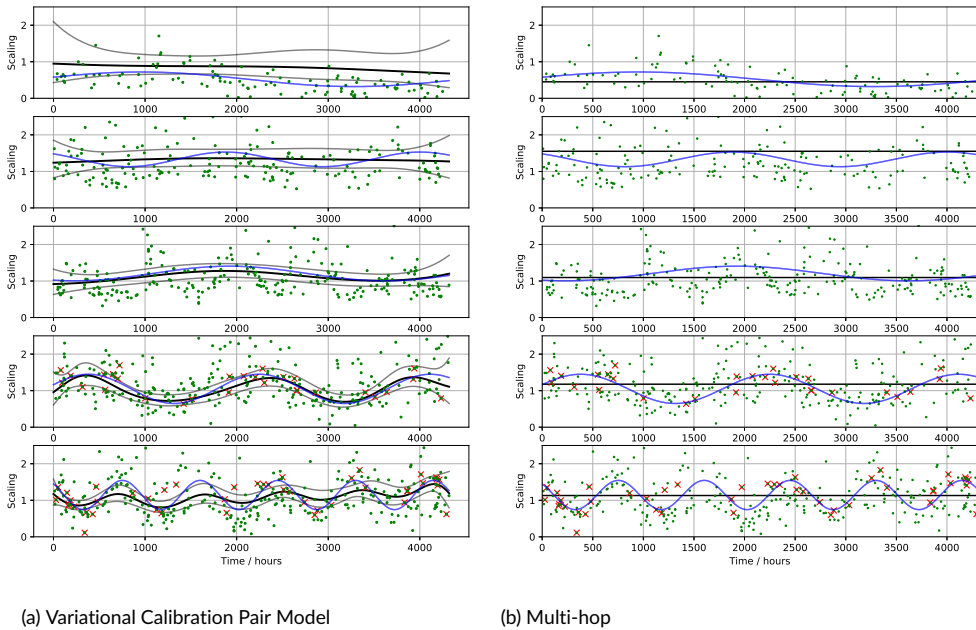


FIGURE 4 Predictions for the scaling in the synthetic sensor example for both the Variational Pair Model and the Multi-hop method, showing five of the ten low-cost sensors. Noise scale = $10\mu\text{g m}^{-3}$. Upper three plots in each, static sensors. Lower two plots, mobile sensors. Black line, posterior mean. Grey lines, 95% of the posterior density. Blue line, true scaling of each sensor. Green dots, ratio of colocation observation pairs involving this sensor. Red crosses, ratio of colocations observations between this sensor and a reference sensor.

of the low-cost mobile sensor has not been properly detected by the model in all the low-cost sensors, leading to this drift appearing in the predictions for the other sensors.

To explore the effectiveness of the two models, we added different amounts of noise to the synthetic data, and

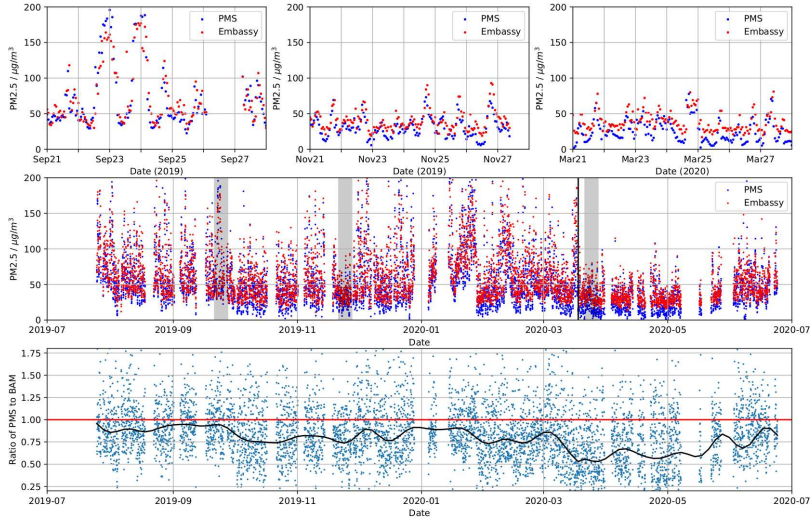


FIGURE 5 Upper plots, enlargements of three weeks (highlighted in grey in the middle plot), showing correlation between the low cost sensor and the reference instrument. Middle plot, the full period of measurements. Discontinuity on 18th March, 2020, due to the covid-19 lock down (vertical black line). Lower plot, the ratio of the two sensors. A GP with an exponentiated quadratic (EQ) kernel has been fitted to the log ratio and the exponent of its posterior mean plotted (pre- and post- lock down modelled with independent GPs). Times are in UTC.

reran the analysis, although we kept the variational calibration method's hyperparameters as those trained with a noise scale of $10 \mu\text{g m}^{-3}$, to see how well it adapts to misspecified noise. Table 1 reports the NMSE for synthetic data with four added noise-scales for the two methods. The multi-hop method is allowed to optimise its parameters for each noise scale and, as might be expected, its window size reduces as the noise is reduced (for noise scales of 2, 5, 10 and $20 \mu\text{g m}^{-3}$ the selected window sizes were 292, 1184, 3010 and 4800 hours, respectively). We can also see that the accuracy of the multi-hop method approaches the calibration pair model's, as the noise is reduced. This is probably because the calibration pair model makes use of multiple pathways through multiple colocations and over time to support a prediction, whilst the multi-hop method is constrained to use a single path through the network for a given prediction, making it more vulnerable to the effect of noise. The use of multiple pathways in the pair model gives it considerable robustness against noisy data.

2.2 | Kampala example

2.2.1 | The Kampala data

Recently a network of low-cost particulate air pollution sensors has been deployed across Kampala. These consist of pairs of PMS5003 (Plantower) optical particle counters (OPCs) mounted at fixed, static locations and on motorbike taxis, known as boda bodas (Figure 6b). Before discussing the network we will briefly inspect the measurements of a static OPC mounted 10m from the US Embassy's regularly calibrated Beta Attenuation Monitor (BAM), which provides hourly ground truth measurements of PM2.5 pollution. This brief comparison will give some insight into the calibration problem: What the calibration function might consist of, and how it might vary over time.

Figure 5 illustrates the strong correlation between the two sensors, but also highlights some differences. The

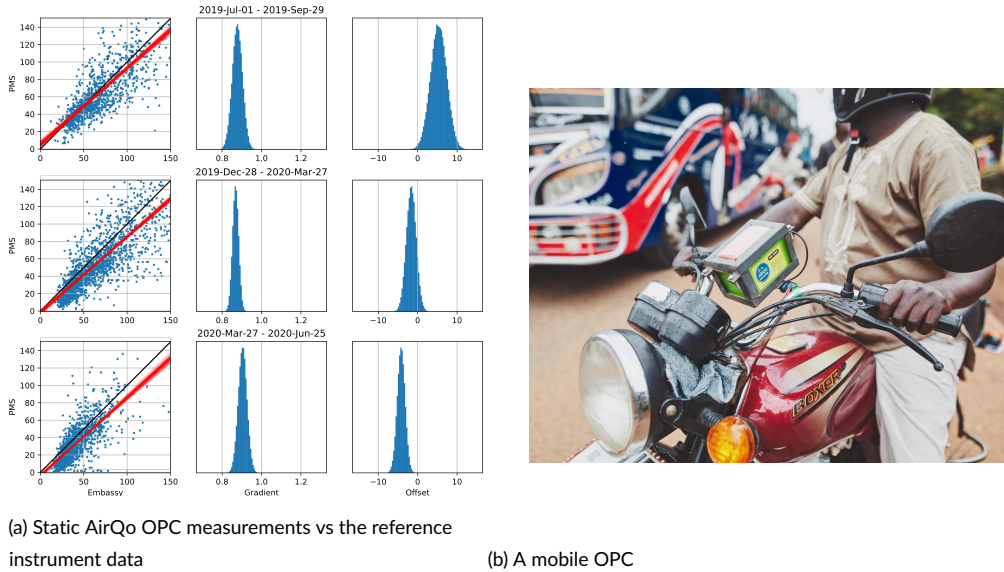


FIGURE 6 (left) AirQo OPC measurements vs the reference instrument at the US embassy. Each row corresponds to 90 days of measurements. The left scatter plot shows the raw connection between the two sensors. A black line indicates equality while red lines are sampled from the linear fit's posterior probability distribution. Middle and right plots are the gradient and offset marginal posterior distributions respectively. The gradient distribution remains roughly constant, while the offset appears to 'drift' downwards over time. (right) One of the mobile low-cost sensors mounted on the front of a motorbike taxi ('boda boda'), photographed in Kampala by AirQo (www.airqo.net).

earlier plot from September seems to show very similar measurements, with a ratio near one, on average. But later measurements seem to have a greater discrepancy with the low cost OPC sensor significantly underestimating the true pollution. The ratio itself is not enough to explain this change. Figure 6a illustrates how an offset can explain some of the difference. In all three the gradient of the linear fit appears to remain between 0.85 and 0.95, however the offset maybe drifts from positive to negative, explaining the increasing disparity in the previous graphs.

2.2.2 | Calibration over the network

51 sensors were in the network in the second-half of 2020, consisting of 2 reference sensors (BAM 1022 Particulate Monitor, Met One), 4 mobile OPCs and 45 static OPCs (PMS5003, plantower). Between the 15th July, 2020 and 3rd February, 2021 (203 days) there were 6,118,977 measurements of PM_{2.5} recorded. We use this dataset for the analysis. We built a dataset of 433,935 'encounters', consisting of every pair of measurements that occurred within 40m and 30 minutes of each other. We averaged ten minute periods and removed observations outside the 10 – 300 $\mu\text{g}/\text{m}^3$ range, as these we found to sometimes be associated with faults. This left 40,432 records (approximately 200/day). The majority of these are associated with permanently co-located pairs of static sensors. In Figure 7, (a) shows the physical locations of the sensors and (b) the number of colocations between them. The two reference sensors '-24516' (on Makerere University Campus) and '-24517' (in Nakawa), are approximately 5km apart.

To test the models we only defined -24516 as being a reference sensor and let the models estimate the calibration required of -24517. We gave the models all the colocation data (none was held out) as we are interested in the

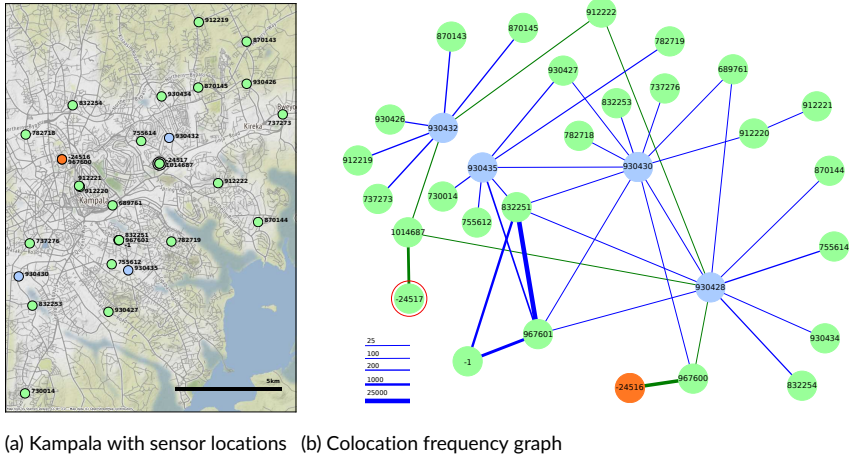


FIGURE 7 The AirQo air pollution monitoring network (sensors that have been colocated). The four blue nodes are the mobile sensors. The orange node is the reference sensor at Makerere University, the circled node is the reference sensor whose calibration we are testing. (a) Locations of sensors, the latest locations of mobile sensors are plotted, with one not in the city limits. (b) The network of colocations between sensors (from July 2020 to February 2021). The line thicknesses indicates the number of colocations in that time. The green lines indicate two possible paths from the reference to the test sensor.

prediction of the calibration scaling for the test sensor. There are co-location observations at the test sensor over a period of 69 days from the 203 days total.

We would expect the test sensor to have a calibration scaling of exactly 1.0 as it's a high quality, carefully maintained and calibrated BAM sensor. Our Gaussian process model has a prior mean of the log-ratio of zero, which is exactly a gain of 1.0, while the multi-hop model doesn't have a prior. To avoid our prior giving an unfair advantage and to simulate the effect of a drifting calibration, as is seen in the low-cost OPCs (Section 2.2.1), we artificially scaled the measurements of the reference sensor to slowly increase over time (starting at $1.0\times$ on the first day of observations at the test sensor, and increasing to $1.83\times$ by the 69th day). In summary: The algorithms being tested must reconstruct this drift using the network of colocations from the known reference sensor. For our calibration pair model we select hyperparameter values that we believe are appropriate for this dataset (lengthscale for static and mobile sensors = 100 days, EQ kernel variance = $9 (\log(\mu\text{g}/\text{m}^3))^2$, Bias kernel variance = $3 (\log(\mu\text{g}/\text{m}^3))^2$, likelihood (ratio) noise variance = $0.2 (\log(\mu\text{g}/\text{m}^3))^2$). This approach is necessary, as discussed previously, to allow us to include our prior knowledge around lengthscales associated with sensor drift and failure. We again compare to the multihop method described in Section 1.4. For comparison with our method, we optimised the parameters of the competing multi-hop method (the window size and weight ratio between time and colocation) on the dataset we also test with. Effectively giving an upper bound on its capability.

2.2.3 | Kampala Synthetic-Drift Results

After running the grid search optimisation, the optimum multi-hop configuration had a window size of just 24 hours and an edge weight ratio of 0.21 (it roughly equally weights a colocation with a reference sensor from 5 days earlier and via an intermediate sensor during the same 24 hours). We computed error metrics between the test values of the reference sensor and the calibration predictions. To review the aim: The two calibration methods used the

Metric	No calibration	Multi-hop calibration	Calibration Pair Model	
MAE	12.71	8.67	5.545 (0.47)	$\mu\text{g m}^{-3}$
NMSE	0.321	0.180	0.078 (0.011)	

TABLE 2 Kampala experimental results, with synthetic drift added. Using no calibration, the multi-hop calibration and the pair model (bracketed values are standard error from 20 runs. Not required for multi-hop as it's deterministic). MAE, Mean Absolute Error; NMSE, Normalised Mean Squared Error.

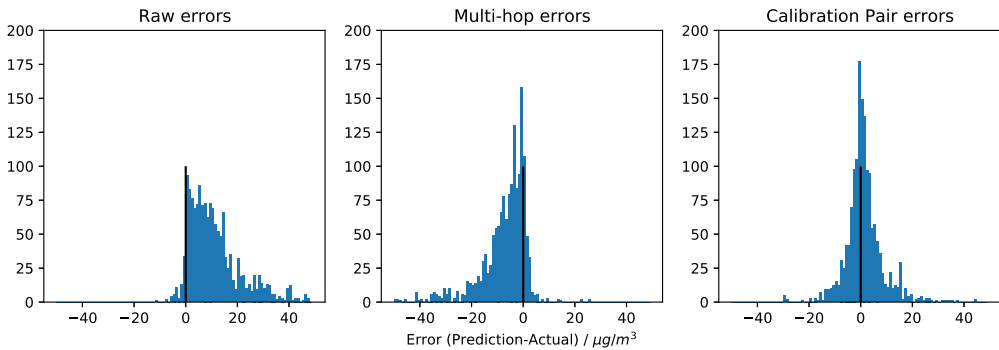



FIGURE 8 The distribution of errors in pollution predictions at the test sensor. As we give the three models the measured pollution, in principle the error could be zero if the model correctly estimates the calibration. The left graph is the distribution of predictions if we just use the measurements as the predictions. The added, synthetic, drift causes this to over-estimate. The central graph is using the multi-hop approach, this appears to underestimate. The right graph is for the calibration pair model. Runs of the variational inference optimisation end with slightly varying results here. This is a typical example (with a MAE about the average).

other reference instrument and the network of colocations to estimate the calibration of the test instrument (with the added synthetic drift). Table 2 outlines the results for these approaches, while Figure 8 shows the distribution of errors from the different approaches. The 'raw measurement' errors are caused by our artificial scaling (without this the raw measurements would have no error).

We find the multi-hop method does help improve the accuracy of the predictions to an extent, but the variational method performs considerably better. Besides accuracy, the variational pair-model method provides uncertainty estimates for its predictions, and the associated 95% confidence intervals include the true scaling function. Variational inference is associated with underestimates of the variance of the posterior (Blei et al., 2017) so this result is reassuring. The co-location data is very noisy. To illustrate this, the green markers in Figure 9 are the ratios between the test sensor and other sensor(s) that visit. Note that we are not assuming that these are just noisy samples from the true latent scaling - as these OPC sensors will themselves have some scaling associated with their estimates. But plotting them gives the reader a sense as to the scale of the noise in the dataset.

2.2.4 | Real Calibration

As a final brief experiment we tried calibrating the low cost OPC sensor (1014687) that is colocated with the test reference sensor, then using the calibration combined with the OPC observations to predict the values measured by the test sensor. The sensor was new and we only had 69 days of data associated. We found the actual scaling was



(a) Multi-hop scaling predictions

(b) Variational pair model scaling predictions

FIGURE 9 The predictions of the scaling parameter at the test sensor. Red line, our synthetically generated scaling drift we are trying to predict. Black line, the prediction. Green dots, ratios between the test sensor and other sensor(s) that visit. The variational prediction posterior mean varied quite a lot per run of the algorithm, probably due to high uncertainty, leaving the choice of mean difficult to determine.

the OPC was underestimating the reference sensor by only about 6% which equates to an error of about $3 \mu\text{g}/\text{m}^3$. The raw data had a MAE of $15.08 \mu\text{g}/\text{m}^3$ while the multi-hop and variational pair model approaches had a MAE of 16.07 and 14.87 respectively. Neither approach really had much effect on the MAE. The earlier demonstration with the more significant artificial scaling was however modelled successfully, giving us confidence in the approach.

3 | DISCUSSION

In this paper we've proposed a method to calibrate a network of sensors using observations consisting of pairs of potentially bias or transformed observations in which a few observers are assumed to provide ground truth. We applied this to our main dataset - of static and mobile air pollution sensors in Kampala, adding a synthetic drift term to illustrate the potential issues in longer term data. We found the calibration pair method models the drift appropriately, making the predictions far more accurate. It also, importantly, provides uncertainty estimates on the calibration. For the real Kampala data, without the synthetic drift added, we found (in Section 2.2.4) it had only a 1.4% improvement in the MAE (over simply using the raw data values, and an 8% improvement on the multi-hop approach). Potentially over longer periods, we might expect the approach to become more relevant (as demonstrated by its utility when the drift was added).

Comparing the results of the multi-hop method and the variational calibration pair method, it seems that, for larger networks, the latter does better on both the synthetic and real data. It was somewhat surprising that it had much of an advantage in the Kampala case (especially with the synthetic drift added). In simulations we typically found the performance was considerably better than the multi-hop method when there were many sources of data to integrate. In the synthetic dataset for example, there were four reference sensors. So, with only one reference sensor, we anticipated the two approaches to perform similarly. However not only is the number of reference sensors relevant here, but also the number of paths between the reference sensor and the test sensor. Figure 7b indicates two routes exist from the reference sensor to the test sensor in the real dataset. Finally it isn't simply the path from the reference sensor to the test sensor that supports the variational result. The network graph doesn't show the full complexity of data availability over time, but importantly we envisage that the other sensors visited provide support. By also being visited, they can help constrain the calibration, indirectly.

3.1 | Calibration in context

Many studies exist in which low-cost air pollution sensors are co-located with reference sensors, and calibration functions are then derived to, in principle, allow the low-cost sensor to then be used without the reference sensor (Lewis et al., 2016; Liu et al., 2017; Zimmerman et al., 2018; Barcelo-Ordinas et al., 2018; Crilley et al., 2018; Badura et al., 2019; Wang et al., 2019; Datta et al., 2020; Lee et al., 2020; Ferrer-Cid et al., 2020). This can also extend to calibrating remote observation data (for example Shaddick et al., 2018). It has been noted many times that on-site field calibration is better than lab-based (Rai et al., 2017), and that sensor and concept drift due to ageing, contamination or changes in the underlying pollution (e.g. due to changes in the source or environmental) means that regular recalibration is necessary (Crilley et al., 2018; Marathe et al., 2021; Castell et al., 2017, Section 3.2.3; and the results in this paper, Section 2.2.1). Rai et al. (2017) wrote,

...it is necessary to perform a field calibration for each sensor individually. Moreover, the calibration parameters might change over time depending on the meteorological conditions and the location, i.e., once the nodes are deployed it will be difficult to determine if they are under- or over-estimating the pollutant concentrations.

Although it is widely agreed regular field recalibration is essential, the complexity of achieving this makes it potentially the main barrier to large-scale, low-cost sensor deployment (Rai et al., 2017). Several attempts have been made to look at how to do field-recalibration using mobile reference sensors (Hasenfratz et al., 2012; Arfire et al., 2015). But using low-cost mobile sensors to act as 'intermediaries' (Tsuji et al., 2004; Hasenfratz et al., 2012) is far more practical. The low-cost sensors (e.g. the plantower pms5003) are typically smaller and often will be more resilient to being moved, and can be replaced more cheaply if damaged. In the AirQo Project in Kampala, moving the reference instruments regularly would have been very challenging. An alternative would be to swap the static low-cost sensors between reference collocation sites and their static locations. However such access required site permissions and expertise, while the mobile low-cost sensors could visit the static stations simply by the motor-bike taxi drivers parking next to the sensor. This simplicity provided by permanently mobile low-cost sensors is discussed in Zhao et al. (2021).

It is arguably only worth deploying, calibrating and maintaining the low-cost sensors if they can provide useful information to policy makers. Castell et al. (2017) suggest that they are limited to non-policy/science applications. They however looked at data from a city with relatively low-pollution. The average PM_{2.5} in Kampala was roughly 30-60 $\mu\text{g}/\text{m}^3$. The ten minute-averaged sample error in our experiment for the low cost sensor (with a calibration computed across the network) was about 15 $\mu\text{g}/\text{m}^3$. Once averaged to daily or annual estimates this would almost certainly be reduced below the 50% threshold in the EU Data Quality Objectives (DQO) (for daily PM_{2.5} estimates). (Janssen et al., 2017). This is independent of the calibration approach and is probably possible in Kampala simply due to the high background pollution, reflecting the findings of Castell et al. (2017), who found their relative accuracy improved at those times and locations with most pollution. This suggests achieving the EU's DQOs may be easier in high-pollution environments. It would be interesting to consider whether, in high-pollution locations, low-cost sensors could provide estimates of sufficient accuracy that they can be used for policy-making decisions. In-field calibration would still be necessary. Borrego et al. (2018) provide further support for the potential use of low-cost sensors. They found that, with on-site calibration some types of low-cost sensors (although not PM_x) were able to reach the appropriate EU DQOs, and note that this poses 'a considerable challenge in terms of method scalability.' Hopefully our approach can support this scaling challenge. An important feature of our method is that, by using a probabilistic approach to calibration and spatiotemporal modelling, one should also be able to state where (and when)

in the domain the model reaches sufficient confidence.

A related field is the detection of sensor failure in a network of sensors. Miskell et al. (2016) used proxies (reference stations in an area with a similar land use) to detect potential sensor failure. Our method might be useful in the detection of drift, or similar signs of hardware failure. For example combining our method for calibration with a spatiotemporal model with leave one-sensor-out cross-validation. Poor prediction accuracy may be due to poor modelling assumptions or inference, e.g. we haven't included some aspect of the local environment; sensor failure, e.g. the inlet is clogged; or training data errors, e.g. if neighbouring sensors are malfunctioning. For more nuance it might be better to follow other approaches to fault detection. Peng et al. (2017), for example, distinguish between types of fault.

Combining sensor estimates using a Bayesian approach is not novel. Talampas and Low (2012) take into account each sensor's uncertainty to combine their measurements using a Bayesian approach, and find it performs much better than simply averaging. In a different field Perälä et al. (2020) use Gaussian processes for another type of calibration problem: combining expert estimates, potentially with biases, using a hierarchical GP. This could be applicable, in particular if different types of sensor can be grouped.

3.1.1 | Remote Calibration

Calibration (usually) requires that the two sensors are observing the same level of pollution. In this paper and discussion we have assumed that this meant visiting the same location at the same time. However, if the pollution in two locations is likely to be very similar, then sensors in those two locations could also be calibrated. Tsujita et al. (2005) identify times (and locations) in which NO₂ might be similar (low values, to allow baseline calibration). A similar approach is taken in Mueller et al. (2017, also for NO₂) and in Moltchanov et al. (2015, multiple pollutants) in which low concentrations at night are considered to be similar and allow calibration to take place. These three papers used this period for calibrating the offset (as they were calibrating at low concentrations). This raises two issues: first, whether a simple locally-generated zero-pollution source could instead be used to determine the zero-offset, rather than relying on low-pollution periods in the atmosphere; second is that this doesn't support the calibration of other parameters (the gradient, or other polynomial weights, for example) which require higher concentrations.

We briefly mentioned in Section 1.1 that 'the full problem consists of both estimating the calibration functions and modelling the pollution over space and time...Modelling these two aspects jointly may have advantages.' We didn't explore in this paper the potential capabilities of the joint model (in which the pollution is modelled explicitly in time and space), instead we focused on the 'pair model' (in which the pollution remains a hidden, latent random variable) as we found it, in unreported work, more reliable. Specifically, the joint model had high correlations in the posterior, mediated via both the prior and each co-location, via the likelihood function. To explain: We have a latent process describing the true pollution over space and time. We have additional processes that each describe components (offset, scaling, etc) of the calibration function of each sensor. If we consider two different sensor observations, these are a priori independent but are coupled via two different likelihoods that combine the latent modelled spatiotemporal pollution concentrations at each of those times/locations and each of the sensors' calibration components. One way to consider the problem is the number of ways the same observations could be explained: This joint model has far more model freedom. The pair model of this paper simplifies this challenge and was found to scale and behave in a far more robust manner, with very reliable optimisation (when using VI).

If the joint model were made to work, a key advantage would be the potential for remote calibration: if there were two non-colocated sensors which the model has confidence are being exposed to similar levels of pollution (e.g. if they are both located by the same road), then this joint-model will be able to use their observations to continuously

calibrate them, as long as the spatio-temporal model is well specified.

The combination of this paper's probabilistic calibration framework, probabilistic spatio-temporal models of pollution and the remote-calibration described in (e.g. Tsujita et al., 2005) could be an extremely powerful way to maintain and quantify calibration across the network; minimising the need for mobile sensors and colocation.

3.1.2 | Calibration Function

We have not discussed particularly the choice of calibration function. This can vary from a simple scaling parameter, to the complexities of a neural network (e.g. Spinelle et al., 2014; De Vito et al., 2020). In the examples in this paper we use a simple scaling function to allow us to compare to the multi-hop calibration approach. However our method can use other (parametric) functions. We have included, in the supplementary, a toy-example to demonstrate a scaling-plus-offset calibration function. Our method does have a limit to the number of parameters it could model. In unpublished work, we consider a similar, categorical calibration problem. In that, each 'sensor' had over a dozen parameters and the tool still was able to perform successful inference. However, the number of parameters in even a 'shallow neural network' (as in De Vito et al., 2020) may be challenging. Also if there are multiple ways the model can explain the observation, the posterior needs to be able to account for increasing correlation between the parameters (the supplementary Figure 14 shows this for even the simple scaling-plus-offset example). Finally, the calibration function in our model can include environmental inputs. We have not exploited this feature, but this is often an input to the calibration function (for example Crilley et al., 2018; De Vito et al., 2020), and something to consider in future work.

3.2 | The problem of spatially structured bias

Sousan et al. (2016) showed how different sources of particulate pollution lead to different responses by optical particle counters. Castell et al. (2017) and Crilley et al. (2018) also discuss this issue. Sandradewi et al. (2008) propose that one deliberately use the change in the optical properties of the pollution to infer its source. We have assumed that the calibration function remains the same over space, and similar over short periods of time. If the former assumption doesn't hold we will need to colocate reference instrument at each low-cost sensor location. If the latter assumption doesn't hold, then we can't use low-cost sensors at all. Chu et al. (2020) suggest one could model the former, spatial-non stationarity. Presumably the temporal non-stationarity could be addressed in a similar way. They were investigating PM_{2.5} in a high humidity climate, so potentially applicable to our work in Kampala. Low cost OPC sensors are affected by high humidity, so the variation in humidity over Taiwan lead to changing biases. We could use these relevant features (such as humidity) in the calibration function. Unfortunately the type of pollution (road dust, diesel exhaust or charcoal burning) isn't so easily measured, so some sort of spatio-temporal component might need to be introduced to the calibration function.

In the worst case we could imagine that one part of the city has pollution that is underestimated by all the network's low-cost sensors (e.g. a source such as burning tyres), thus all calibration that occurs between low-cost sensors in that area will fail to detect this bias. The only solution is to colocate a reference instruments in those locations. Carefully distributing the small number of reference instruments to sample a broad range of land-use types in the city is therefore prudent.

3.3 | Choice of model and choice of inference

We have (in unreported work) experimented with using MCMC to approximate the variational distribution (Hensman et al., 2015) but found it failed to achieve timely mixing (probably due to the correlations in the posterior) when the network was more than two or three edges deep, so we concentrated on the VI approach.

The variational calibration pair model was quite successful in modelling the posterior variance. One reason might be that approximating distribution included a full covariance matrix, a necessity due to the strong correlations between sensor parameters in the posterior. By avoiding the mean field approximation we might also have avoided some of the extreme underestimation of the variance (see Figure 1 in Blei et al., 2017, for an explanation). We did (unreported) experiment with using the mean-field approximation in the joint model and found the variance often ‘collapsed’, giving highly confident predictions, supporting this explanation.

One issue that might impact inference in larger networks is how the number of inducing points would roughly scale with time(T) \times sensors(S) \times parameters(C), as we assume the number of inducing points is proportional to time.

Using fixed hyperparameters for our model was driven partly by domain knowledge that past stability in calibration does not guarantee future stability, but we also found them somewhat difficult to estimate using black box variational inference, an experience shared by others, e.g. Nguyen and Bonilla (2014) report that they ‘have found this to be problematic, ineffectual, and time-consuming.’

The method allows the modeller to select what aspects of the model are shared amongst sensors: For example if it is believed that all the sensors will share the same parameter associated with humidity, then that component can be used by all the functions. If each sensor is independent but the calibration function components share the same hyperparameters then the same kernel can be used for the different models. By ensuring appropriate component sharing between sensors, the uncertainty can be reduced, accuracy increased and the model made more efficient (as fewer processes need to be approximated). This will depend on fabrication variance, deployment conditions and other, application specific factors.

3.4 | Heteroscedastic likelihood

In the likelihood used, the variance of the normal distribution was a fixed value. But some sensors (and locations) may have differing amounts of noise, and just as the calibration is expected to vary, so too will this likelihood variance. In unpublished work we have explored adding heteroscedastic noise (see also, for example Lázaro-Gredilla and Titsias, 2011). We found it was quite effective if, for the variational approximation of noise, one just use a Dirac distribution (rather than a multivariate Gaussian) and so just specify the mean. The motivation being that estimating the uncertainty in the noise is unrealistic. Our implementation allows either a fixed likelihood variance, a point estimate ‘Dirac’ approximation, or a full multivariate Gaussian distribution, but in the examples in the paper we used the fixed version, as it seemed to perform sufficiently for our datasets.

3.5 | Summary and future work

In summary, the calibration of a network through co-located observations of pairs of (potentially drifting) sensors can be performed by modelling each sensor’s calibration function parameters using one or more Gaussian processes, and providing a likelihood function computed for two observations. This provides a principled and robust Bayesian approach to handling this type of data. This will enable those designing sensor deployments to potentially use pair-

wise calibration as part of the design solution, in which a reference sensor does not need to visit every low-cost sensor directly.

To put this paper in context, Cui et al. (2021) describe a four stage calibration process for an air pollution monitoring system: (1 and 2) lab based, (3) field based reference sensor colocation (4) field based with occasional visits of mobile reference sensor. The calibration pair method would support the last of these four stages, allowing complex network structures, involving low-cost mobile sensors.

Future work includes optimising the choice of co-locations to minimise error: specifically choosing the best order to visit the static sensors in Kampala. There might be some sensors that are more useful to calibrate (due to their importance in the spatiotemporal model or location in the network). Cost is also a constraint: In our case, mobile-unit visits to the sensors had different (financial) costs, depending on their distance from the driver's base (known locally as their 'stage'). In particular the reference sensors were quite distant.

We do not consider uncertainty in the multi-hop model. We anticipate in future work looking at performing probabilistic inference over a simplified graph. We think this could be a good compromise (averaging over longer periods and making a relatively simplified graph network), potentially leading to faster inference than the variational Gaussian process approach, while still leading to good estimates of calibration and uncertainty, albeit in a more constrained modelling space.

The main limitation to future work in developing approaches towards networks of calibration the area of air pollution is that of data. Many low-cost network sensors exist (e.g. Parmar et al., 2017; Rai et al., 2017; Nyarku et al., 2018; Feinberg et al., 2019; Zhao et al., 2021) but few perform regular co-location calibrations. Possibly this is because there are no well developed, easily used, computational tools to perform the in-field calibration discussed in this paper. Hopefully methods, such as the one outlined in this paper, will help solve this chicken-and-egg problem.

The example networks used in this paper (both synthetic and real) were not very deep. We could envisage that, with larger datasets (for example those collected on www.zooniverse.org) this network could become more complex, both potentially benefiting more from this method, but also providing a more challenging inference problem as the number of GPs grows. One exciting possibility is to apply the method to the beewalk.org.uk project, a large, citizen-science initiative in which volunteers walk transects every month and count the number of each bumblebee species. If beewalkers were to occasionally walk other walkers' routes, could the data quality be improved through pair-wise calibration?

3.6 | Conclusion

The problem of calibrating low-cost, large-scale, air pollution sensor networks is quickly becoming acute: A number of large scale deployments of low-cost air pollution networks have happened recently, including one running on 260 cars in Beijing (Zhao et al., 2021). The difficulty in field-calibration and bounding the accuracy of estimates from these deployments suggests that methods that can provide automatic low-cost recalibration across the network are becoming increasingly necessary.

The calibration pair method described enables sensor deployments with complex networks of colocations. It allows the uncertainty in the predictions from these networks to be quantified, providing guaranteed confidence bounds and robust insights that mean the results from such networks can start being used by policy makers.

acknowledgements

This work was part funded by google.org and an EPSRC GCRF grant, 'Physically-informed probabilistic modelling of air pollution in Kampala using a low cost sensor network' [EP/T00343X/1]. We are grateful to Pablo A. Alvarado for his support at the early stages and walking through the derivation of stochastic VI in Alvarado and Stowell (2017)'s supplementary material.

conflict of interest

The authors declare that they have no conflict of interest, and that there is no financial interest to report.

code and data

The variational and multi-hop calibration methods are available as a python module which can be downloaded and installed from <https://github.com/lionfish0/calibration>, which includes a demonstration jupyter notebook, and the notebooks used to make the figures etc in this paper.

Kampala sensors: This data is owned and controlled by AirQo and is not available for public release. In particular, the raw boda boda data can reveal locations of individual journeys and the home locations of the taxi drivers. Sserunjogi et al. (2022) describe the AirQo data collection in detail.

references

- Alvarado, P. A. and Stowell, D. (2017) Efficient learning of harmonic priors for pitch detection in polyphonic music. *arXiv preprint arXiv:1705.07104*.
- Arfire, A., Marjovi, A. and Martinoli, A. (2015) Model-based rendezvous calibration of mobile sensor networks for monitoring air quality. In *2015 IEEE SENSORS*, 1–4. IEEE.
- Badura, M., Batog, P., Drzeniecka-Osiadacz, A. and Modzel, P. (2019) Regression methods in the calibration of low-cost sensors for ambient particulate matter measurements. *SN Applied Sciences*, **1**, 1–11.
- Barcelo-Ordinas, J. M., Garcia-Vidal, J., Doudou, M., Rodrigo-Muñoz, S. and Cerezo-Llavero, A. (2018) Calibrating low-cost air quality sensors using multiple arrays of sensors. In *2018 IEEE Wireless Communications and Networking Conference (WCNC)*, 1–6. IEEE.
- Blei, D. M., Kucukelbir, A. and McAuliffe, J. D. (2017) Variational inference: A review for statisticians. *Journal of the American statistical Association*, **112**, 859–877.
- Borrego, C., Ginja, J., Coutinho, M., Ribeiro, C., Karatzas, K., Sioumis, T., Katsifarakis, N., Konstantinidis, K., De Vito, S., Esposito, E. et al. (2018) Assessment of air quality microsensors versus reference methods: The EuNetAir Joint Exercise–Part II. *Atmospheric environment*, **193**, 127–142.
- Bychkovskiy, V., Megerian, S., Estrin, D. and Potkonjak, M. (2003) A collaborative approach to in-place sensor calibration. In *Information processing in sensor networks*, 301–316. Springer.
- Castell, N., Dauge, F. R., Schneider, P., Vogt, M., Lerner, U., Fishbain, B., Broday, D. and Bartonova, A. (2017) Can commercial low-cost sensor platforms contribute to air quality monitoring and exposure estimates? *Environment international*, **99**, 293–302.
- Chu, H.-J., Ali, M. Z. and He, Y.-C. (2020) Spatial calibration and PM_{2.5} mapping of low-cost air quality sensors. *Scientific reports*, **10**, 1–11.

- Crilly, L. R., Shaw, M., Pound, R., Kramer, L. J., Price, R., Young, S., Lewis, A. C. and Pope, F. D. (2018) Evaluation of a low-cost optical particle counter (alphasense OPC-N2) for ambient air monitoring. *Atmospheric Measurement Techniques*, **11**, 709–720.
- Csiba, D. and Richtárik, P. (2018) Importance sampling for minibatches. *The Journal of Machine Learning Research*, **19**, 962–982.
- Cui, H., Zhang, L., Li, W., Yuan, Z., Wu, M., Wang, C., Ma, J. and Li, Y. (2021) A new calibration system for low-cost sensor network in air pollution monitoring. *Atmospheric Pollution Research*, **12**, 101049.
- Datta, A., Saha, A., Zamora, M. L., Buehler, C., Hao, L., Xiong, F., Gentner, D. R. and Koehler, K. (2020) Statistical field calibration of a low-cost PM2.5 monitoring network in baltimore. *Atmospheric Environment*, **242**, 117761.
- De Vito, S., Di Francia, G., Esposito, E., Ferlito, S., Formisano, F. and Massera, E. (2020) Adaptive machine learning strategies for network calibration of iot smart air quality monitoring devices. *Pattern Recognition Letters*, **136**, 264–271.
- Delaine, F., Lebental, B. and Rivano, H. (2019) In situ calibration algorithms for environmental sensor networks: A review. *IEEE Sensors Journal*, **19**, 5968–5978.
- Dyo, V., Ellwood, S. A., Macdonald, D. W., Markham, A., Mascolo, C., Pásztor, B., Scellato, S., Trigoni, N., Wohlers, R. and Yousef, K. (2010) Evolution and sustainability of a wildlife monitoring sensor network. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, 127–140.
- Feinberg, S. N., Williams, R., Hagler, G., Low, J., Smith, L., Brown, R., Garver, D., Davis, M., Morton, M., Schaefer, J. et al. (2019) Examining spatiotemporal variability of urban particulate matter and application of high-time resolution data from a network of low-cost air pollution sensors. *Atmospheric environment*, **213**, 579–584.
- Ferrer-Cid, P., Barcelo-Ordinas, J. M., Garcia-Vidal, J., Ripoll, A. and Viana, M. (2020) Multisensor data fusion calibration in IoT air pollution platforms. *IEEE Internet of Things Journal*, **7**, 3124–3132.
- Fonollosa, J., Fernandez, L., Gutiérrez-Gálvez, A., Huerta, R. and Marco, S. (2016) Calibration transfer and drift counteraction in chemical sensor arrays using direct standardization. *Sensors and Actuators B: Chemical*, **236**, 1044–1053.
- van de Giesen, N., Hut, R. and Selker, J. (2014) The trans-African hydro-meteorological observatory (TAHMO). *Wiley Interdisciplinary Reviews: Water*, **1**, 341–348.
- GPYOpt authors (2016) GPYOpt: A Bayesian optimization framework in python. <http://github.com/SheffieldML/GPYOpt>.
- Hasenfratz, D., Saukh, O. and Thiele, L. (2012) On-the-fly calibration of low-cost gas sensors. In *European Conference on Wireless Sensor Networks*, 228–244. Springer.
- Hensman, J., Fusi, N. and Lawrence, N. D. (2013) Gaussian processes for big data. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*, 282–290.
- Hensman, J., Matthews, A. G., Filippone, M. and Ghahramani, Z. (2015) MCMC for variationally sparse Gaussian processes. In *Advances in Neural Information Processing Systems*, 1648–1656.
- Janssen, S., Guerreiro, C., Víaene, P., Georgieva, E. and Thunis, P. (2017) Guidance document on modelling quality objectives and benchmarking. https://fairmode.jrc.ec.europa.eu/document/fairmode/WG1/Guidance_MQO_Bench_vs2.1.pdf.
- Khedo, K. K., Perseedoss, R., Mungur, A. et al. (2010) A wireless sensor network air pollution monitoring system. *International Journal of Wireless & Mobile Networks*, **2**.
- Kizel, F., Etzion, Y., Shafran-Nathan, R., Levy, I., Fishbain, B., Bartonova, A. and Broday, D. M. (2018) Node-to-node field calibration of wireless distributed air pollution sensor network. *Environmental Pollution*, **233**, 900–909.
- Lázaro-Gredilla, M. and Titsias, M. K. (2011) Variational heteroscedastic Gaussian process regression. In *ICML*.

- Lee, H., Kang, J., Kim, S., Im, Y., Yoo, S. and Lee, D. (2020) Long-term evaluation and calibration of low-cost particulate matter (PM) sensor. *Sensors*, **20**, 3617.
- Lewis, A. C., Lee, J. D., Edwards, P. M., Shaw, M. D., Evans, M. J., Moller, S. J., Smith, K. R., Buckley, J. W., Ellis, M., Gillot, S. R. et al. (2016) Evaluating the performance of low cost chemical sensors for air pollution research. *Faraday discussions*, **189**, 85–103.
- Liu, D., Zhang, Q., Jiang, J. and Chen, D.-R. (2017) Performance calibration of low-cost and portable particular matter (PM) sensors. *Journal of Aerosol Science*, **112**, 1–10.
- Liu, X., Jayaratne, R., Thai, P., Kuhn, T., Zing, I., Christensen, B., Lamont, R., Dunbabin, M., Zhu, S., Gao, J. et al. (2020) Low-cost sensors as an alternative for long-term air quality monitoring. *Environmental research*, **185**, 109438.
- Maag, B., Zhou, Z., Saukh, O. and Thiele, L. (2017) Scan: Multi-hop calibration for mobile sensor arrays. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, **1**, 1–21.
- Mao, F., Khamis, K., Krause, S., Clark, J. and Hannah, D. M. (2019) Low-cost environmental sensor networks: recent advances and future directions. *Frontiers in Earth Science*, **7**, 221.
- Marathe, S., Nambi, A., Swaminathan, M. and Sutaria, R. (2021) Currentsense: A novel approach for fault and drift detection in environmental IoT sensors. In *Proceedings of the International Conference on Internet-of-Things Design and Implementation*, 93–105.
- Markert, J.-F., Budde, M., Schindler, G., Klug, M. and Beigl, M. (2018) Privacy-preserving collaborative blind macro-calibration of environmental sensors in participatory sensing. *EAI Endorsed Trans. Internet Things*, **18**, 1–8.
- Miskell, G., Salmond, J., Alavi-Shoshtari, M., Bart, M., Ainslie, B., Grange, S., McKendry, I. G., Henshaw, G. S. and Williams, D. E. (2016) Data verification tools for minimizing management costs of dense air-quality monitoring networks. *Environmental science & technology*, **50**, 835–846.
- Miskell, G., Salmond, J. A. and Williams, D. E. (2018) Solution to the problem of calibration of low-cost air quality measurement sensors in networks. *ACS sensors*, **3**, 832–843.
- Moltchanov, S., Levy, I., Etzion, Y., Lerner, U., Broday, D. M. and Fishbain, B. (2015) On the feasibility of measuring urban air pollution by wireless distributed sensor networks. *Science of The Total Environment*, **502**, 537–547.
- Mueller, M., Meyer, J. and Hueglin, C. (2017) Design of an ozone and nitrogen dioxide sensor unit and its long-term operation within a sensor network in the city of zurich. *Atmospheric Measurement Techniques*, **10**, 3783–3799.
- Nguyen, T. V. and Bonilla, E. V. (2014) Automated variational inference for Gaussian process models. *Advances in Neural Information Processing Systems*, **27**, 1404–1412.
- Nyarku, M., Mazaheri, M., Jayaratne, R., Dunbabin, M., Rahman, M. M., Uhde, E. and Morawska, L. (2018) Mobile phones as monitors of personal exposure to air pollution: Is this the future? *PLoS One*, **13**, e0193150.
- Parmar, G., Lakhani, S. and Chattopadhyay, M. K. (2017) An IoT based low cost air pollution monitoring system. In *2017 International Conference on Recent Innovations in Signal processing and Embedded Systems (RISE)*, 524–528. IEEE.
- Peng, Y., Qiao, W., Qu, L. and Wang, J. (2017) Sensor fault detection and isolation for a wireless sensor network-based remote wind turbine condition monitoring system. *IEEE Transactions on Industry Applications*, **54**, 1072–1079.
- Perälä, T., Vanhatalo, J. and Chrysafi, A. (2020) Calibrating expert assessments using hierarchical Gaussian process models. *Bayesian analysis*, **15**, 1251–1280.
- Pohl, S., Garvelmann, J., Wawerla, J. and Weiler, M. (2014) Potential of a low-cost sensor network to understand the spatial and temporal dynamics of a mountain snow cover. *Water Resources Research*, **50**, 2533–2550.

- Rai, A. C., Kumar, P., Pilla, F., Skouloudis, A. N., Di Sabatino, S., Ratti, C., Yasar, A. and Rickerby, D. (2017) End-user perspective of low-cost sensors for outdoor air pollution monitoring. *Science of The Total Environment*, **607**, 691–705.
- Salimbeni, H. and Deisenroth, M. (2017) Doubly stochastic variational inference for deep Gaussian processes. *Advances in neural information processing systems*, **30**.
- Sandradewi, J., Prévôt, A. S., Szidat, S., Perron, N., Alfarra, M. R., Lanz, V. A., Weingartner, E. and Baltensperger, U. (2008) Using aerosol light absorption measurements for the quantitative determination of wood burning and traffic emission contributions to particulate matter. *Environmental science & technology*, **42**, 3316–3323.
- Saukh, O., Hasenfratz, D. and Thiele, L. (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*, 274–285.
- Shaddick, G., Thomas, M. L., Green, A., Brauer, M., van Donkelaar, A., Burnett, R., Chang, H. H., Cohen, A., Dingenen, R. V., Dora, C. et al. (2018) Data integration model for air quality: a hierarchical approach to the global estimation of exposures to ambient air pollution. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, **67**, 231–253.
- Sousan, S., Koehler, K., Hallett, L. and Peters, T. M. (2016) Evaluation of the alphasense optical particle counter (OPC-N2) and the grimm portable aerosol spectrometer (pas-1.108). *Aerosol Science and Technology*, **50**, 1352–1365.
- Spinelle, L., Gerboles, M., Villani, M. G., Alexandre, M. and Bonavitaola, F. (2014) Calibration of a cluster of low-cost sensors for the measurement of air pollution in ambient air. In *SENSORS, 2014 IEEE*, 21–24. IEEE.
- Sserunjogi, R., Ssematimba, J., Okure, D., Ogenrwot, D., Adong, P., Muyama, L., Nsimbe, N., Bbaale, M. and Bainomugisha, E. (2022) Seeing the air in detail: hyperlocal air quality dataset collected from spatially distributed AirQo network. *Data in Brief*, **44**, 108512.
- Stedman, J., Abbott, J., Willis, P. and Bower, J. (2006) Review of background air quality data and methods to combine these with process contributions. *Environment Agency for England and Wales: Bristol*.
- Talampas, M. C. R. and Low, K.-S. (2012) Maximum likelihood estimation of ground truth for air quality monitoring using vehicular sensor networks. In *TENCON 2012 IEEE Region 10 Conference*, 1–6. IEEE.
- Tsujita, W., Ishida, H. and Moriizumi, T. (2004) Dynamic gas sensor network for air pollution monitoring and its auto-calibration. In *SENSORS, 2004 IEEE*, 56–59. IEEE.
- Tsujita, W., Yoshino, A., Ishida, H. and Moriizumi, T. (2005) Gas sensor network for air-pollution monitoring. *Sensors and Actuators B: Chemical*, **110**, 304–311.
- Wang, Y., Du, Y., Wang, J. and Li, T. (2019) Calibration of a low-cost PM2.5 monitor using a random forest model. *Environment international*, **133**, 105161.
- Wilson, A. G., Knowles, D. A. and Ghahramani, Z. (2012) Gaussian process regression networks. In *Proceedings of the 29th International Conference on International Conference on Machine Learning*, 1139–1146.
- Xiang, Y., Bai, L. S., Pledrahita, R., Dick, R. P., Lv, Q., Hannigan, M. and Shang, L. (2012) Collaborative calibration and sensor placement for mobile sensor networks. In *2012 ACM/IEEE 11th International Conference on Information Processing in Sensor Networks (IPSN)*, 73–83. IEEE.
- Zhao, B., Yu, L., Wang, C., Shuai, C., Zhu, J., Qu, S., Taiebat, M. and Xu, M. (2021) Urban air pollution mapping using fleet vehicles as mobile monitors and machine learning. *Environmental Science & Technology*, **55**, 5579–5588.
- Zimmerman, N., Presto, A. A., Kumar, S. P., Gu, J., Haurlyliuk, A., Robinson, E. S., Robinson, A. L. and Subramanian, R. (2018) A machine learning calibration model using random forests to improve sensor performance for lower-cost air quality monitoring. *Atmospheric Measurement Techniques*, **11**, 291–313.

Figure Captions

Figure 1:

Calibration pair model. *A priori* the vectors of GP random variables in F are independent, but become correlated in the posterior due to colocated observations connecting them together in pairs.

Figure 2:

Multi-hop: The ratio of edge weights for neighbouring times (green) over weights for colocated sensors (blue) affects the choice of shortest path to a reference instrument. A low ratio will lead to shorter chains of colocation but relying on colocations from different time windows (dashed grey path), a high ratio will lead to a preference for colocations from nearby times, but with longer chains of colocations pairs (solid grey path).

Figure 3:

Synthetic sensor example: Colocation visits over a representative 350 hours of the total 4320 hours: Visits to the reference (blue) and low-cost (green) static sensors by the mobile ones (red), indicate by vertical black lines.

Figure 4:

Predictions for the scaling in the synthetic sensor example for both the Variational Pair Model and the Multi-hop method, showing five of the ten low-cost sensors. Noise scale = $10 \mu\text{g m}^{-3}$. Upper three plots in each, static sensors. Lower two plots, mobile sensors. Black line, posterior mean. Grey lines, 95% of the posterior density. Blue line, true scaling of each sensor. Green dots, ratio of colocation observation pairs involving this sensor. Red crosses, ratio of colocations observations between this sensor and a reference sensor.

Subfigure (a): Variational Calibration Pair Model

Subfigure (b): Multi-hop

Figure 5:

Upper plots, enlargements of three weeks (highlighted in grey in the middle plot), showing correlation between the low cost sensor and the reference instrument. Middle plot, the full period of measurements. Discontinuity on 18th March, 2020, due to the covid-19 lock down (vertical black line). Lower plot, the ratio of the two sensors. A GP with an exponentiated quadratic (EQ) kernel has been fitted to the log ratio and the exponent of its posterior mean plotted (pre- and post- lock down modelled with independent GPs). Times are in UTC.

Figure 6:

(left) AirQo OPC measurements vs the reference instrument at the US embassy. Each row corresponds to 90 days of measurements. The left scatter plot shows the raw connection between the two sensors. A black line indicates equality while red lines are sampled from the linear fit's posterior probability distribution. Middle and right plots are the gradient and offset marginal posterior distributions respectively. The gradient distribution remains roughly constant, while the offset appears to 'drift' downwards over time. (right) One of the mobile low-cost sensors mounted on the front of a motorbike taxi ('boda boda'), photographed in Kampala by AirQo (www.airqo.net).

- Subfigure (a): Static AirQo OPC measurements vs the reference instrument data

- Subfigure (b): A mobile OPC

Figure 7:

The AirQo air pollution monitoring network (sensors that have been colocated). The four blue nodes are the mobile sensors. The orange node is the reference sensor at Makerere University, the circled node is the reference sensor whose calibration we are testing. (a) Locations of sensors, the latest locations of mobile sensors are plotted, with one not in the city limits. (b) The network of colocations between sensors (from July 2020 to February 2021). The line thicknesses indicates the number of colocations in that time. The green lines indicate two possible paths from the reference to the test sensor.

- Subfigure (a): Kampala with sensor locations
- Subfigure (b): Colocation frequency graph

Figure 8:

The distribution of errors in pollution predictions at the test sensor. As we give the three models the measured pollution, in principle the error could be zero if the model correctly estimates the calibration. The left graph is the distribution of predictions if we just use the measurements as the predictions. The added, synthetic, drift causes this to over-estimate. The central graph is using the multi-hop approach, this appears to underestimate. The right graph is for the calibration pair model. Runs of the variational inference optimisation end with slightly varying results here. This is a typical example (with a MAE about the average).

Figure 9:

The predictions of the scaling parameter at the test sensor. Red line, our synthetically generated scaling drift we are trying to predict. Black line, the prediction. Green dots, ratios between the test sensor and other sensor(s) that visit. The variational prediction posterior mean varied quite a lot per run of the algorithm, probably due to high uncertainty, leaving the choice of mean difficult to determine.

- Subfigure (a): Multi-hop scaling predictions
- Subfigure (b): Variational pair model scaling predictions

Table Captions

Table 1:

NMSE using raw data, the multi-hop calibration and the variational calibration pair model, for synthetic data with four levels of added noise. Bracketed values are the standard error of the mean based on predictions generated from 10 synthetic datasets (for each, the multi-hop was optimised on the test data, while the calibration pair model was optimised once on training data with noise scale = $10\mu\text{g m}^{-3}$). NMSE, Normalised Mean Squared Error; NLPD, Negative Log Predictive Density.

Table 2:

Kampala experimental results, with synthetic drift added. Using no calibration, the multi-hop calibration and the pair model (bracketed values are standard error from 20 runs. Not required for multi-hop as it's deterministic). MAE, Mean Absolute Error; NMSE, Normalised Mean Squared Error.