
Optimising Placement of Pollution Sensors in Windy Environments

Sigrid Passano Hellan **Christopher G. Lucas** **Nigel H. Goddard**
School of Informatics
University of Edinburgh, UK
{s.p.hellan, c.lucas, nigel.goddard}@ed.ac.uk

Abstract

Air pollution is one of the most important causes of mortality in the world. Monitoring air pollution is useful to learn more about the link between health and pollutants, and to identify areas for intervention. Such monitoring is expensive, so it is important to place sensors as efficiently as possible. Bayesian optimisation has proven useful in choosing sensor locations, but typically relies on kernel functions that neglect the statistical structure of air pollution, such as the tendency of pollution to propagate in the prevailing wind direction. We describe two new wind-informed kernels and investigate their advantage for the task of actively learning locations of maximum pollution using Bayesian optimisation.

Air pollution is one of the most important causes of death globally. Particulate matter with diameter less than $2.5\ \mu\text{m}$, $\text{PM}_{2.5}$, is estimated to cause 3-4 million deaths yearly [1, 2]. Through monitoring one can learn more about the impact of air pollution on public health and identify areas requiring intervention. Traditional monitoring networks consist of a few high-quality reference sensors [3], but many low-cost pollution sensors are now available, e.g. [4], which can be used to build higher-resolution models to complement the traditional networks. This work aims to improve monitoring networks by improving the placement of low-cost, potentially noisy sensors to monitor average concentrations of $\text{PM}_{2.5}$ and identify areas of high average concentration, assuming that sensors can be placed incrementally. This problem – of finding the maxima of a black-box function given sparse and noisy observations – is well-matched to Bayesian optimisation (BO). Most machine learning research on air pollution tends not to consider the problem of efficient sensor placement, e.g. [5, 6, 7, 8, 9]. To our knowledge, previous work on environmental sensing using BO and related methods, e.g. [10, 11, 12], has not considered wind.

Wind is known to move pollutants downwind [13, p. 78-80], and work on gas mapping has shown the importance of incorporating this effect [14, 15]. For instance, [15] used a gas localisation task to show that a method incorporating wind velocity outperforms BO using wind-insensitive kernels. Because we are focusing on time-averaged pollution concentrations, detailed wind information and fluid simulations are not essential: The covariance structure of pollution is driven by the prevailing wind direction, which can be characterised with a single parameter. The kernels presented here allow such information to be incorporated into BO, and have the advantage that wind need not be measured separately. This is desirable from a cost perspective, as additional wind sensors are not needed.

The key question is whether these kernels improve our ability to find the location(s) of highest pollution in an area, with a minimum number of sensor placements, i.e., samples of time-averaged pollution concentrations. To answer this question using real data, we used satellite images of nitrogen dioxide (NO_2) concentrations (see Supplementary Material for details). NO_2 is a health-related pollutant in its own right [16], and we assume that the dispersion of $\text{PM}_{2.5}$ is comparable. While satellite data cannot replace surface-level sensor data, it provides a good testbed given the spatial regularity and density of the observations, and it allows us to set aside complex effects of urban topography.

1 Wind-informed kernels for Bayesian optimisation

In Bayesian optimisation [17] a probabilistic model, in our case a GP, is fitted to the set of observed data points. This model is then used to compute an acquisition function, and the next sample is taken at the point which maximises this function. The new sample is added to the set of data points and the process is repeated. In this paper, the upper confidence bound (UCB) [18] was used as the acquisition function, $\alpha_{\text{UCB}}(\mathbf{x}; \beta, \mathcal{D}, M) = \mu(\mathbf{x}; \mathcal{D}, M) + \beta\sigma(\mathbf{x}; \mathcal{D}, M)$, with $\beta = 1$. $\mu(\mathbf{x}; \mathcal{D}, M)$ and $\sigma(\mathbf{x}; \mathcal{D}, M)$ are the posterior mean and standard deviation of a point \mathbf{x} given data \mathcal{D} and model M . An entropy-based acquisition function [19] and other values of β were tested, without marked improvement. To compute $\mu(\cdot)$ and $\sigma(\cdot)$, covariances need to be computed between the data points using a kernel function. A thorough introduction to GPs can be found in [20].

The motivation for the wind-informed kernels is to exploit the directional structure of wind-driven pollution transport. Intuitively, one would expect to see the highest levels of pollution downwind from a source, and if one recorded a high pollution concentration in a location one would expect to find a source upwind from it. In this work we do not aim to locate sources, but to look at concentrations, and so do not differentiate between upwind and downwind movement.

The standard radial basis function (RBF) was modified to create the wind-informed kernels. The RBF kernel calculates the covariance between two points \mathbf{x}_i and \mathbf{x}_j as $k_R(\boldsymbol{\tau}; l, \sigma_R) = \sigma_R^2 \exp(-\boldsymbol{\tau}^T \boldsymbol{\tau} / l^2)$ where $\boldsymbol{\tau} = \mathbf{x}_i - \mathbf{x}_j$, l is the length scale and σ_R^2 the signal variance. The new kernels $k_S(\cdot)$, *Sum*, and $k_P(\cdot)$, *Product*, are given in Eqs. (1) and (2) and visualised in Fig. 1. The kernels are a combination of a standard RBF kernel and another RBF kernel on the distance orthogonal to the wind direction. This distance is given by $\tilde{\tau}(\boldsymbol{\tau}; \gamma) = \|\boldsymbol{\tau} - (\boldsymbol{\tau}^T \mathbf{b}) \mathbf{b}\|$ where $\mathbf{b} = [\cos(\gamma) \ \sin(\gamma)]^T$ and γ is the wind direction hyperparameter. Further, l_D is the directional length scale, σ_D^2 the directional signal variance, σ_P^2 the signal variance for $k_P(\cdot)$, $A = \begin{bmatrix} \sin(\gamma)^2 & -\sin(\gamma) \cos(\gamma) \\ -\sin(\gamma) \cos(\gamma) & \cos(\gamma)^2 \end{bmatrix}$ and $E = l_D^2 / l^2 \mathbb{I} + A$. Finally, $\theta_S = \{l, l_D, \sigma_R, \sigma_D, \gamma\}$ and $\theta_P = \{l, l_D, \sigma_P, \gamma\}$ denote sets of hyperparameters. The kernel hyperparameters, together with the noise hyperparameter σ_n^2 , are fitted in the GP tuning step.

$$k_S(\boldsymbol{\tau}; \theta_S) = k_R(\boldsymbol{\tau}; l, \sigma_R) + k_R(\tilde{\tau}(\boldsymbol{\tau}; \gamma); l_D, \sigma_D) = \sigma_R^2 \exp\left(\frac{-\boldsymbol{\tau}^T \boldsymbol{\tau}}{l^2}\right) + \sigma_D^2 \exp\left(\frac{-\boldsymbol{\tau}^T A \boldsymbol{\tau}}{l_D^2}\right) \quad (1)$$

$$k_P(\boldsymbol{\tau}; \theta_P) = k_R(\boldsymbol{\tau}; l, \sigma_P) k_R(\tilde{\tau}(\boldsymbol{\tau}; \gamma); l_D, 1) = \sigma_P^2 \exp\left(\frac{-\boldsymbol{\tau}^T E \boldsymbol{\tau}}{l_D^2}\right) \quad (2)$$




Figure 1: Visualisation of wind-informed kernels from Eqs. (1) and (2). The top row shows covariance as function of distance, and the bottom row shows example images sampled from the kernels. $\sigma_D^2 = 0.5$, $l = 105$, $l_D = 35$, $\gamma = \pi/4$, $\sigma_P^2 = 1$, $\sigma_n^2 = 0.01$. For Sum $\sigma_R^2 = 0.5$, for RBF $\sigma_R^2 = 1$.

These kernels are sums and products of RBF kernel functions that are valid for all inputs, and are thus valid kernel functions themselves [20, p. 83,95].

2 Data

The data set consists of 1083 images of 28x28 pixels, with each pixel giving the NO₂ concentration in mol/m² within a ca. 7x7km box on the ground reaching the upper troposphere. The images are from October and November 2018 and cover landmasses across the world. More information is given in the Supplementary Material. We removed images with more than 10 % missing values, and split the data into subsets. The split was done as NO₂ concentrations varied greatly between images, and locating peaks of high pollution is more important. Three subsets were created as follows: the images were ordered by the highest NO₂ concentration present, and the top, median and bottom 50 images put into sets called 'Strong', 'Median' and 'Weak', respectively. See Fig. 2. Three adjacent sets of 10 images each were created for tuning and normalisation. A final set of 100 images, dubbed 'Selection', was created by picking every ninth of the remaining ordered images. The first three subsets were created to inspect performance on different degrees of difficulty, and the last to test the overall performance.




Figure 2: Example NO₂ concentration images from the Strong, Median and Weak subsets. White pixels indicate missing data. Note the difference in scale between the subsets, and the pollution plumes from point sources in two of the examples. Such plumes were present in many images, particularly in the Strong subset, and are what the wind-informed kernels are meant to capture.

3 Results

We used log concentrations rather than raw concentrations as our data – and pollution concentrations in general [21] – are approximately log-normal, satisfying the assumption of Gaussian marginal and residual distributions in GP models. Images were normalised using means and standard deviations from the associated tuning set. The rationale for this was that the true mean and standard deviation for a deployment region would likely be unknown, but values from similar areas available. The tuning sets were also used to develop priors for the hyperparameters, except for γ for which a uniform prior was used. For each tuning image, gradient descent with 200 hyperparameter initialisations was used to maximise the log marginal likelihood, $\log(p(Y^i|M))$, where M denotes the model, including kernel and hyperparameter choices, and Y^i is the pixel values in image i . A log normal distribution was then fitted to these hyperparameter samples using the maximum likelihood estimator and this was used as a prior. For the Selection subset, the Median tuning set was used for normalisation and prior computation. The models obtained on the tuning sets were compared using the Bayesian information criterion (BIC) [22]. It was highest for the Sum kernel, followed by the Product kernel, c.f. the Supplementary Material. This shows that the wind-informed kernels are able to capture more of the structure of the pollution in the images. For the Weak subset little difference was seen.

The kernels were evaluated on the task of locating the pixel with the maximum pollution concentration. Results on two of the subsets, Strong and Selection, are presented. For comparison, our baselines were (1) Bayesian optimisation with an RBF kernel (RBF) and (2) a policy of choosing points uniformly at random (Random). The latter was repeated 100 times per image and the mean taken. The experiments were initiated with 10 randomly selected pixel values. Fig. 3 shows the performance as a function of iterations done, showing the Sum kernel performs better than the RBF kernel on the Strong subset. For the selection subset, the Random baseline beats the RBF baseline. This indicates that the problem is hard for BO, as would be expected from the example images in Fig. 2. The Product kernel outperforms the other kernels on the distance metric, while the ratio metric is very similar for all methods. Fig. 4 looks more closely at the performance on the Selection subset as a function of the maximum pollution concentration present in each image. It shows that the performance improves with higher levels of pollution and also that over most of the concentration values the Product kernel is better than the RBF.




Figure 3: Comparison of kernel performance on Strong and Selection subsets. \hat{x} is the estimated maximiser and x^* the true maximiser. \hat{y} and y^* are the true concentration values at \hat{x} and x^* , respectively. Both metrics are shown as means with standard deviation of the mean in dotted lines. The Sum kernel reaches slightly better values than the RBF kernel on the Strong subset, and the Product kernel outperforms the other kernels on the distance metric on the Selection subset.




Figure 4: Inspection of performance after 100 iterations over the range of concentrations in the Selection subset. Each line shows the running average of the 20 previous values. \hat{x} , x^* , \hat{y} and y^* as defined for Fig. 3. As can be seen the Product kernel outperforms RBF across most concentration.

4 Discussion and future work

By augmenting a standard RBF kernel with a directional element that captures the prevailing wind direction, we can better capture the covariance structure in NO_2 pollution, as indicated by the BIC scores. The results also show that this better estimates pollution maxima and the locations of those maxima for some of the images. The improvement was less clear than we expected from the BIC scores. We hypothesise that this is due to the hyperparameter tuning, which the results were sensitive to. For an example, see the Supplementary Material. Therefore, we plan to improve our priors by basing them on a larger tuning set and using hierarchical modelling. Additionally, the wind-informed kernels should be tested on more data sets, and particularly on data sets of urban air pollution. As wind is also important in moving air pollution over short distances one would expect improvement, but this might be distorted by other factors, like buildings. The approach used, of incorporating a data transformation into the covariance function, can be extended to other kernels, e.g. the set of Matérn kernels, or combined in other ways, e.g. a combination of Sum and Product kernels.

Acknowledgments and Disclosure of Funding

The authors would like to thank Dr Douglas Finch at the University of Edinburgh for extracting the data and helpful correspondence.

This work was supported in part by the EPSRC Centre for Doctoral Training in Data Science, funded by the UK Engineering and Physical Sciences Research Council (grant EP/L016427/1) and the University of Edinburgh.

References

- [1] 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, Martin Randall, 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, May 2017. ISSN 0140-6736. doi: 10.1016/S0140-6736(17)30505-6. URL <https://www.sciencedirect.com/science/article/pii/S0140673617305056>.
- [2] J. Lelieveld, J. S. Evans, M. Fnais, D. Giannadaki, and A. Pozzer. The contribution of outdoor air pollution sources to premature mortality on a global scale. *Nature*, 525(7569):367–371, September 2015. ISSN 1476-4687. doi: 10.1038/nature15371. URL <https://www.nature.com/articles/nature15371>. Number: 7569 Publisher: Nature Publishing Group.
- [3] Marco Carminati, Giorgio Ferrari, and Marco Sampietro. Emerging miniaturized technologies for airborne particulate matter pervasive monitoring. *Measurement*, 101:250–256, April 2017. ISSN 0263-2241. doi: 10.1016/j.measurement.2015.12.028. URL <http://www.sciencedirect.com/science/article/pii/S0263224115006880>.
- [4] Xiaoting Liu, Rohan Jayaratne, Phong Thai, Tara Kuhn, Isak Zing, Bryce Christensen, Riki Lamont, Matthew Dunbabin, Sicong Zhu, Jian Gao, David Wainwright, Donald Neale, Ruby Kan, John Kirkwood, and Lidia Morawska. Low-cost sensors as an alternative for long-term air quality monitoring. *Environmental Research*, 185:109438, June 2020. ISSN 0013-9351. doi: 10.1016/j.envres.2020.109438. URL <http://www.sciencedirect.com/science/article/pii/S0013935120303315>.
- [5] Michael T. Smith, Joel Ssematimba, Mauricio A. Alvarez, and Engineer Bainomugisha. Machine Learning for a Low-cost Air Pollution Network. *NeurIPS 2019 Workshop on Machine Learning for the Developing World*, December 2019. URL <http://arxiv.org/abs/1911.12868>. arXiv: 1911.12868.
- [6] Xiuming Liu, Teng Xi, and Edith Ngai. Data Modelling with Gaussian Process in Sensor Networks for Urban Environmental Monitoring. In *2016 IEEE 24th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS)*, pages 457–462, September 2016. doi: 10.1109/MASCOTS.2016.45. ISSN: 2375-0227.
- [7] Hongbin Liu, Chong Yang, Mingzhi Huang, Dongsheng Wang, and ChangKyoo Yoo. Modeling of subway indoor air quality using Gaussian process regression. *Journal of Hazardous Materials*, 359: 266–273, October 2018. ISSN 0304-3894. doi: 10.1016/j.jhazmat.2018.07.034. URL <http://www.sciencedirect.com/science/article/pii/S0304389418305648>.
- [8] Haolin Fei, Xiaofeng Wu, Chunbo Luo, and Yang Luo. Accurate Air Quality Prediction: A Physical-Temporal Collection Model. In *ICLR Workshop on AI for Earth Science*, 2020. URL https://www.researchgate.net/publication/341445468_ACCURATE_AIR_QUALITY_PREDICTION_A_PHYSICAL-_TEMPORAL_COLLECTION_MODEL_CONFERENCE_SUBMISSIONS.
- [9] Congcong Wen, Shufu Liu, Xiaojing Yao, Ling Peng, Xiang Li, Yuan Hu, and Tianhe Chi. A novel spatiotemporal convolutional long short-term neural network for air pollution prediction. *Science of The Total Environment*, 654:1091–1099, March 2019. ISSN 0048-9697. doi: 10.1016/j.scitotenv.2018.11.086. URL <http://www.sciencedirect.com/science/article/pii/S0048969718344413>.
- [10] R. Marchant and F. Ramos. Bayesian optimisation for Intelligent Environmental Monitoring. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2242–2249, October 2012. doi: 10.1109/IROS.2012.6385653.
- [11] Philippe Morere, Roman Marchant, and Fabio Ramos. Sequential Bayesian optimization as a POMDP for environment monitoring with UAVs. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6381–6388, May 2017. doi: 10.1109/ICRA.2017.7989754.
- [12] A. Singh, F. Ramos, H. D. Whyte, and W. J. Kaiser. Modeling and decision making in spatio-temporal processes for environmental surveillance. In *2010 IEEE International Conference on Robotics and Automation*, pages 5490–5497, May 2010. doi: 10.1109/ROBOT.2010.5509934.

- [13] Perry J Samson. Atmospheric Transport and Dispersion of Air Pollutants Associated with Vehicular Emissions. In Ann Y. Watson, Richard R. Bates, and Donald Kennedy, editors, *Air Pollution, the Automobile, and Public Health*. National Academies Press (US), 1988. URL <https://www.ncbi.nlm.nih.gov/books/NBK218142/>. Publication Title: Air Pollution, the Automobile, and Public Health.
- [14] Matteo Reggente and Achim J. Lilienthal. Using local wind information for gas distribution mapping in outdoor environments with a mobile robot. In *2009 IEEE SENSORS*, pages 1715–1720, October 2009. doi: 10.1109/ICSENS.2009.5398498. ISSN: 1930-0395.
- [15] Martin Asenov, Marius Rutkauskas, Derryck Reid, Kartic Subr, and Subramanian Ramamoorthy. Active Localization of Gas Leaks Using Fluid Simulation. *IEEE Robotics and Automation Letters*, 4(2):1776–1783, April 2019. ISSN 2377-3766. doi: 10.1109/LRA.2019.2895820. Conference Name: IEEE Robotics and Automation Letters.
- [16] Kai Luo, Runkui Li, Wenjing Li, Zongshuang Wang, Xinming Ma, Ruiming Zhang, Xin Fang, Zhenglai Wu, Yang Cao, and Qun Xu. Acute Effects of Nitrogen Dioxide on Cardiovascular Mortality in Beijing: An Exploration of Spatial Heterogeneity and the District-specific Predictors. *Scientific Reports*, 6(1): 38328, December 2016. ISSN 2045-2322. doi: 10.1038/srep38328. URL <https://www.nature.com/articles/srep38328>. Number: 1 Publisher: Nature Publishing Group.
- [17] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proceedings of the IEEE*, 104(1):148–175, January 2016. ISSN 1558-2256. doi: 10.1109/JPROC.2015.2494218. Conference Name: Proceedings of the IEEE.
- [18] Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias Seeger. Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design. *IEEE Transactions on Information Theory*, 58(5):3250–3265, May 2012. ISSN 0018-9448, 1557-9654. doi: 10.1109/TIT.2011.2182033. URL <http://arxiv.org/abs/0912.3995>. arXiv: 0912.3995.
- [19] José Miguel Hernández-Lobato, Matthew W Hoffman, and Zoubin Ghahramani. Predictive Entropy Search for Efficient Global Optimization of Black-box Functions. *Advances in Neural Information Processing Systems 27*, pages 918–926, 2014.
- [20] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [21] GJ Cats and AAM Holtslag. Prediction of air pollution frequency distribution—Part I. The lognormal model. *Atmospheric Environment (1967)*, 14:255–258, 1980. URL <https://www.sciencedirect.com/science/article/pii/0004698180902851>.
- [22] Gideon Schwarz. Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464, 1978.

A Supplementary Material

Supplementary material for *Optimising Placement of Pollution Sensors in Windy Environments* by Sigrid Passano Hellan, Christopher G. Lucas and Nigel H. Goddard presented at the AI for Earth Sciences Workshop at NeurIPS 2020.

A.1 Data

The data were collected by the TROPOspheric Measuring Instrument (TROPOMI) aboard the Sentinel 5P satellite,¹ and were extracted and shared with us by Dr Douglas Finch at the University of Edinburgh. In the original data source, the NO₂ data used is given in the column labelled 'level 2 nitrogen dioxide tropospheric'. Images covering mostly oceans were excluded.

A.2 BIC scores

The Bayesian information criterion (BIC) scores of the tuning sets referred to in Section 3 are given in Table 1. The values are given as differences because the obtained BIC scores varied a lot between problems, e.g. between 306 and -1014 for the Sum kernel on the Strong subset. As can be seen, the relative differences vary a lot less.

Table 1: Differences in Bayesian information criterion (BIC) scores obtained on tuning sets. Mean (standard deviation of mean). The wind-informed kernels were better able to explain the data for the Strong and Median subsets. There was little difference for the Weak subset, where the images were very noisy.

	Sum - Product	Sum - RBF	Product - RBF
Strong tune	13.43 (3.69)	19.92 (3.44)	6.49 (2.75)
Median tune	2.69 (1.95)	10.02 (4.30)	7.33 (3.73)
Weak tune	0.35 (0.58)	-1.12 (1.72)	-1.47 (1.34)

A.3 Implementation notes

A small amount of noise was added to sampled values in the BO loop for numerical stability. At each iteration in the BO loop 100 hyperparameter seeds were sampled from the priors and used as starting points for the GP tuning. The setting which maximised the posterior log marginal likelihood was then used. When estimating the hyperparameter priors, Bessel's correction was used for the variance estimate.

A.4 Hyperparameter prior sensitivity

Fig. 5 shows the results on the Strong subset when a slightly different prior was used for the hyperparameters, as referred to in Section 4 of the main paper.

¹Copernicus, <https://www.copernicus.eu/en>




Figure 5: Comparison of kernel performance on Strong subsets, when biased estimates of variances of hyperparameter priors are used, i.e. not using Bessel's correction. Compare to Fig. 3 in body of paper. With these settings, the Product kernel performs best. Comparing the figures, the Product kernel has improved but the RBF worsened. \hat{x} is the estimated maximiser and x^* the true maximiser. \hat{y} and y^* are the true concentration values at \hat{x} and x^* , respectively.