



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/220989/>

Version: Preprint

Preprint:

Smith, M.T., Ssematimba, J., Álvarez, M.A. et al. (Submitted: 2019) Machine Learning for a low-cost air pollution network. [Preprint - arXiv] (Submitted)

<https://doi.org/10.48550/arXiv.1911.12868>

© 2019 The Author(s). For reuse permissions, please contact the Author(s).

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.

Machine Learning for a Low-cost Air Pollution Network

Michael T. Smith

Department of Computer Science
University of Sheffield, UK
m.t.smith@sheffield.ac.uk

Joel Ssematimba

Department of Computer Science
Makerere University, Uganda
semajoel@gmail.com

Mauricio A. Álvarez

Department of Computer Science
University of Sheffield, UK
mauricio.alvarez@sheffield.ac.uk

Engineer Bainomugisha

Department of Computer Science
Makerere University, Uganda
baino@cis.mak.ac.ug

Abstract

Data collection in economically constrained countries often necessitates using approximate and biased measurements due to the low-cost of the sensors used. This leads to potentially invalid predictions and poor policies or decision making. This is especially an issue if methods from resource-rich regions are applied without handling these additional constraints. In this paper we show, through the use of an air pollution network example, how using probabilistic machine learning can mitigate some of the technical constraints. Specifically we experiment with modelling the calibration for individual sensors as either distributions or Gaussian processes over time, and discuss the wider issues around the decision process.

1 Introduction

We consider the example of a deployment of an air pollution monitoring network in Kampala, an East African city. Air pollution contributes to over three million deaths globally each year [Lelieveld et al., 2015]. Kampala has one of the highest concentrations of fine particulate matter (PM 2.5) of any African city Mead [2017]. Unfortunately, there is no programme for monitoring air pollution in the city due to the high cost of the equipment required. Hence we know little about its distribution or extent. Lower cost devices do exist, but these do not, on their own, provide the accuracy required for decision makers. In our case study, the Kampala network of sensors consists largely of low cost optical particle counters (OPCs) that give estimates of the PM2.5 particulate concentration. These are known to experience bias depending on humidity [Badura et al., 2018] and degrade relatively quickly due to dust and clogging. This network of sensors will soon be supplemented with three reference instruments (certified by MCERTS or equivalent). It is useful to briefly consider the additional issues in the Kampala network, compared to (for example) the LAQN in London. (1) The low-cost OPCs increasingly overestimate PM2.5 in humid conditions. (2) There is considerably more dust (coarse particulates) in the environment in Kampala, leading to sensor degradation. (3) In Kampala pollution exists from additional sources (e.g. road-surfaces, cooking, rubbish-burning, diesel generators, etc). (4) The PM2.5 estimate provided by the OPC is based on assumptions around particle size distributions which are likely to be inaccurate in Kampala. Regular calibration is clearly necessary, but the sensors typically can't be regularly moved. Thus we will be performing in-situ calibration using a set of mobile sensors installed on motor-bike taxis. This is a similar concept to that described in Kizel et al. [2018], in which sensors are calibrated in a chain. That model becomes somewhat intractable as the network becomes more complex, and fails to account for the time since calibration. Closely working with the Kampala Capital City Authority (KCCA) we have identified




Figure 1: *Left*, effect of clogging on measured particulate pollution. No co-located reference data was available, but ambient pollution was known to remain mainly in the 10-40 $\mu\text{g}/\text{m}^3$ range over the whole period. Note how measured value reduces due to dust entering the sensor. Maintenance was conducted in May, leading to a recovery of measurements. *Right* photo of dust clogging the fan.




Figure 2: *Left* Construction of the covariance matrix from three simulated sensors. The coregionalisation structure this produces is in the left most matrix. The second sensor is near the third initially and is then near the first sensor. This leads to the covariance in the central matrix. *Right* The M sensor calibration weights at all observed time points w_i are modelled with a sparse set of pseudo-observations z_s at time points t_s . The posterior mean predictions $w_i(t)$ of this sparse GP are used to scale the latent GP f 's posterior N predictions to produce the observations y . The two GPs hyperparameters are fixed in this model θ_w and θ_y .

a series of specific requirements for the model output. To summarise: the model should allow a prediction to be made at any location and should attempt to quantify its uncertainty. We chose to use Gaussian process regression (GPR) as this allows us to specify strong priors around the expected spatio-temporal structure of pollution and also provides the necessary uncertainty quantification. Our approach is to assume the low-cost OPC sensors merely measure a scaling to the true pollution value f by a weight w_i specific for that time and sensor i . So a given measurement y is given by, $y(t) = w_i(t)f(t) + \epsilon$. This calibration is both uncertain and is known to drift over time. For example Figure 1 illustrates the effect of clogging which causes gradual degradation of the sensor. We therefore develop a model to handle the uncertainty in the measurements.

1.1 Model

As this model and system is still being developed we use a numerical approach (using Hamiltonian Monte Carlo, HMC) to allow the model to be adjusted quickly. This also leads to greater sustainability (due to reduced expertise required). However the scalability of the current system is an issue so some of the posterior distribution estimation will need to be via a variational approximation in a deployed solution.

First consider several co-located sensors (one reference and several low-cost). A coregionalised Gaussian process regression model would be suitable for modelling this case [Álvarez and Lawrence, 2011], using a rank-1 coregionalisation matrix $\mathbf{a}\mathbf{a}^\top$, with the reference instrument's weight $a_1 = 1$ fixed. Standard maximum likelihood (ML) hyperparameter optimisation leads to weights that reflect the calibration adjustment required to correct for the biases in the low-cost sensors. We extend this by considering the case in which the sensors can be in different locations and (some) can move. Modelled with a kernel element-wise product, in which the distance (in space and time) is modelled with a standard EQ kernel, multiplied (element-wise) with the coregionalisation covariance matrix (see figure 2). ML optimisation is then able to find calibration values which reflect the sensor biases,




Figure 3: Left, Simulation data with two instruments. Reference instrument, blue crosses; bias OPC instrument, black crosses ($3\times$ correct value). For simplicity the correct value at the two instrument locations remains the same (unknown to the model), indicated by black line. Red circles are samples from the MCMC. Thick blue line shows their median and thin blue lines one standard error. Right, spatiotemporal location of both sensors (note the OPC moves away after time zero).

even if those sensors are not co-located with the reference instrument - so long as mobile instruments have visited the pairs of instruments. This approach fails however to properly quantify the uncertainty in the calibration, for example due to the changes to the pollution distribution (e.g. road-dust in the dry season). To mitigate this we model the weight hyperparameters a with a series of independent GP priors. For computational efficiency and to aid mixing, we introduced sparsity in these weights, so that, rather than each observation requiring a weight that also is a random variable from a GP, we use a series of latent virtual time points, t_s , and observations z , and use these to produce the posterior mean vectors $\mathbf{w}^{(j)}$ which is then used to scale the predictions from the GP at the observation locations. Figure 2 illustrates this slightly more complex model.

2 Results

The results below, unlike the example above, are based on simulated data, as reference instruments to validate the method are yet to be deployed. Implemented with Tensorflow Probability.

Two sensor demo (GP prior) Figure 3 shows with the use of a simulation the effect of the GP prior on the scaling weight by considering two simulated sensors (a reference instrument and a low-cost OPC). From time -1 to time zero they are collocated. After time zero the low-cost instrument moves away from the reference. The lengthscale of the scaling GP (exponentiated quadratic kernel) prior has been chosen to be only 4 time units, to demonstrate how the uncertainty in the calibration grows once the reference instrument stops providing support.

Sensor network (Gaussian prior) We next simulate the more complex situation in which four distant low-cost units need calibrating by using a pair of low-cost mobile sensors to ‘transport’ the calibration signal from the reference instrument. Due to the additional complexity we initially use a simple Gaussian prior on the weights (so the calibration is assumed to be fixed over time). Figure 4 illustrates the locations of the simulated sensors and the route that the mobile sensors make between them. Figure 4 shows the results. Most locations are fairly well estimated, with the model selecting the correct calibration. Note that mobile sensor (2) is more precisely characterised than (1). This might be because it spent longer at the reference instrument. The upshot though is that weights for sensors (5) and (6), which are visited by sensor (1), are not very accurately estimated. It could be that the prior’s variance needs increasing (currently 25) to reduce its effect, as the data is not very informative.

3 Discussion

The described probabilistic modelling approach to this issue appears to be relatively robust (the model worked without any careful parameter selection) and straight-forward. The main issue is when the

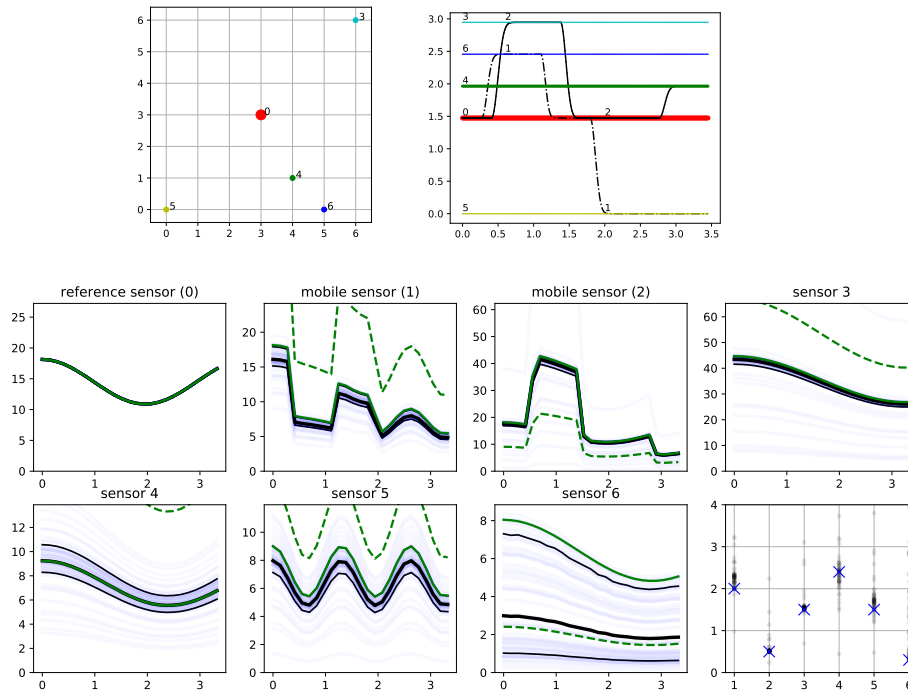


Figure 4: *Upper left*, a map of the four low-cost instruments (labelled 3-6, and the reference instrument 0. *Upper right*, shows how the two mobile sensors (black lines) move between the five static sensors. *Lower plot* Simulation of seven sensors. Measured value indicated with dashed green line. True pollution value indicated with solid green line. Samples from the MCMC based on the measured values indicated with faint blue lines (median, black thick line; standard-error, thin black lines). Figure at lower-right shows MCMC samples of these scaled weights.

calibration is ‘chained’ (not shown) which leads to very large uncertainty towards the end of the chain. This may just reflect the true limitations of a long calibration chain. The method described assumes that the pollution characteristics will not vary: i.e. the bias of a sensor is independent of location. This assumption can be tested by occasional visits with reference instrumentation across the city. Our next steps are to deploy reference instruments and test this method with the motor-bike taxis and low-cost OPCs already in-place. Some issues still need to be resolved: (1) Data collection biases. Specifically the focus on ambient (outdoor) air pollution. Gender roles in the society mean this is a gendered issue. We are actively working with partners to incorporate an indoor element to the monitoring. (2) Opportunity cost of implementation. The money and time spent developing and deploying the network may have been better spent on other development issues. However, training and mentoring are central to the project, with the intention that the research group will reach an international research standard. These indirect benefits may even exceed the direct results, for research projects, where supporting tertiary education is a key development outcome. (3) Sustainability. How long will the system last? Long-term cost? Who can maintain it? (4) The mobile sensors are mounted on motorbike taxis. The routes they take could conceivably contain private data. We are developing differential privacy methods to obscure this in the predictions. (5) (Ab)use of these results. Who will use the data? It is conceivable that it might be used as a reason or excuse to constrain an activity, such as a solid-fuel cooking, which a vulnerable group depends on. The alternative is to with-hold or cancel the monitoring (which also may be unethical).

Conclusion Many similar deployments of low-cost sensors exist as part of the move to ‘smart-cities’. But the poor quality of the data collected limits its use for policy making. In this paper we suggest a method to quantify these uncertainties, allowing predictions to be made to aid policy making and monitor interventions.

Acknowledgments Project funded by Google, USAID: Development Impact Lab and the UK EPSRC.

References

- Mauricio A Álvarez and Neil D Lawrence. Computationally efficient convolved multiple output gaussian processes. *Journal of Machine Learning Research*, 12(May):1459–1500, 2011.
- Marek Badura, Piotr Batog, Anetta Drzeniecka-Osiadacz, and Piotr Modzel. Evaluation of low-cost sensors for ambient pm_{2.5} monitoring. *Journal of Sensors*, 2018, 2018.
- Fadi Kizel, Yael Etzion, Rakefet Shafran-Nathan, Ilan Levy, Barak Fishbain, Alena Bartonova, and David M Broday. Node-to-node field calibration of wireless distributed air pollution sensor network. *Environmental Pollution*, 233:900–909, 2018.
- Jos Lelieveld et al. The contribution of outdoor air pollution sources to premature mortality on a global scale. *Nature*, 525(7569):367, 2015.
- Nick Van Mead. Pant by numbers: the cities with the most dangerous air – listed. *The Guardian*, Feb 2017.