
Gaussian Processes for Monitoring Air-Quality in Kampala

Clara Stoddart

Dept. of Computing
Imperial College London
London, UK

Lauren Shrack

Dept. of Computer Science
MIT
MA, USA

Richard Sserunjogi

Dept. of Computer Science
Makerere University / AirQo
Kampala, Uganda

Usman Abdul-Ganiy

Dept. of Computer Science
Makerere University / AirQo
Kampala, Uganda

Engineer Bainomugisha

Dept. of Computer Science
Makerere University / AirQo
Kampala, Uganda

Deo Okure

Dept. of Computer Science
Makerere University / AirQo
Kampala, Uganda

Ruth Misener

Dept. of Computing
Imperial College London
London, UK

Jose Pablo Folch *

Dept. of Mathematics
Imperial College London
London, UK

Ruby Sedgwick *

Depts. of Materials,
Bioengineering, and Computing
Imperial College London
London, UK

Abstract

Monitoring air pollution is of vital importance to the overall health of the population. Unfortunately, devices that can measure air quality can be expensive, and many cities in low and middle-income countries have to rely on a sparse allocation of them. In this paper, we investigate the use of Gaussian Processes for both *nowcasting* the current air-pollution in places where there are no sensors and *forecasting* the air-pollution in the future at the sensor locations. In particular, we focus on the city of Kampala in Uganda, using data from AirQo’s network of sensors. We demonstrate the advantage of removing outliers, compare different kernel functions and additional inputs. We also compare two sparse approximations to allow for the large amounts of temporal data in the dataset.

1 Introduction

The World Health Organisation recognises air quality as the single biggest environmental threat to human health, estimated to cause 7 million deaths annually [1], with the negative affects of air pollution predicted to get more severe with climate change [2–4]. Notably, particulate matter with a diameter of less than or equal to $2.5\mu\text{m}$, known as $PM_{2.5}$, has been associated with increased all-cause, cardiovascular, respiratory and lung cancer mortality [5, 6]. $PM_{2.5}$ has also been found to have short-term global heating effects [7–9].

In light of its dangerous impact, monitoring $PM_{2.5}$ levels in cities is incredibly important, but air pollution sensors are often prohibitively expensive. AirQo [<https://www.airqo.net/about>] are a research group based in Uganda, aiming to empower communities across Africa with information about air quality, in order to increase awareness and encourage citizens to take action against pollution. To address the current lack of air quality information in sub-Saharan Africa, AirQo has developed and deployed a sparse network of low-cost air quality monitors in multiple cities [10].

*equal contribution

We aim to accurately predict $PM_{2.5}$ levels in Kampala, using data from AirQo’s network of air quality sensors. Due to the sparse placement of the sensors, we require the predictive model to be data efficient and give a measure of uncertainty, which is useful for decision making. Gaussian processes (GPs) fulfil both these criteria, so we investigate their use for air pollution prediction, comparing the performance of a range of kernels, input factors and sparse approximations to find the best GP model for both nowcasting and forecasting. Data and code is available through the link in Appendix B.

2 Background and Related Work

2.1 Gaussian Processes

A Gaussian Process [11] is a stochastic process such that for any number of n inputs, $x_1, \dots, x_n \in \mathcal{X}$ the function $(f(x_1), \dots, f(x_n))$ follows a multivariate Gaussian distribution. The distribution is defined by a mean function, $\mu_0 : \mathcal{X} \rightarrow \mathbf{R}$ and a positive definite covariance function or *kernel* $\kappa_0 : \mathcal{X} \times \mathcal{X} \rightarrow \mathbf{R}$. When trying to model an unknown function, f , it is common to assume we have data of the form $y = f(x) + \sigma\epsilon$ where $\sigma > 0$, and ϵ is noise modelled using a standard Gaussian distribution. GP regression assumes the noise model, and puts a GP *prior* on the unknown function f . The posterior of f given a dataset $D = \{(x_i, y_i)\}_{i=1}^N$ is also a GP whose mean and covariance can be calculated analytically. See Appendix C for further details.

2.2 Kernel functions

The kernel function specifies the covariance of the Gaussian distribution, and therefore it defines the class of functions that it can learn. We focus on the squared exponential and periodic kernels:

$$\kappa_{SE}(x, x') = \exp\left(-\frac{\|x - x'\|^2}{2\ell^2}\right); \kappa_{Period}(x_i, x_j) = \exp\left(-\frac{1}{2} \sum_{k=1}^d \left(\frac{\sin(\frac{\pi}{p}(x_i^k - x_j^k))}{p}\right)^2\right) \quad (1)$$

where ℓ represents the length-scale parameter, which determines how smooth the function is, and p is the period parameter. The squared exponential assigns similarity to inputs that are close together, while the periodic assigns similarity between inputs that are in similar parts of the period. We can build a kernel from the product or sum of kernel functions, allowing for flexible combinations of different kernels when looking to combine different properties. For example, using a squared exponential kernel spatially and using a periodic kernel temporally.

2.3 Scalability of Gaussian Processes

Computing a GP posterior requires inverting a matrix, which typically incurs cubic cost in the number of training data points, i.e. $O(N^3)$. The current state-of-the-art solution to scaling GPs is the use of sparse approximations [12]. For these approximations, M pseudo-inputs or *inducing points* are selected such that the posterior of the GP is approximately equal to the posterior if it were to be trained in the full dataset. This results in a computational cost of order $O(M^2N)$, where $M \ll N$. Spatio-temporal variational Gaussian processes [13] provide similar results to standard variational GPs but manage to reduce the scaling in the temporal dimension from cubic to linear.

2.4 Related Work

GPs have shown promising results in predicting air quality [14–16]. Reggente et al. found GP regression outperformed Bayesian linear regression when forecasting particulate matter [14]. Jang et al. built on this by using a stochastic variational Gaussian process model (SVGP) [17] to predict particulate concentration using data from 131 monitoring sites with hourly measurements. They found that GPs outperformed linear regression and vector auto-regressive moving-average processes [15]. However both these studies used other pollutants as inputs to their models, which aren’t measured by AirQo’s low-cost sensors so are unavailable to us. They also only focus on forecasting, where we are also interested in nowcasting.

Cheng et al. [16] develop a nowcasting GP using data from a sensor network of 51 $PM_{2.5}$ monitoring devices deployed over an area of $30km$ by $30km$. Similar to our nowcasting aim, they sought to interpolate air-quality at areas between monitoring sites, restricting their analysis to the squared

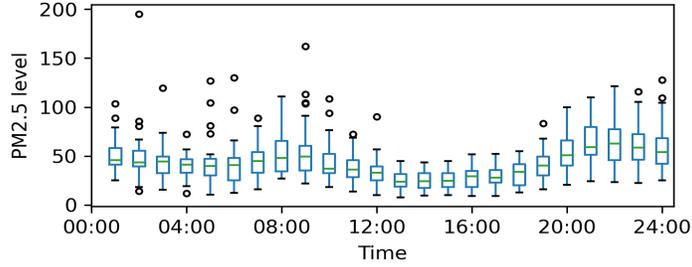


Figure 1: A box plot of $PM_{2.5}$ levels over November 2021 for the site Makerere University Weather Station. We can see a high quantity of outliers in the data; as well as a clear periodic pattern.

exponential kernel. Their results show that GPs outperform linear interpolation and cubic spline models.

3 Predicting Air Pollution in Kampala

3.1 Dataset

We use a dataset, provided by AirQo, of 66 air-quality sensors deployed around Kampala, using measurements from November 2021 [10]. The sensors provide hourly readings of $PM_{2.5}$ levels.

The dataset contains a lot of extreme values, as seen in Figure 1, with many readings outside of the interquartile range, as well as many complete outliers. These outliers can be explained by sensor errors, or sudden spikes in $PM_{2.5}$ levels. These sudden spikes in pollution don't necessarily follow any spatial or temporal pattern and it is unclear if they can be predicted from the data. Indeed, spikes may occur if, for example, a car accident causes unusual levels of traffic.

There also appears to be a daily periodic pattern with higher $PM_{2.5}$ levels in the morning, at around 8am, and later in the day, at around 9pm. This pattern holds across most sites, as seen in Appendix D. This motivates us to look at using a periodic kernel to predict the pollution levels.

3.2 Building the Models

In our experiments we investigate the effect of:

- (I) Adding periodicity to the kernel. We do this by testing a model with a squared exponential kernel against a model with a kernel that is a sum of a squared exponential kernel applied to the spatial dimensions and a periodic kernel applied to the temporal dimensions.
- (II) Removing extreme values from the training set, under the assumption that they are caused by rare external factors that may bias the model during normal times.
- (III) Our base models only use the location and time as inputs to get predictions. We investigate the effect of including additional inputs to the model: wind speed, wind direction, wind gusts, humidity, temperature, and precipitation.
- (IV) Due to the cubic scalability of GPs, we are unable to use the whole dataset to fit standard models and use 1000 randomly selected points. We investigate the use of sparse approximations to allow the training of the model on the whole dataset.

All models are built using the GPflow package [18]. For the periodic kernels, we encode a daily period and a weekly period (we do not consider monthly fluctuations as the dataset is for one month).

3.3 Model Predictive Performance

For each model we investigate two different tasks. *Nowcasting*, which refers to predicting pollution levels at *the current time*, in an *unknown location*; we test using leave-one-out cross validation on the sensor sites. The second task is *forecasting*, where we predict the pollution levels at the *current*

Table 1: Comparison of different Gaussian Process models for **nowcasting**.

Periodic	Outliers Removed	Additional Inputs	Sparse	Min RMSE	Average RMSE	Max RMSE
✗	✗	✗	✗	17.12	27.65	86.36
✓	✗	✗	✗	11.03	25.49	87.26
✓	✓	✗	✗	15.49	25.06	86.01
✓	✓	✓	✗	12.98	24.94	86.50
✓	✓	✓	SVGP	11.25	23.63	85.86
✗	✓	✗	ST-SVGP	9.01	22.34	87.04

Table 2: Comparison of different Gaussian Process models for **forecasting**.

Periodic	Outliers Removed	Additional Inputs	Sparse	Min RMSE	Average RMSE	Max RMSE
✗	✗	✗	✗	10.52	22.99	141.77
✓	✗	✗	✗	9.65	17.75	44.27
✓	✓	✗	✗	5.31	15.66	27.18
✓	✓	✓	✗	5.89	15.17	30.15
✓	✓	✓	SVGP	2.86	14.57	27.06
✗	✓	✗	ST-SVGP	5.56	15.34	32.85

location, in the *future*; we test by holding out the last day’s data. See Appedix E for implementation details. We use root mean squared error (RMSE) on the test sets as a performance metric.

The effect of adding periodicity We note a very small improvement for the nowcasting task. The amount of improvement suggests that previous trends have little effect on current prediction and it is the latest measurements that are more important. For the forecasting task, however, the average RMSE notably drops across sites, with the maximum RMSE decreasing significantly.

Removing outliers Removing outliers has a small impact on nowcasting results, but stronger for forecasting. It reduces the minimum, average and the maximum RMSE. The maximum RMSE is reduced significantly, explained by the aperiodicity of the spikes creating bias in the (otherwise periodic) model. We do not remove outliers from the test set for fair comparison.

Additional variables Wind direction, wind speed, humidity, precipitation and temperature, are included as variables into the RBF kernel, treating them the same as spatial inputs. The additional variables don’t seem to lead to any clear improvement in model performance. All additional variables are measured at one weather station in Kampala; so we would hope that with more fine-grained meteorological information results might be more significant.

Sparse approximations are vital to allow training of GPs on full large spatio-temporal datasets. In both nowcasting and forecasting we see small improvements across all metrics.

ST-SVGP Spatio-temporal variational GPs [13] reduce the temporal complexity significantly. However, we are unable to include additional inputs, as the model relies on creating a dense spatio-temporal grid. We also lose the ability to use a periodic kernel, since it cannot easily be written in the state-space form required to use ST-SVGP. We observe improved performance for nowcasting, but worsened performance for forecasting; most likely because of the loss of periodicity.

4 Conclusion

In our evaluation of GP models for monitoring air quality in Kampala, we found the models to be better at forecasting than nowcasting. For forecasting, periodicity and outlier removal improved performance. For nowcasting, the maximum RMSE was similar amongst all nowcasting models, suggesting difficulty in predicting certain locations without more sensors and perhaps the need for a non-stationary kernel [19]. We hope this work can help guide pollution predictive models and sensor placement in the future [20].

References

- [1] World Health Organisation. WHO global air quality guidelines: particulate matter (PM_{2.5} and PM₁₀) ozone, nitrogen dioxide, sulfur dioxide and carbon monoxide, 2021.
- [2] Sourangsu Chowdhury, Sagnik Dey, and Kirk R. Smith. Ambient PM_{2.5} exposure and expected premature mortality to 2100 in India under climate change scenarios. *Nature Communications*, 9:318, January 2018.
- [3] Marcos Lorrán Paranhos Leão, Linjie Zhang, and Flavio Manoel Rodrigues da Silva Júnior. Effect of particulate matter (PM_{2.5} and PM₁₀) on health indicators: climate change scenarios in a Brazilian metropolis. *Environmental Geochemistry and Health*, 45(5):2229–2240, May 2023.
- [4] Song Liu, Jia Xing, Daniel M. Westervelt, Shuchang Liu, Dian Ding, Arlene M. Fiore, Patrick L. Kinney, Yuqiang Zhang, Mike Z. He, Hongliang Zhang, Shovan K. Sahu, Fenfen Zhang, Bin Zhao, and Shuxiao Wang. Role of emission controls in reducing the 2050 climate change penalty for PM_{2.5} in China. *Science of The Total Environment*, 765:144338, April 2021.
- [5] Jie Chen and Gerard Hoek. Long-term exposure to PM and all-cause and cause-specific mortality: A systematic review and meta-analysis. *Environment International*, 143:1, 2020.
- [6] Aaron J Cohen, Michael Brauer, Richard Burnett, H Ross Anderson, Joseph Frostad, Kara Estep, Kalpana Balakrishnan, Bert Brunekreef, Lalit Dandona, Rakhi Dandona, Valery Feigin, Greg Freedman, Bryan Hubbell, Amelia Jobling, Haidong Kan, Luke Knibbs, Yang Liu, Randall Martin, Lidia Morawska, C Arden Pope, Hwashin Shin, Kurt Straif, Gavin Shaddick, Matthew Thomas, Rita van Dingenen, Aaron van Donkelaar, Theo Vos, Christopher J L Murray, and Mohammad H Forouzanfar. Estimates and 25-year trends of the global burden of disease attributable to ambient air pollution: an analysis of data from the global burden of diseases study 2015. *The Lancet*, 389(10082):1907–1918, 2017.
- [7] Eleanor J. Highwood and Robert P. Kinnersley. When smoke gets in our eyes: The multiple impacts of atmospheric black carbon on climate, air quality and health. *Environment International*, 32(4):560–566, May 2006.
- [8] T. C. Bond, S. J. Doherty, D. W. Fahey, P. M. Forster, T. Berntsen, B. J. DeAngelo, M. G. Flanner, S. Ghan, B. Kärcher, D. Koch, S. Kinne, Y. Kondo, P. K. Quinn, M. C. Sarofim, M. G. Schultz, M. Schulz, C. Venkataraman, H. Zhang, S. Zhang, N. Bellouin, S. K. Guttikunda, P. K. Hopke, M. Z. Jacobson, J. W. Kaiser, Z. Klimont, U. Lohmann, J. P. Schwarz, D. Shindell, T. Storelvmo, S. G. Warren, and C. S. Zender. Bounding the role of black carbon in the climate system: A scientific assessment. *Journal of Geophysical Research: Atmospheres*, 118(11): 5380–5552, 2013.
- [9] Szopa, S., V. Naik, B. Adhikary, P. Artaxo, T. Berntsen, W.D. Collins, S. Fuzzi, L. Gallardo, A. Kiendler-Scharr, Z. Klimont, H. Liao, N. Unger, and P. Zanis. Short-lived Climate Forcers. In *In Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, pages 817–922. 2023.
- [10] Richard Sserunjogi, Joel Ssematimba, Deo Okure, Daniel Ogenrwot, Priscilla Adong, Lillian Muyama, Noah Nsimbe, Martin Bbaale, and Engineer Bainomugisha. Seeing the air in detail: Hyperlocal air quality dataset collected from spatially distributed AirQo network. *Data in brief*, 44:108512, October 2022.
- [11] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.
- [12] Edward Snelson and Zoubin Ghahramani. Sparse gaussian processes using pseudo-inputs. In *Advances in Neural Information Processing Systems*, volume 18, 2005.
- [13] Oliver Hamelijnck, William Wilkinson, Niki Loppi, Arno Solin, and Theodoros Damoulas. Spatio-temporal variational gaussian processes. *Advances in Neural Information Processing Systems*, 34:23621–23633, 2021.
- [14] Matteo Reggente, Jan Peters, Jan Theunis, Martine Van Poppel, Michael Rademaker, Prashant Kumar, and Bernard De Baets. Prediction of ultrafine particle number concentrations in urban environments by means of gaussian process regression based on measurements of oxides of nitrogen. *Environmental Modelling & Software*, 61:135–150, 2014.

- [15] JoonHo Jang, Seungjae Shin, Hyunjin Lee, and Il-Chul Moon. Forecasting the concentration of particulate matter in the seoul metropolitan area using a gaussian process model. *Sensors*, 20 (14):3845, 2020.
- [16] Yun Cheng, Xiucheng Li, Zhijun Li, Shouxu Jiang, and Xiaofan Jiang. Fine-grained air quality monitoring based on gaussian process regression. In *Neural Information Processing: 21st International Conference*, pages 126–134, 2014.
- [17] James Hensman, Nicolò Fusi, and Neil D Lawrence. Gaussian processes for big data. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*, pages 282–290, 2013.
- [18] Alexander G. De G. Matthews, Mark Van Der Wilk, Tom Nickson, Keisuke Fujii, Alexis Boukouvalas, Pablo León-Villagrà, Zoubin Ghahramani, and James Hensman. GPflow: A Gaussian Process Library Using Tensorflow. *JMLR*, 18(1):1299–1304, 2017.
- [19] Zeel B. Patel, Palak Purohit, Harsh M. Patel, Shivam Sahni, and Nipun Batra. Accurate and Scalable Gaussian Processes for Fine-Grained Air Quality Inference. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(11):12080–12088, June 2022.
- [20] Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9(2), 2008.
- [21] Michalis Titsias. Variational Learning of Inducing Variables in Sparse Gaussian Processes. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, volume 5, pages 567–574, 2009.
- [22] Priscilla Adong, Engineer Bainomugisha, Deo Okure, and Richard Sserunjogi. Applying machine learning for large scale field calibration of low-cost pm2. 5 and pm10 air pollution sensors. *Applied AI Letters*, 3(3):76, 2022.
- [23] Michael Thomas Smith, Magnus Ross, Joel Ssematimba, Pablo A Alvarado, Mauricio Alvarez, Engineer Bainomugisha, and Richard Wilkinson. Modelling calibration uncertainty in networks of environmental sensors. *arXiv preprint arXiv:2205.01988*, 2022.
- [24] Alexander G de G Matthews, Mark Van Der Wilk, Tom Nickson, Keisuke Fujii, Alexis Boukouvalas, Pablo León-Villagrà, Zoubin Ghahramani, and James Hensman. Gpflow: A gaussian process library using tensorflow. *J. Mach. Learn. Res.*, 18(40):1–6, 2017.
- [25] Historical Weather Data for Kampala 2021-11-01 to 2021-12-01. URL <https://www.visualcrossing.com/weather-data>.

A Acknowledgments

We would firstly like to acknowledge the AI for Social Good seminar at Dagstuhl (seminar no. 22091) where the collaboration began, in particular we would like to thank Raghu Rajan and Asma Atamna for valuable feedback in previous iterations of the project. Richard S, Usman AG, Engineer B and Deo O received support from Belgium through the Wehubit programme implemented by Enabel and from Google LLC. Jose Pablo F is funded by EPSRC through the Modern Statistics and Statistical Machine Learning (StatML) CDT (grant no. EP/S023151/1) and by BASF SE, Ludwigshafen am Rhein. Ruby S was supported by the UKRI AI for Healthcare CDT (grant no. EP/S023283/1).

B Data Availability

The full code and dataset will be available at <https://github.com/claramst/gps-kampala-airquality>.

C Further Background on Gaussian Processes

A Gaussian Process [11] is defined as a stochastic process whose finite dimensional distributions are jointly Gaussian. That is, for some input space \mathcal{X} , and any number of n inputs, $x_1, \dots, x_n \in \mathcal{X}$ then $(f(x_1), \dots, f(x_n))$ follows a multivariate Gaussian distribution. The distribution can be completely defined by a mean function, $\mu_0 : \mathcal{X} \rightarrow \mathbf{R}$ and a positive definite covariance function or *kernel* $\kappa_0 : \mathcal{X} \times \mathcal{X} \rightarrow \mathbf{R}$.

When trying to model an unknown function, f , it is common to assume we have data of the form:

$$y = f(x) + \sigma\epsilon \tag{2}$$

where $\sigma > 0$, and ϵ is noise modelled using a standard Gaussian distribution. Gaussian Processes regression assumes the model above, and puts a Gaussian Process *prior* on the unknown function f . It can then be shown that the posterior of f given a dataset $D = \{(x_i, y_i)\}_{i=1}^N$ is also a Gaussian Process, with mean and covariance functions given by:

$$\mu(\mathbf{f}_*) = K(X_*, X)[K(X, X) + \sigma^2 I]^{-1}(\mathbf{y} - \mu_0(X)) \tag{3}$$

$$\sigma(\mathbf{f}_*) = K(X_*, X_*) - K(X_*, X)[K(X, X) + \sigma^2 I]^{-1}K(X, X_*) \tag{4}$$

where X_* is a matrix of target locations, \mathbf{f}_* the corresponding target outputs, X is the training inputs, \mathbf{y} is the training observations, $[\mu_0(X)]_i = \mu_0(X_i)$ and $[K(X, X')]_{i,j} = \kappa_0(x_i, x'_j)$. The hyper-parameters of the kernel and the noise level in the likelihood can all be learnt from data by maximizing the log marginal likelihood of the model.

The matrix inversion in equations (3) and (4) lead to the cubic scaling in GPs. When using sparse approximations to alleviate the problem, it is common to treat the pseudo-inputs as variational parameters [21] and learn them along the kernel hyper-parameters.

D Further Data Description and Analysis

Low-cost air quality monitors are prone to errors as they may be affected by weather conditions, and their overall accuracy degrades over time. AirQo use machine learning models (random forests and lasso regression) to calibrate the data in order to mitigate against errors [22, 23]. We use the calibrated $PM_{2.5}$ levels after AirQo’s processing. Some sensors have a number of missing or null readings due to errors, meaning that not all sites have the same number of readings for a given time frame. The full dataset for November 2021 contained readings from 68 sites across Kampala, however, 2 of the sites had less than 100 readings across the whole month so they were removed.

Figure 2 displays the mean $PM_{2.5}$ levels for each of the 66 sensors across the city. We can observe the clear periodic pattern, with pollution peaks around 8am and 9pm.

E Implementation Details

For the standard GP, we use an RBF kernel. Due to the size of the dataset, we cannot train using the full dataset. Therefore we randomly select 1000 training points before fitting non-sparse GPs;

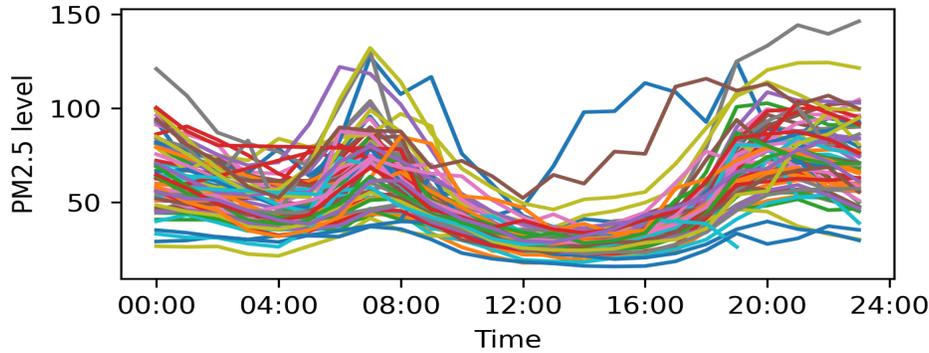


Figure 2: Mean $PM_{2.5}$ levels at each time for the 66 sites in Kampala over September 2021. We can see clear periodic patterns related to the daily commute of people in the city.

due to the randomness in the process we repeat each experiment 4 times and take an average – we do this for all non-sparse methods. We built all models using GPFlow [24], with the exception of spatio-temporal variational Gaussian processes (ST-SVGP) [13] where we used the implementation provided by the authors.

The effect of adding periodicity is tested by using a linear combination of an RBF kernel on the spatial dimensions and a periodic kernel on temporal ones. The periodic part is created using a product of a periodic kernel with a daily period and one with a weekly period (we do not consider monthly fluctuations as the dataset is for one month only).

Removing outliers We remove observations that are over 1.5 times the inter-quartile range away from the mean and comparing model performance. We do not remove outliers from the test set for fair comparison.

Additional variables are treated as spatial variables in an RBF kernel. There are publicly accessible sources of this information, and so we used data from the Visual Crossing Weather data service [25]. We do note that this data is sourced from only one weather data station in Kampala, and we take this information to be accurate across the whole city. In reality, wind speed may be slightly different on the other side of the city.

Sparse approximations A vital decision when using sparse approximations is the number of inducing points to use – the more inducing points the more accurate the approximation, but the computational cost increases as well. For the standard sparse variational approximation, **SVGP**, we decided on 100 inducing points. We selected this after seeing diminishing gains for nowcasting (we tested on 10, 50, 100, and 200 points). This allowed the sparse model to train on the full dataset in approximately two hours. For this model we included additional variables.

ST-SVGP We were able to comfortably use 66 inducing points, which effectively corresponds to using 66×720 inducing points in the SVGP model. Additional variables could be included as spatial ones, however, this would cause the spatio-temporal grid to become too large and make the method computationally infeasible.