

Applications of Mathematics

Gabor Tajnafoi; Rossella Arcucci; Laetitia Mottet; Carollanne Vouriot; Miguel Molina-Solana; Christopher Pain; Yi-Ke Guo

Variational Gaussian process for optimal sensor placement

Applications of Mathematics, Vol. 66 (2021), No. 2, 287–317

Persistent URL: <http://dml.cz/dmlcz/148835>

Terms of use:

© Institute of Mathematics AS CR, 2021

Institute of Mathematics of the Czech Academy of Sciences provides access to digitized documents strictly for personal use. Each copy of any part of this document must contain these *Terms of use*.



This document has been digitized, optimized for electronic delivery and stamped with digital signature within the project *DML-CZ: The Czech Digital Mathematics Library* <http://dml.cz>

VARIATIONAL GAUSSIAN PROCESS
FOR OPTIMAL SENSOR PLACEMENTGABOR TAJNAFOI, ROSSELLA ARCUCCI, LAETITIA MOTTET,
CAROLANNE VOURIOT, London,
MIGUEL MOLINA-SOLANA, Granada,
CHRISTOPHER PAIN, YI-KE GUO, London

Received November 16, 2019. Published online February 12, 2021.

Abstract. Sensor placement is an optimisation problem that has recently gained great relevance. In order to achieve accurate online updates of a predictive model, sensors are used to provide observations. When sensor location is optimally selected, the predictive model can greatly reduce its internal errors. A greedy-selection algorithm is used for locating these optimal spatial locations from a numerical embedded space. A novel architecture for solving this big data problem is proposed, relying on a variational Gaussian process. The generalisation of the model is further improved via the preconditioning of its inputs: Masked Autoregressive Flows are implemented to learn nonlinear, invertible transformations of the conditionally modelled spatial features. Finally, a global optimisation strategy extending the Mutual Information-based optimisation and fine-tuning of the selected optimal location is proposed. The methodology is parallelised to speed up the computational time, making these tools very fast despite the high complexity associated with both spatial modelling and placement tasks. The model is applied to a real three-dimensional test case considering a room within the Clarence Centre building located in Elephant and Castle, London, UK.

Keywords: sensor placement; variational Gaussian process; mutual information

MSC 2020: 65Z05, 68T99

This work is supported by the EPSRC Grand Challenge grant Managing Air for Green Inner Cities (MAGIC) EP/N010221/1, the EP/T003189/1 Health assessment across biological length scales for personal pollution exposure and its mitigation (INHALE), the EP/T000414/1 PREdictive Modelling with QuantIfication of UncERTainty for MultiphasE Systems (PREMIERE) and the Leonardo Centre for Sustainable Business at Imperial College London.

1. INTRODUCTION

Indoor Air Quality (IAQ) impacts health, comfort and quality of life [24], and three basic strategies have been proposed to improve it: control of pollution sources, use of natural/mechanical ventilation, and cleaning of air. In the building context, the management and development of smart monitoring tools can support an adequate IAQ within them (e.g. by automatically opening windows or starting a mechanical air cleaning system). Sensors coupled with indoor pollutant forecasting models can tackle bad IAQ by implementing one of the previously-cited strategies before the indoor pollutant concentration reaches dangerous and adverse levels.

In order to achieve accurate, online updates of the predictive model, sensors can be used to provide observations. Spatio-temporal models such as Data Assimilation (DA) provide online learning and forecasting of sensor observations by means of updating the model's internal view through the incorporation of collected data [5], [13]. In this context, sensor positioning has gained relevance [20], [26] as it is crucial to ensure a good quality and usefulness of monitored data. Optimal sensor positioning tools pin-point the discrete spatial locations that possess most conditional information of all other spatial points, thus improving the predictive accuracy of prediction models [3]. Hence, sensor placement can be seen as an optimisation problem [20], [26].

Early attempts on sensor placement used geometric approaches, supported by the assumption that sensors measure spatial features with a fixed sensing radius [17]. This geometric approach does not take into account the nonlinear dynamic behaviour of air motion, so, in order to tackle this problem, parametric models [1], nonparametric Gaussian Process (GP) [11] and ensemble Kalman-filters [28] approaches were subsequently proposed.

The main work on this field implements sensor placement using a GP in a 2D space [26]. The time complexity of the placement algorithm is $O(N^4)$, where N denotes the side of the computational domain. The GP is trained on data collected from fixed sensors located in a room, and therefore N is relatively small. The placement algorithm only selects the best sensors in the set already provided. In [15], the problem of finding the optimal sensor placement for pollutant dispersion within an urban environment is addressed using a weak constraint GP model [4]. Also, the model uses temporal sequences of data from fluid dynamic simulations, therefore facing a big data problem. The resulting big data problem is addressed by introducing a domain decomposition approach.

In this work, for the first time, a sensor placement model is developed for a 3D domain representing an indoor real case scenario. In our case, N is on a scale such that the use of a GP is unfeasible. To address this problem we developed a combination of deep learning, probabilistic frameworks and variational methods,

reducing the complexity associated with training, inference and optimisation and thus enabling us to achieve optimal placement results in a real case scenario. The complexity of our model is, in fact, $O(klM^4)$, where k is the number of sensors, l is the number of iterations needed to optimise the position and M is such that $M \ll N$.

The rest of this paper is organized as follows: the next section introduces the background to this work and our main contributions. Section 3 presents the mathematical formulation of our proposed model, variational Gaussian process optimal sensor (VGPosp). Section 4 describes the direct application of VGPosp to a real indoor environment. The manuscript ends with some conclusions and further work.

2. BACKGROUND AND MAIN CONTRIBUTIONS

Nonparametric models consist in learning a Gaussian Process (GP) associated with the phenomenology considered (e.g. pollution levels in indoor environments). In general, GPs are highly appropriate to study environmental problems as they allow for learning complex, high-dimensional correlations with uncertainty quantification. Indeed, nonparametric expressiveness is an advantage over parametric models that are more prone to the curse of dimensionality [30]. A GP, sometimes referred as the Bayesian interpretation of neural networks, is fully determined by only two parameters, namely the mean-function and covariance-function, regardless of their dimensionality [37]. Three main GP methodologies can be identified: the Traditional GP [20], [26], the Sparse GP [35] and the Variational GP [42], [41].

Traditional GP is a stochastic process based on prior distribution over functions, and it has been successfully applied for indoor optimal sensor positioning [20], [26]. However, Traditional GP suffers from the high complexities associated with spatial modelling $O((mN)^3)$, where N denotes the size of input sensor potential locations and m the number of physical variables, which explains why the work presented in [20], [26] was only carried out in two dimensions.

Sparse Variational Process (SGP) tackles the inconvenient $O((mN)^3)$ computational complexity associated with Traditional GP [21]. This method constructs an approximation based on a small subset of size \hat{N} , namely inducing points. This optimisation results in a reduced complexity $O((mN)\hat{N}^2)$, enabling the scalability of training data-points from the previous limit of a few thousand to the range of millions [22]. In general, sparsity can be achieved by working on a low-rank representation of the full kernel [29]. The key idea is to approximate the prior or modify the likelihood function, thus creating a model selection problem solving the optimisation for the approximation of the truth [35]. However, the main criticism of SGPs

is that they learn unknown hyperparameters, potentially leading them to underestimate variance and thus over-fitting [42], [41].

Alternatively, variational Gaussian process (VGP), a variational method for SGP, was developed [42], [41] to deal with the approximation of model components that are hard to compute. Inducing points are variational parameters selected by minimising the Kullback-Leibler (KL) divergence [42], [41]. The kernel hyperparameters and inducing points are jointly optimised by maximising a lower bound (Evidence Lower BOund (ELBO)) of the variational distribution over the function’s latent values [42], [41], [25]. The key innovation is that the likelihood and the GP prior are not modified, separating the model and the inference. The variational posterior iteratively approaches the true posterior.

The datasets used to train a GP usually comprise data coming from monitoring sensors during extensive field experiment periods. This training dataset usually suffers from being non-Gaussian distributed, rendering it unusable for GP learning [33]. In this regard, several methodologies to precondition, normalise and render the training dataset Gaussian can be mentioned: Variational Autoencoder [14], Autoregressive Flows [45], Normalising Flows [38], Masked Autoencoder for Distribution Estimation (MADE) [16], and Masked Autoregressive Flows (MAF) [33]. The MAF approach is a stack of MADE networks [16], [33] and has proved its competitiveness over the other methodologies in terms of accuracy [33].

Even if VGP can deal with non-Gaussian distributed data, the preconditioning phase of learning nonlinear, invertible transformations between the conditional input distributions and the output Gaussian family of distributions enables greater generalisation of the learned spatial model.

At this stage, the actual sensor placement problem can then be addressed using the trained GP. Indeed, the placement algorithm is solved in an embedded space that is predicted by a GP [42], [41]. In other words, the GP serves as a numerical setting for the optimisation problem: conditional predictions are used to generate the covariance matrix taken as input by the placement algorithm. The complexity associated with the placement task is $O(N^4)$, where N denotes the size of input sensor potential locations.

Mutual Information (MI) [20], [26] and minimum cross entropy [12], [36] are some of the metrics traditionally used. The use of minimum cross entropy tends to maximise the distance between sensors. In indoor environment problems, this results in having sensors located near the boundary of the domain, i.e. near the walls, thus losing information monitored [36], [26]. On the other hand, information gain or Mutual Information [20] shifts the amount of information captured by a single random variable to the information each random variable has of the other unobserved one. More specifically, considering a finite set of possible placement locations by max-

imising the objective metric, it evaluates how well a given smaller subset of sensor locations describes the values of the unselected other locations. This paper considers the optimisation problem of sensor placement in indoor environments by separating the problem into the learning of a spatial model, i.e. Gaussian Process training, and the optimisation algorithm itself, i.e. optimal sensor placement. It is demonstrated that with a combination of deep learning, probabilistic frameworks and variational methods, the complexity associated with training, inference and optimisation can be significantly reduced in order to achieve optimal placement results. This paper builds on existing 2D sensor placement algorithm [20], [26] and the latest VGP spatial modelling technologies [42], [41]. Its value is found primarily in the pairing of technologies that in turn improve the existing methods of sensor placement.

The choice of the technologies used in this work are detailed and argued in the following points:

- ▷ **Preprocessing input distribution.** A Masked Autoregressive Flow (MAF) is used to normalise the training dataset suffering from being non-Gaussian distributed, [33] making our methodology greatly generalised as well as improving the accuracy.
- ▷ **Spatial model.** A major challenge facing scalable sensor placement is overcome by deploying a variational Gaussian process (VGP), using a low rank approximation that is far more scalable and also addresses the question of model generalisation. In particular, this helps to tackle the limiting $O((mN)^3)$, high polynomial time complexities associated with GPs to $O((mN)\hat{N}^2)$ where \hat{N} denotes the number of approximate posterior samples computed in the VGP. Using a VGP is a good trade-off between efficiency and accuracy [42], [41].
- ▷ **Sensor placement algorithm.** The Mutual Information (MI) based placement algorithm [20], [26] is extended with a Markov-Chain Monte Carlo (MCMC) wrapper to fine-tune the sensor placement and tackle the time complexity $O(N^4)$ and achieve $O(klM^4)$, where k is the number of sensors, l is the number of iteration needed to optimise the position and M is such that $M \ll N$. The use of MCMC leads to similar placement results in a fraction of the computation time required when not using it. Error propagation through the system does exist due to this approximation; however, it is shown in the paper to be a worthwhile trade-off.

The technologies used in this paper are general and are not limited to the test case of sensor placement in indoor environments, even though their integrated implementation was designed accordingly. In fact, the core underlying technology used in the spatial model, i.e. variational Gaussian process, has been successfully deployed for a multitude of other domains ranging from kriging to robotics [10].

In addition to the new pairing of technologies proposed in this paper, the novelties of this work also lie in:

- ▷ **The simulated training data.** In order to achieve scalable and reduced cost deployment of sensor placement optimisation, simulation data was used for model training, replacing the expensive option of collecting large amounts measurements from installed sensors. As an example, the training dataset used in [26] consists of 52 sensors, located on a 2D plane. In this paper, the simulated training dataset is much larger, i.e. more reliable, and consist of 10,000 sensor locations distributed in 3D.
- ▷ **The increase in dimensionality.** This paper increases the dimensionality of the learning problem in order to capture further correlations between hidden features as well as the output features, resulting in 3-dimensional spatial placement. The 3D placement made it possible for the model to capture more realistic and complex physical phenomena such as thermal stratification for example.
- ▷ **The fine-tuning of sensor placement.** The MCMC wrapper algorithm is used to increase the overall Mutual Information captured. The base set of potential sensor location is fine-tuned to include other regions in the continuous space having higher MI. This means the selection pipeline is a more optimal set of possible placement coordinates to choose from. Additionally, the implementation is easy to be customised and makes our methodology generalisable.
- ▷ **The fully parallel and scalable implementation.** The methodology presented in this paper is done through a multi-threaded parallel implementation. The computational graph based implementation, using the TensorFlow library [19], greatly speeds up the computational times.

3. THE VARIATIONAL GAUSSIAN PROCESS FOR OPTIMAL SENSOR PLACEMENT (VGPOSP) MODEL

This section introduces the theoretical concepts and mathematical formulation that were developed and implemented as part of the proposed model architecture.

Let X be the solution of a dynamic system:

$$(3.1) \quad \dot{X} = F(X, t),$$

where t denotes the time and $X = [X_1, \dots, X_m]$ denotes a vector of m state variables¹ such that $X_i \in \mathbb{R}^N$ for all $i = 1, \dots, m$. In the following, $X_t = X(t)$ denotes the

¹For example, as shown in Section 4, for fluid dynamic simulations in indoor environment, $X = [T, P, C]$, where T is the temperature, P the pressure and C the pollutant concentration.

solution of the dynamic system at time t and $X_{i,t} = X_i(t)$ the i th state variable at time t .

Given a temporal sequence X_{t_1}, \dots, X_{t_n} of n solutions of the dynamical system defined in equation (3.1), with $X_{t_i} = [X_{1,t_i}, \dots, X_{m,t_i}]$ for all $i = 1, \dots, n$, the variational Gaussian process for optimal sensor placement (VGPosp) model consists of the following main steps described in the sub-sections 3.1 to 3.3.

▷ Section 3.1 - **Preprocessing and preconditioning.**

- Preconditioning of input vector distributions using Masked Autoregressive Flows (MAF).
- Converting time-series to vector value distributions at distinct spatial regions.

▷ Section 3.2 - **Variational Gaussian process.**

- Sampling algorithm, input vectors spatial training.
- Variational Gaussian process (VGP) training.

▷ Section 3.3 - **Placement algorithm.**

- Selection of a set S of considered input coordinates.
- Generating target vectors with VGP inference at coordinates.
- Covariance matrix of values indexed by set S .
- Greedy selection algorithm of coordinates with maximal Mutual Information (MI).
- Markov-Chain Monte Carlo (MCMC) based fine-tuning of coordinates.

3.1. Pre-processing and preconditioning.

The temporal sequence X_{t_1}, \dots, X_{t_n} requires a pre-processing to be suitable for the spatial model and the inference step during the placement algorithm. The pre-processing consists of the following steps, also described in Figure 1:




Figure 1. Pre-processing and pre-conditioning steps.

- (1) X is first normalised to a mean value of 0 and standardised to achieve a standard deviation of each feature of 1.
- (2) Secondly, a Masked Autoregressive Flow (MAF) is implemented to learn invertible, nonlinear transformations between the non-Gaussian distributed features and the target Gaussian family of distributions.

Firstly, the mean of the temporal sequence for each state variable X_i , $i = 1, \dots, m$, is computed, i.e. the vector

$$(3.2) \quad \bar{X} = [\bar{X}_1, \dots, \bar{X}_m],$$

where

$$(3.3) \quad \bar{X}_i = \frac{\sum_{j=1}^n X_{i,t_j}}{n} \quad \forall i = 1, \dots, m.$$

The vector \bar{X} in equation (3.2) is used to train the Masked Autoregressive Flows (MAF) in order to make our input variables Gaussian distributed. MAF is a neural network that executes a normalising flow nonlinear transformation at each neuron. MAF can be also computed as a stack of autoregressive Masked Autoencoder for Distribution Estimation (MADE) networks [33], [16], where each model uses the vector \bar{X} in equation (3.2). The Autoregressive property of MAF defined from time-series analysis predicts a future value of a variable from a linear combination of its past values. Each MADE learns the distribution of the state variables.

The MAF model can be defined as follows [45]:

$$(3.4) \quad \bar{X}_N = f_N(\bar{X}_{N-1}, \bar{X}_{N-2}, \dots, \bar{X}_1),$$

where f_N has a polynomial form such that [33]:

$$(3.5) \quad \bar{X}_N = \theta_0 + \theta_1 \bar{X}_{N-1} + \theta_2 \bar{X}_{N-2} + \dots + \theta_p \bar{X}_{N-p},$$

where θ_i are the polynomial coefficients.

Normalising flows apply a sequence of N invertible, differentiable transformation functions f_N in $p(\bar{X})$. A base distribution $p(\bar{X}')$ is specified most likely from the family of Gaussian distributions [33], where p denotes the probability distribution. The procedure begins with this initial distribution $p(\bar{X}')$ [38].

$$(3.6) \quad p(\bar{X}) = p(\bar{X}') \left| \det \frac{\partial f^{-1}}{\partial \bar{X}'} \right|.$$

A chain rule can then be applied to the conditionals of a joint distribution.

$$(3.7) \quad p(\bar{X}) = \prod_{i=1}^N p(\bar{X}_i | \bar{X}_{i-1}, \bar{X}_{i-2}, \dots, \bar{X}_1) = \prod_{i=1}^N p(\bar{X}_i | \bar{X}_{<i}).$$

After each nonlinear transformation the distribution becomes more complex. Sampling from this transformed distribution is done via the flow of a straightforward

sample from the original $p(\bar{X})$ Gaussian distribution through the nonlinear transformations. The entropy of the resultant distribution is computed with the logarithm of the transformations:

$$(3.8) \quad \log p(\bar{X}_N) = \ln p(\bar{X}_0) - \sum_{i=1}^N \ln \left| \det \frac{\partial f_N}{\partial \bar{X}_{i-1}} \right|,$$

$$(3.9) \quad p(Y) = \prod_{i=1}^N (f_N^{-1}(Y)) \left| \det \frac{\partial f_N^{-1}}{\partial Y} \right|.$$

The vector

$$(3.10) \quad Y = [Y_1, \dots, Y_m] \in \mathbb{R}^{m \times N} \quad \text{with} \quad Y_j \in \mathbb{R}^N \quad \forall j = 1, \dots, m$$

is the set of normalised state variables which are the input of the variational Gaussian process introduced in next section.

3.2. VGP training and inference. Given the data set $\{Y_j\}_{j=1}^m$ of $Y_j \in \mathbb{R}^N$ as defined in equation (3.10), the m source-target pairs $\mathcal{D} = \{(Y_j, T)\}_{j=1}^m$, where T is a target² [9]. We aim to learn a function over all source-target pairs:

$$(3.11) \quad T = g(Y_j),$$

where $g: \mathbb{R}^N \rightarrow \mathbb{R}^N$ is unknown. Let the function g decouple as $g = (g_1, \dots, g_N)$, where each $g_i: \mathbb{R}^N \rightarrow \mathbb{R}$. A GP regression [37] estimates the functional form of g by placing a prior,

$$(3.12) \quad p(g) = \prod_{i=1}^N \mathcal{GP}(g_i; 0, \Sigma_{ij}),$$

where Σ_{ij} denotes a covariance evaluated over pairs of inputs $Y_i Y_j \in \mathbb{R}^N$ [37]:

$$(3.13) \quad \Sigma_{ij} = \frac{1}{N} \sum_k Y_{ik} Y_{jk}^\top.$$

A variational Gaussian process (VGP) is a Bayesian nonparametric variational model that admits arbitrary structures to match posterior distributions. As described in the following steps, the VGP generates approximate posterior samples Z by generating latent inputs, warping them with random nonlinear mappings, and using

² In some application, $T = Y_j$ for a fixed j can be assumed.

the warped inputs as parameters to a mean-field distribution. The random mappings are drawn conditionally on variational parameters. The VGP specifies a generative process for posterior latent variables Z . At the first step it draws latent input $\xi \in \mathbb{R}^N$: $\xi \sim \mathcal{N}(0, I)$ and a nonlinear mapping $g: \mathbb{R}^N \rightarrow \mathbb{R}^N$ conditioned on \mathcal{D} : $g \sim \prod_{i=1}^N \mathcal{GP}(0, \Sigma_{\xi\xi}) | \mathcal{D}$. Then it draws approximate posterior samples $Z \in \text{supp}(p)$: $Z = (Z_1, \dots, Z_{\hat{N}}) \sim \prod_{i=1}^{\hat{N}} q(g_i(\xi))$. Marginalising over all latent inputs and nonlinear mappings, the VGP is [9]:

$$(3.14) \quad q_{\text{VGP}}(Z; \theta, \mathcal{D}) = \iint \left[\prod_{i=1}^{\hat{N}} q(Z_i | g_i(\xi)) \right] \left[\prod_{i=1}^{\hat{N}} \mathcal{GP}(g_i; 0, \Sigma_{\xi\xi}) | \mathcal{D} \right] \mathcal{N}(\xi; 0, I) \, d\xi \, d\xi.$$

The VGP is parametrised by kernel hyperparameters θ and variational data [9], [37]. The random function interpolates the values in the variational data, which are optimised to minimise the Kullback-Leibler divergence [27]. It defines a measure between two probability density functions: $q_{\text{VGP}}(Z; \theta, \mathcal{D})$ and $q^*(Z|Y)$, where $q^*(Z|Y)$ is the posterior distribution [18].

$$(3.15) \quad \begin{aligned} D_{\text{KL}}(q_{\text{VGP}}(Z; \theta, \mathcal{D}) || q^*(Z|Y)) &= \mathbb{E}_q \left[\log \frac{q_{\text{VGP}}(Z; \theta, \mathcal{D})}{q^*(Z|Y)} \right] \\ &= \mathbb{E}_q [\log q_{\text{VGP}}(Z; \theta, \mathcal{D}) - \log q^*(Z|Y)], \end{aligned}$$

where $\mathbb{E}[f(x, \theta)] = \int_{\theta} f(x, \theta) \, d\theta$, where f denotes a distribution function (see [25]). An approximating distribution is chosen from a predefined family of distributions with parameters: θ ,

$$(3.16) \quad q_{\theta}(Z|Y) = \arg \min_{\theta} D_{\text{KL}}(q_{\text{VGP}}(Z; \theta, \mathcal{D}) || q^*(Z|Y)).$$

When the KL-divergence converges to 0 (see the Theorem of the Universal approximation in [9]), the posterior is approximately the same as the learned posterior. Minimising KL-divergence is done via the objective function, defined by the Evidence Lower Bound (ELBO), parameterised by the parameters θ . The maximisation of this function is equivalent to a minimising KL with a difference measured by a constant factor [38]. The objective is summed over all data-points [25]:

$$(3.17) \quad \text{ELBO}(\theta) = \sum_{i=1}^{\hat{N}} \mathbb{E}_{q_{\theta}(Z|X_i)} [\log q_{\text{VGP}}(Z; C, Y_i) - \log q_{\theta}(Z|Y_i)].$$

When applied to the marginal probability of the evidence [9], [41], the objective lower bound on the marginal likelihood can be quantified by the ELBO:

$$(3.18) \quad \text{ELBO}(\theta) = \mathbb{E}_q [\log q_{\text{VGP}}(Z; \theta, X)] - \mathbb{E}_q [\log q_{\theta}(z|x)].$$

The posterior is therefore the sum of the ELBO and the KL term:

$$(3.19) \quad \log q(Z) = \text{ELBO}(\theta) + D_{\text{KL}}(q_\theta(Z|Y) || q^*(Z|Y)).$$

An approximate solution in the mean-field [43] is sought to learn parameters of the marginal likelihood [7], obtaining an approximate posterior distribution of the true posterior. The mean-field approximation of the variational inference allowed for the approximate $q_\theta(Z|Y)$ distribution to be considered as a factor of \widehat{N} independent latent variable partitions $q_\theta(Z|Y_i)$.

$$(3.20) \quad q^*(Z|Y) \approx q_\theta(Z|Y) = \prod_{i=1}^{\widehat{N}} q_\theta(Z|Y_i)$$

then it yields $g \sim q^*(Z|Y)$ [44].

3.3. Placement algorithm. In this section, the placement algorithm of the VGPosp model is introduced. The algorithm computes optimal coordinates for sensor placement following three main steps as described in Figure 2 and detailed in Algorithm 1.

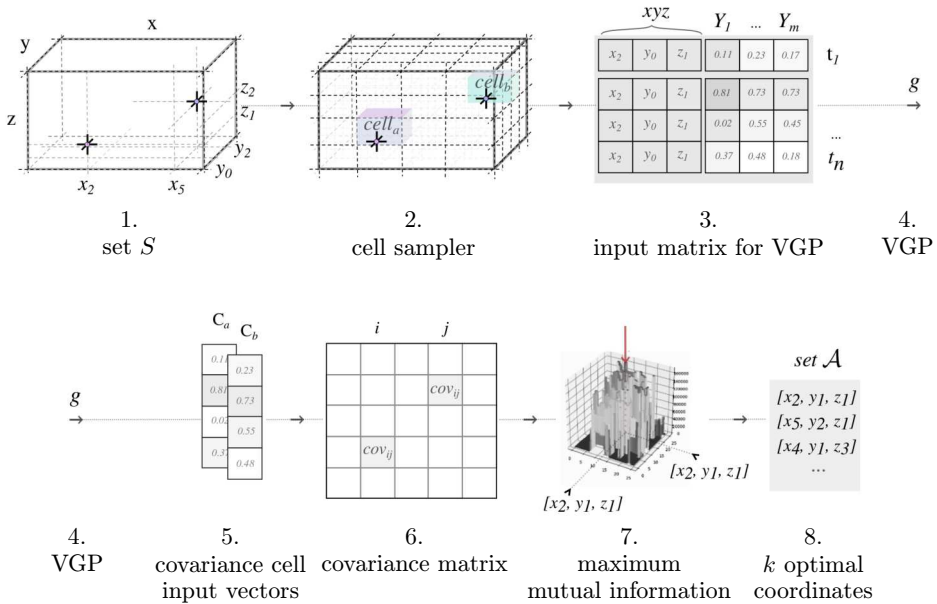


Figure 2. Graphical representation of the main steps of the variational Gaussian process for optimal sensor placement (VGPosp) model.

Algorithm 1. THE VARIATIONAL GAUSSIAN PROCESS FOR THE OPTIMAL SENSOR PLACEMENT (VGPOSP) ALGORITHM.

input: number of sensors: k , space of coordinates: S , maximum number of iterations: l ; ε

$\mathcal{A}^* = \emptyset$ ▷ initialise the set of optimal sensor locations
 $\Sigma = \emptyset$ ▷ initialise the covariance matrix
 $\delta_{\bar{y}} = 0$ ▷ initialise the mutual information parameter

```

while  $ii < k$  do
  while  $it < l$  do
    ▷ compute the covariance matrix
    for  $i \in S$  do
      for  $t_k \in [t_0, t_n]$  do
         $T_{ik} = g(Y_{ik})$  ▷ VGP function in 3.11
       $T_i = [T_{i0}, T_{i1}, \dots, T_{in}]$ 
      for  $j \in S$  do
        for  $t_k \in [t_0, t_n]$  do
           $T_{jk} = g(Y_{jk})$  ▷ VGP function in 3.11
         $T_j = [T_{j0}, T_{j1}, \dots, T_{jn}]$ 
      ▷ using equation (3.13)
      compute  $\Sigma_{ij}$ 
     $\mathcal{A}, \delta_{i^*} \leftarrow$  Algorithm 2 ( $\Sigma, S, k$ ) ▷ estimate the mutual information
    if  $\delta_{i^*} > \delta_{\bar{y}}$  or  $\mathcal{A}^* = \emptyset$  then
       $\mathcal{A}^* = \mathcal{A}$ 
       $\delta_{\bar{y}} = \delta_{i^*}$ 
     $ii++$ 
   $it++$ 
return  $\mathcal{A}^*$ 

```

The set of coordinates \mathcal{V} initially consists of N grid points: $|\mathcal{V}| = N$. Our placement algorithm is mainly based on the Mutual Information (MI) based placement algorithm [20], [26] which is extended in this paper with a MCMC wrapper to fine-tune the sensor placement and tackle the time complexity $O(N^4)$ and achieve $O(klM^4)$, where k is the number of sensors, l is the number of iterations needed to optimise the position and M is a predefined small subset of N such that $M \ll N$. In fact, the first step of Algorithm 1 consists of identifying all the possible locations which

constitute a set S , $|S| = M$ (see points 1 and 2 in Figure 2). The set S is an input of Algorithm 1. Other inputs are the number k of sensors to place and the number l of maximum iterations for the optimisation process.

We implement an optimisation method to identify a set \mathcal{A} as the placement output from the predefined set S of input coordinates [20] such that $|\mathcal{A}| = k$. The second step of Algorithm 1 consists in sampling m state variables from value distributions in Y at corresponding spatial regions/cells (see point 3 in Figure 2). For each state variable in Y (see point 4 in Figure 2), the samples T are produced by the function g which is computed by a VGP as described in Section 3.2 (see points 1–2 in Algorithm 1):

$$\forall j \in S, \quad T_{jk} = g(Y_{jk}), \quad T_j = [T_{j0}, T_{j1}, \dots, T_{jn}].$$

Then the covariance matrix of the values T_j is computed for the locations specified in S using equation (3.13) (see points 5 and 6 in Figure 2). We also define the covariance matrices related to a subset \mathcal{A} such that:

$$(3.21) \quad \Sigma_{i\mathcal{A}} = \begin{bmatrix} \Sigma_{i1} & 0 & 0 & \dots \\ 0 & \Sigma_{i2} & 0 & \dots \\ \vdots & \ddots & \ddots & \vdots \\ \dots & 0 & 0 & \Sigma_{ik} \end{bmatrix},$$

where Σ_{ij} is defined in equation (3.13) and $k = |\mathcal{A}|$. $\Sigma_{i\mathcal{A}}$ is used in the final step of Algorithm 1 which maximise the mutual information (see points 7 and 8 in Figure 2), as described in Algorithm 2, where H is the conditional entropy function defined as (see [26]):

$$(3.22) \quad H(i|\mathcal{A}) = \frac{1}{2} \log \Sigma_{i\mathcal{A}}^2 + \frac{1}{2} (\log(2\pi) + 1)$$

and where δ_i denotes the mutual information parameter [26]:

$$(3.23) \quad \delta_i = \frac{\Sigma_{ii}^2 - \Sigma_{i\mathcal{A}} \Sigma_{\mathcal{A}\mathcal{A}}^{-1} \Sigma_{\mathcal{A}i}}{\Sigma_{ii}^2 - \Sigma_{i\bar{\mathcal{A}}} \Sigma_{\bar{\mathcal{A}}\bar{\mathcal{A}}}^{-1} \Sigma_{\bar{\mathcal{A}}i}}.$$

We developed a MCMC based wrapper that works with a smaller grid of points S , then identifies the vertices of the grid \mathcal{A} that are potential frontiers to be explored further. These optimal coordinates in S are thus adjusted. If the adjustment has improved the overall mutual information, we keep the modification. Using this technique we iteratively adjust S until the mutual information value converges.

Algorithm 2. MAXIMISE MUTUAL INFORMATION (MI) USING LAZY EVALUATIONS.

input: Covariance matrix Σ_{ij} , number of sensors k , set of coordinates S

$\mathcal{A} \leftarrow \emptyset$

```
| for each  $i \in S$  do
|  $\delta_i \leftarrow \infty$ 
|
| for  $j = 1, \dots, k$  do
|   | for each  $i \in \bar{\mathcal{A}} = S \setminus \mathcal{A}$  do
|   |   |  $\text{current}_i \leftarrow \text{false}$ 
|   |   | while  $\text{current}_i == \text{true}$  do
|   |   |   |  $i^* \leftarrow \text{argmax}_{i \in S \setminus \mathcal{A}} \delta_i$ 
|   |   |   |   | if  $\text{current}_{i^*} == \text{true}$  then
|   |   |   |   |   | break
|   |   |   |   |   | Compute  $\Sigma_{i\mathcal{A}}$  ▷ using equation (3.21) for  $i$  and  $\mathcal{A}$ 
|   |   |   |   |   | Compute  $\Sigma_{i\bar{\mathcal{A}}}$  ▷ using equation (3.21) for  $i$  and  $\bar{\mathcal{A}}$ 
|   |   |   |   |   |  $\delta_{i^*} \leftarrow H(i|\mathcal{A}) - H(i|\bar{\mathcal{A}})$  ▷ with  $H$  defined in (3.22)
|   |   |   |   |   |  $\text{current}_{i^*} \leftarrow \text{true}$ 
|   |   |   |  $\mathcal{A} \leftarrow \mathcal{A} \cup i^*$ 
|   |  $\mathcal{A} \leftarrow \mathcal{A} \cup i^*$ 
| return  $\mathcal{A}$ 
```

3.4. Implementation. A code implementing Algorithm 1 and Algorithm 2 using TensorFlow.1.4 [19] is available at <https://github.com/roxarcucci/VGPosp.git>. Instructions for running the tests and algorithms are described in the file README.md. The algorithms were initially implemented in Python and reimplemented later in TensorFlow in order to further improve the efficiency of the run-time through multi-threaded, parallelised execution and automatic scalability to more computational cores. Our implemented model is represented in the form of a computational data-flow graph that is instantiated once a session object is defined. The built-in TensorFlow compiler identifies all dependencies within our algorithms and assigns multi-threaded computational tasks to our resources. TensorBoard creates a visual representation of the nodes and connections, it enables the developer to debug connectivity errors. The Scalars tool that allows the tracking of any metric of interest during model training or optimisation was also used.

4. RESULTS AND DISCUSSION

In this section, two test cases are used to discuss our Variation Gaussian Process for optimal sensor placement (VGPosp):

- ▷ The first test case, named the sine model, is a simplified two-dimensional model of a sine function to test the efficiency, accuracy and precision of VGP compared to GP. The comparison between GP and VGP is done only for the sine model as the complexity of GP is too high to be trained for a real test case.
- ▷ The second test case is a real three-dimensional test case considering a room within the Clarence Centre building located in Elephant and Castle, London, UK. In this test case, the predictive model is the Computational Fluid Dynamics (CFD) software Fluidity. The optimal sensor location is proposed. The benefit of using sensors optimally located is proved by showing how the predictive model error can be more efficiently reduced by using data assimilation technology.

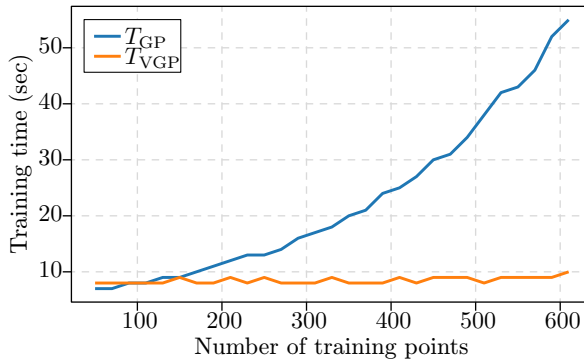


Figure 3. Training time as a function of number of training points when using a Gaussian Process (GP) or a variation Gaussian process (VGP).

4.1. Test case 1: Sine model. This section aims to compare the efficiency, the accuracy and the precision of GP and VGP. The test sine function used is defined in equation (4.1).

$$(4.1) \quad f(x, y) = \sin\left(\frac{2}{3}\pi x\right) + \sin\left(\frac{2}{3}\pi y\right).$$

In order to prove that VGP is more efficient than a GP approach, the training time as a function of the number of training points for both methods is shown in Figure 3. Up to 150 training points, the GP and the VGP method both take a bit less than 10 sec to be trained. However, when using more training points, i.e. more than 150

points, the training time of the GP increases drastically, while the VGP training time stays constant, around 10 sec, independently of the number of points. For example, when considering 600 training points, the training time is divided by 5.5 when using the VGP approach.

The accuracy e and the root mean squared error (also called precision), RMSE of the model \mathcal{M} are:

$$(4.2) \quad e_{\mathcal{M}}(N) = \sum_{i=1}^N |f_{\text{True}}(x_i, y_i) - f_{\mathcal{M}}(x_i, y_i)|,$$

where N is the number of training points, and

$$(4.3) \quad \text{RMSE}_{\mathcal{M}}(N) = \sqrt{\frac{\|F_{\mathcal{M}} - F_{\text{True}}\|_{L^2}}{\|F_{\text{True}}\|_{L^2}}},$$

where L^2 denotes the Euclidean norm, $F_{\mathcal{M}}$ and F_{True} denote the vectors $F_{\mathcal{M}} = [f_{\mathcal{M}}(x_1, y_1), \dots, f_{\mathcal{M}}(x_N, y_N)]$ and $F_{\text{True}} = [f_{\text{True}}(x_1, y_1), \dots, f_{\text{True}}(x_N, y_N)]$, the True model denotes the function in equation (4.1) and the model \mathcal{M} stands for GP or VGP.

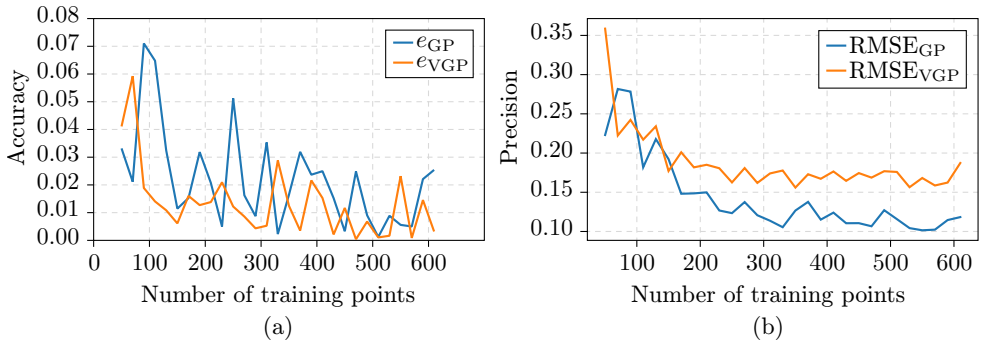


Figure 4. (a) Accuracy e and (b) Precision RMSE as a function of number of training points when using a Gaussian process (GP) and a variation Gaussian process (VGP).

The accuracy (equation (4.2)) and the precision (equation (4.3)) of the two models as a function of the number of training points are shown in Figure 4. From Figure 4a, when using less than 150 training points, the two models highlight the worst accuracy, i.e. the highest values. When using more than 150 training points, it can be seen that the accuracy of GP is lower than 0.03, while the VGP accuracy is lower than 0.02. Globally, the VGP method is slightly more accurate than GP even if the accuracy can be considered as the same order of magnitude. However, looking at Figure 4b,

the GP model is more precise than the VGP. For both models, the RMSE is relatively high when using less than 200 training points. When using more training points, the RMSE reaches a plateau with values for the GP and VGP model of about 0.12 and 0.17, respectively.

Overall, it has been shown that the VGP method is a good trade-off between efficiency, accuracy and precision and will then be used as the assumed tool in the second test case.

4.2. Test case 2: Real test case.

4.2.1. Predictive model: Computational Fluid Dynamics simulation using Fluidity software. The simulated data used to train our VGPosp is obtained using Fluidity, a parallel open-source CFD software (available at <http://fluidityproject.github.io/>). It uses finite elements to solve the following incompressible three dimensional Navier-Stokes equations, continuity equation (4.4) and momentum equation (4.5), on unstructured grids [2]:

$$(4.4) \quad \nabla \cdot \bar{\mathbf{u}} = 0,$$

$$(4.5) \quad \frac{\partial \bar{\mathbf{u}}}{\partial t} + \bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}} = -\frac{1}{\varrho} \nabla \bar{p} + \nabla \cdot [(\nu + \nu_\tau) \nabla \bar{\mathbf{u}}],$$

where $\bar{\mathbf{u}}$ is the resolved velocity (m/s), \bar{p} is the resolved pressure (Pa), ϱ is the fluid density (kg/m³), ν is the kinematic viscosity (m²/s) and ν_τ is the anisotropic eddy viscosity (m²/s).

Turbulence is resolved using Large Eddy Simulation (LES), where the eddies smaller than a scale Δ are parametrised using a subgrid-scale module, while the larger eddies are fully resolved. The subgrid-scale model in Fluidity is based on the Smagorinsky model [39], [8].

The transport of a scalar field C (i.e., a passive tracer or pollutant concentration) in kg/m³ is expressed using the advection-diffusion equation (4.6):

$$(4.6) \quad \frac{\partial C}{\partial t} + \nabla \cdot (\mathbf{u}C) = \nabla \cdot (\overline{\kappa_C} \nabla C) + F,$$

where \mathbf{u} is the velocity vector (m/s), $\overline{\kappa_C}$ is the diffusivity tensor of the pollutant in an excess of air (m²/s) and F represents the source terms (kg/m³/s).

The temperature field T (Kelvin) is expressed using equation (4.7):

$$(4.7) \quad \frac{\partial T}{\partial t} + \nabla \cdot (\mathbf{u}T) = \nabla \cdot (\overline{\kappa_T} \nabla T) + \frac{Q}{\varrho c_p},$$

where \mathbf{u} is the velocity vector (m/s), $\overline{\kappa_T}$ is the thermal diffusivity tensor (m²/s), Q represents thermal source terms (W/m³), ρ is the fluid density (kg/m³) and c_p is the fluid specific heat capacity (J/kg/K). The behaviour of the atmospheric boundary layer in Fluidity is represented using a turbulent inlet velocity based on a synthetic eddy method [34], [23]. Fluidity uses mesh adaptivity where the mesh can be dynamically refined, during the simulation, in areas of physical significance to the user [32].

4.2.2. Test case set up description. The test case considered in this paper is a room within the Clarence Centre building located at London South Bank University (LSBU) near Elephant and Castle in London, UK (Figure 5). The test room has three windows depicting in blue in Figure 5. This test site was used to conduct a one-day field study in January 2018 with the MAGIC project (<http://www.magic-air.uk/>, [40]) during which 7 sensors were monitoring the indoor temperature and CO₂ concentration. The CO₂ sensor (Senseair AB, Sweden) operates on a non-dispersion infrared method that determines CO₂ concentration based on light absorption [31]. As the model is completely general, it could be applied to another kind of concentration analysis that we can find in indoor scenarios such as radon, benzene, NO₂ and others as listed into [24]. The CFD simulation performed in this paper aims to replicate a cross ventilation scenario, where the test room windows on both sides of the building were opened.




Figure 5. The test case room is located in Clarence Centre building in London, UK. The red dots denote the location of sensors during a field experiment and blue rectangles show the location of the three windows.

As shown in Figure 6, the computational domain considered in the numerical simulations includes the entire Clarence building and the test room, as well as the immediate building upwind in order to replicate the local flow conditions near the windows. The mesh is defined such that the resolution is increased in the room

(setting the grid edge length to 0.1 m) and particularly at the openings (grid edge length set to 0.02 m). It progressively decreases in the overall domain to reach an edge length of 10 m away from the room as shown in Figure 6, which gives an overall number of 285,700 cells, i.e. grid points, in the mesh. The total number of nodes within the room is about 1.5×10^5 .




Figure 6. Computational domain and surface mesh of the area of interest showing the Clarence Centre and the upwind building as well as the test case room. The blue arrows denote the wind direction.

The boundary conditions are set to replicate the experimental conditions. A log-law turbulent inlet velocity is imposed upwind, corresponding to a wind direction of 201° . It is parametrised with an incoming wind velocity of 2.58 m/s at 28.5 m. No slip boundary conditions are imposed at the bottom of the domain and on the walls of the test room. The initial temperature is set to 19.5°C inside the room and 9.1°C outside. Before opening the windows, based on sensor data, the average CO_2 concentration in the room is set to $2.58 \times 10^{-3} \text{ kg/m}^3$ (1420 ppm) while $7.2 \times 10^{-4} \text{ kg/m}^3$ (400 ppm), is prescribed outside as a background pollution level.

The simulation was run in parallel on 20 CPUs and for an overall simulation time of 15 min, leading to about 3500 timesteps. In this paper, the target variable is the concentration C of CO_2 within the room. As an example, the evolution of the concentration field on different planes in the room within the room at different times is shown in Figure 7 and Figure 8. At the beginning of the simulation, the outdoor air enters the room gradually and concentration stratification starts to occur after 2 minutes 30 seconds. It can be seen that the concentration in the room starts to reach a steady state after 5 minutes and does not change anymore after 15 minutes of simulations.

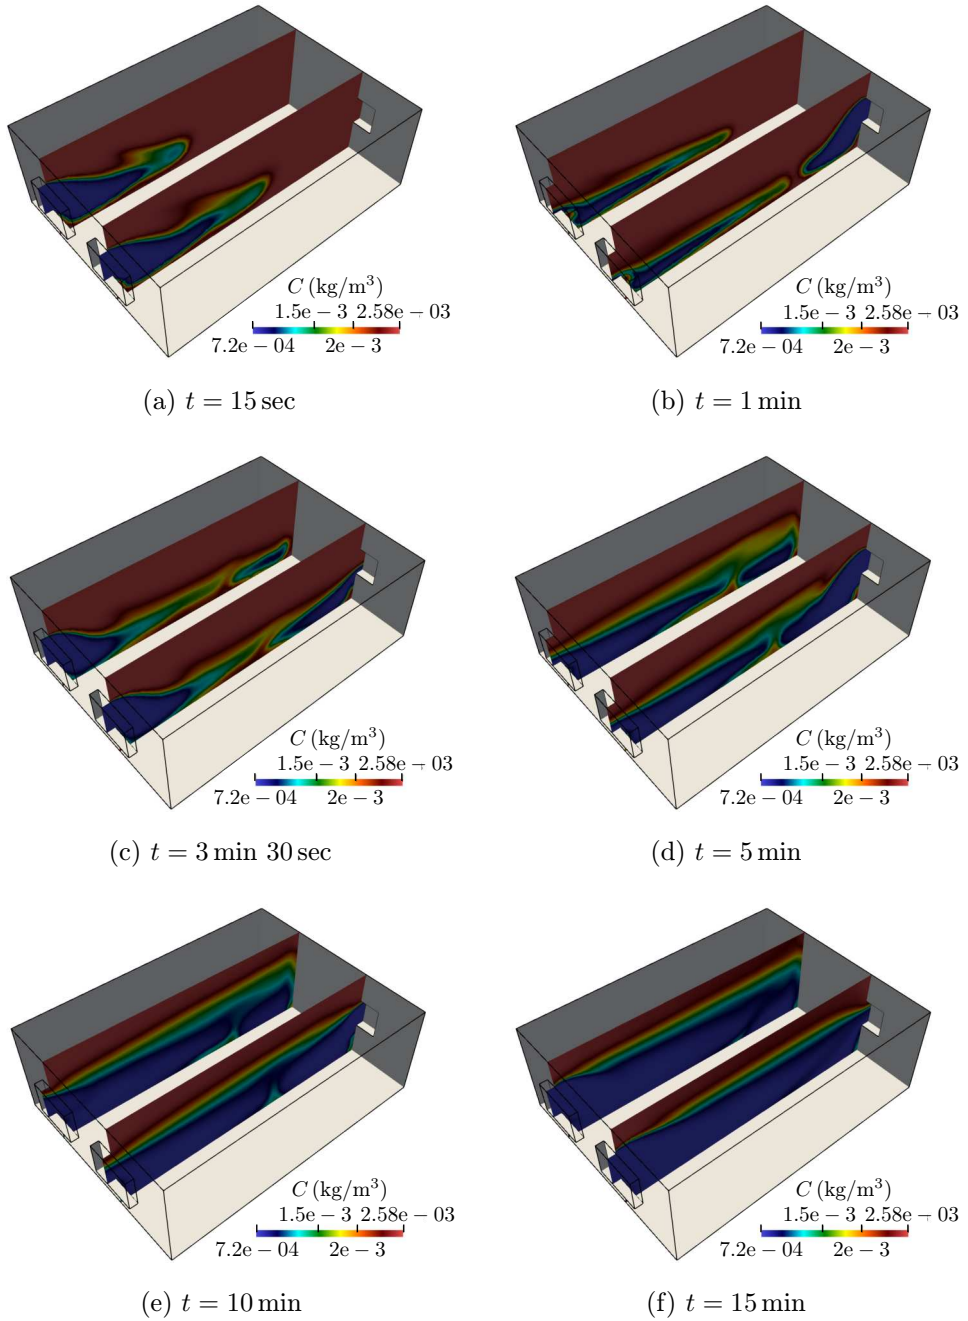


Figure 7. Concentration field of CO_2 on two vertical slices in the room at different times. The scale is between $7.2 \times 10^{-4} \text{ kg/m}^3$ (blue colour) and $2.58 \times 10^{-3} \text{ kg/m}^3$ (red colour).

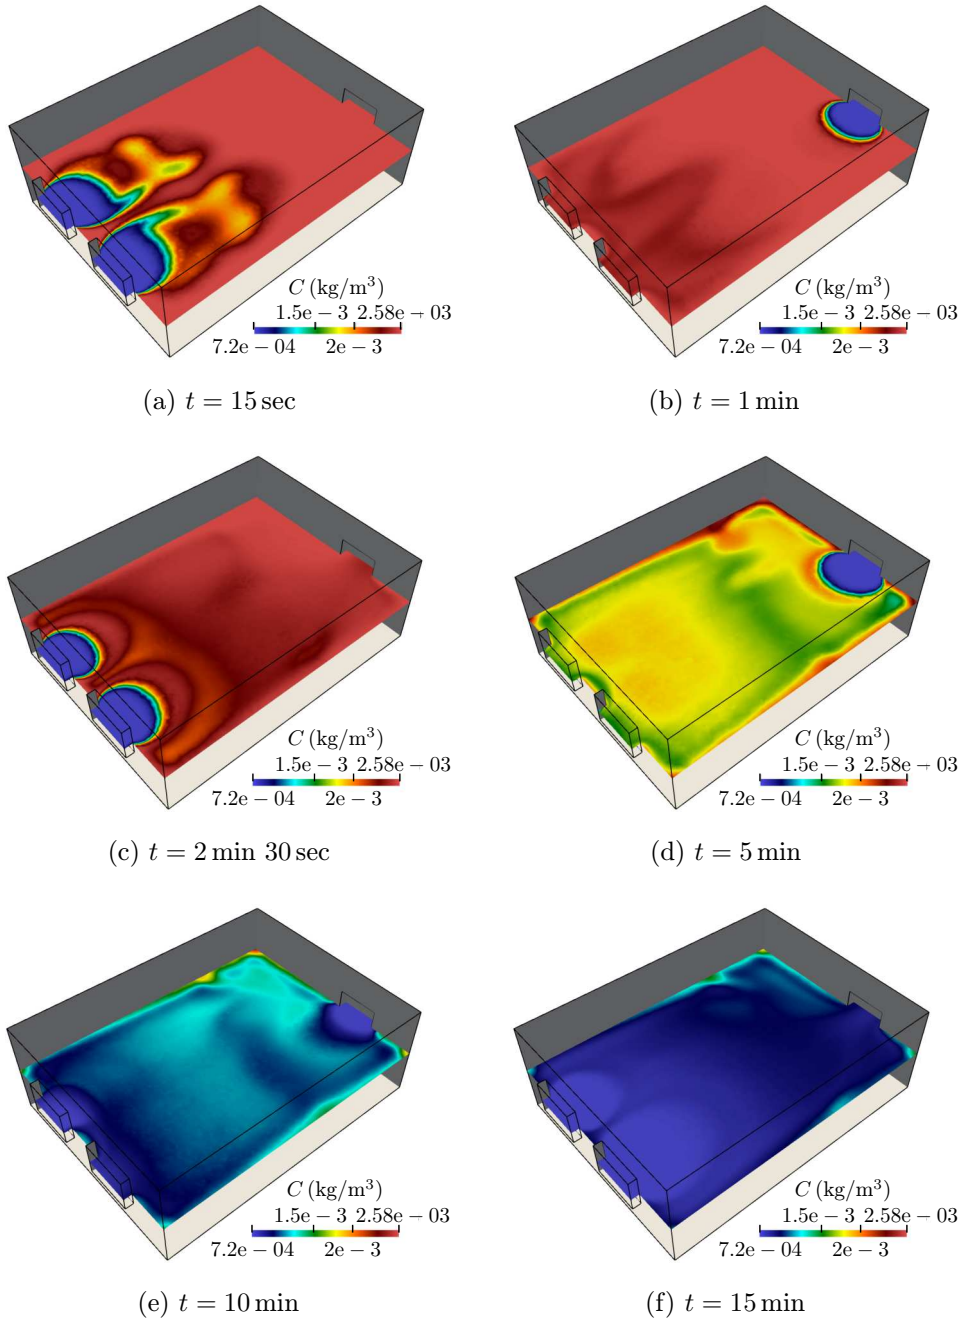
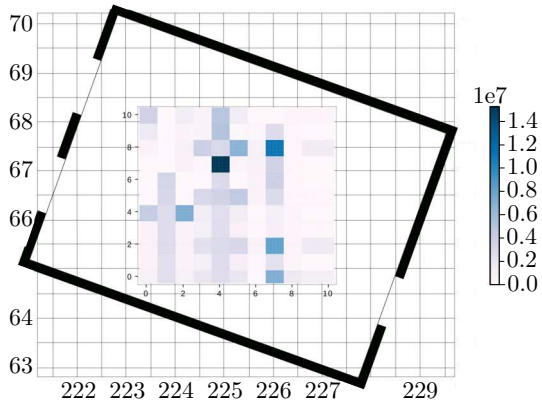




Figure 8. Concentration field of CO₂ on a horizontal plane in the room at different times. The scale is between $7.2 \times 10^{-4} \text{ kg/m}^3$ (blue colour) and $2.58 \times 10^{-3} \text{ kg/m}^3$ (red colour).




(a) Sensor 1




(b) Sensor 2



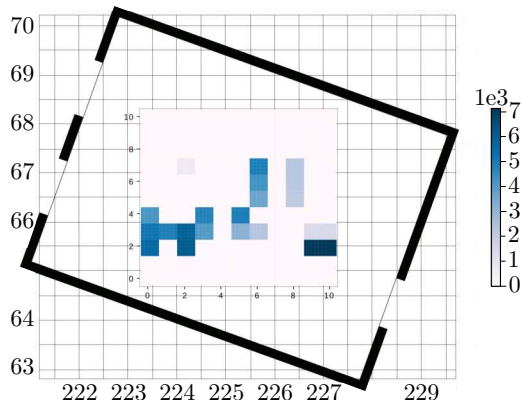
(c) Sensor 3



(d) Sensor 4



(e) Sensor 5



(f) Sensor 6




Figure 9. Mutual Information parameter δ_i on M grid points ($11 \times 11 \times 1$) for $k = 7$ sensors. The colours show the value of δ_i and the scale is different on each sub-figures. Dark blue and light blue colours denote high and low δ_i values, respectively.

4.2.3. Results. The training set consists of 10,000 points, exceeding the amount used in previous work by a factor of 150, see [20], [26]. The quantity and 3D spatial positioning of the training set was sufficient to capture the dynamics of the indoor environment. The CFD results used to train the VGP are taken between 2 min 30 sec and 5 min, period during which the stratification of the concentration is established. The solution of this dynamical system X in (3.1) includes the physical variables $[P, T, C]$, where P is the pressure, T the temperature and C , the physical variable target, is the CO_2 concentration.

The simulated training data had features with non-Gaussian likelihoods, causing potential problems for spatial learning with Gaussian Processes. While the use of a variational Gaussian process helps overcome this issue, we further generalised our training data through the use of a Masked Autoregressive Flow model that transforms the likelihoods of the input features to the family of Gaussian-distributions.

The results of the optimal sensor placement are discussed in the following. Firstly, the scalability issues related to dense grid initialisations of set S are addressed. Secondly, the introduction of the MCMC based fine-tuning algorithm is motivated. Finally, the optimal sensor placement solutions provided by our proposed VGPosp model are presented.

The placement Algorithm 2 is executed on a predefined area of interest, in which a density parameter specifies the grid initialisation that defines set S . Figure 9 demonstrates the internal state, i.e. mutual information parameter, of the placement

algorithm before making a selection. For each coordinate δ_i is computed (Algorithm 2) and the most optimal coordinate is selected into the final selection set \mathcal{A} . For example, in Figure 12a, the first sensor will be chosen to be located where δ_i is the highest, i.e. yellow colour part. From Figure 9, it can be seen that the scale of δ_i is not the same in each sub-figures and becomes narrower. Indeed, the contributions made by earlier selections are higher, explaining why the δ_i quantities decrease after more sensors are added to set S .




Figure 10. Execution time t of Algorithm 1 as a function of the number of initial grid points M . The y -axis is logarithmic.

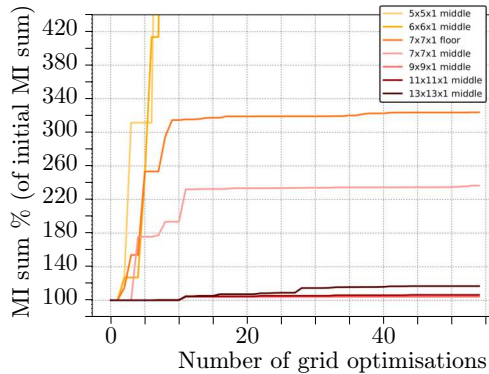


Figure 11. Percentage increase in mutual information parameter after the fine-tuning Algorithm 1.

The running-times shown in Figure 10 are descriptive of the exponential computational cost that is incurred from selecting a larger input set of S . However, it cannot be expected from the selection Algorithm 1 to find the optimal locations in continuous space when defining only distant, discrete points. In order to reduce the time complexity of discovering optimal coordinates in continuous space, we proposed an MCMC-based fine-tuning method of set S . The impact of this procedure

is demonstrated by Figure 11, where the percentage increase in the optimisation criterion, which may miss more optimal regions. For example in the case of a $7 \times 7 \times 1$ grid instantiation, the fine-tuning Algorithm 1 can achieve a 3 order of magnitude

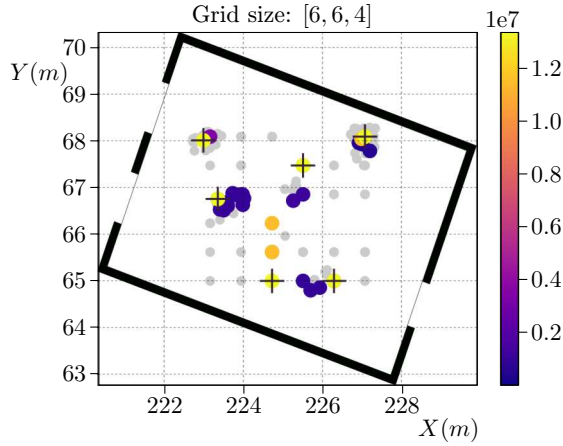


Figure 12. Fine-tuning of grid-points and final selection coordinates. Colour represents the mutual information parameter of placement. The colour of the point shows the value of the Mutual Information. The crosses depict the final optimal sensor locations.




Figure 13. Optimal final location of seven sensors in the room of the Clarence Centre obtained using the variational Gaussian process optimal sensor placement VGPosp in Algorithm 1.

improvement in the overall mutual information, associated with our final selection set *A*. This improvement is achieved after 15 grid optimisation attempts. In Figure 12, a $6 \times 6 \times 4$ grid is considered associated with 40 grid optimisation attempts for each selection. The optimal selection of 7 sensors is plotted with crosses as well

as all modified grid positions that had achieved improved mutual information parameter values. Their colours correspond to the mutual information parameter that the coordinate had achieved. The final optimal coordinates of the room are depicted in Figure 13.

Data Assimilation technology is coupled with the predictive model Fluidity. Data Assimilation uses observed data from sensors to improve and correct the numerical results from the simulation. Seven sensors are assimilated in the predictive model. The DA algorithm and methodology used was previously successfully coupled with Fluidity and is presented in detail in [5], [6]. The accuracy of the DA results is evaluating using the mean squared error:

$$\text{MSE}(C) = \frac{\|C - C^{v,n}\|_{L^2}}{\|C^{v,n}\|_{L^2}},$$

where C is either C^n the Fluidity concentration at time step n , or C^{DA} , the corrected concentration using DA and $C^{v,n}$ is the control variable, i.e. the true observed data. The MSE is computed using the 7 optimal sensor locations shown in Figure 13 and compared with the MSE obtained using 7 sensors located randomly. 2000 random sensor positionings were performed. Assimilating the seven optimally positioned sensors, the error of the predictive model, i.e. Fluidity, is reduced by up to three orders of magnitude: $\text{MSE}(C^n) = 0.17$ and $\text{MSE}(C^{\text{DA}}) = 0.0005$. Moreover, this error is up to two orders of magnitude lower than the ones computed using random sensors placement. In one of the worst random case scenario: $\text{MSE}(C^{\text{DA}}) = 0.023$.

5. CONCLUSION

This work described a novel pipeline for sensor placement, incorporating Masked Autoregressive Flows (MAF) [33] for preconditioning, a variational Gaussian process (VGP) [41], [42] spatial model and a mutual information-based placement algorithm [26]. In this paper, the alterations to the existing placement pipelines are significant, as they introduce multiple layers of approximations in order to reduce time complexities. More specifically, VGP was introduced for the sensor placement pipeline to tackle the $O(N^3)$ complexity associated with traditional GP for spatial modelling. Increased model generalisation was achieved with MAF that learn nonlinear invertible transformations between complicated transformations and the more flexible Normal distribution. All models and algorithms were implemented as

a TensorFlow computational graph to further reduce run-times as opportunities for parallel computation are automatically recognised by the graph compiler. Furthermore, this work proposed and developed two extension algorithms. One focused on incorporating information and sampling environmental features from the time-series for computing mutual information. Secondly, a wrapper algorithm was built to iteratively sub-sample and globally optimise the instantiated set of base grid-coordinates, leading to a three-fold increase in mutual information associated with the final selection. The combination of these two algorithms achieved stability and a global improvement in the selection coordinates. The technologies used in this paper are general and are not limited to the test case of sensor placement in indoor environments, even though their integrated implementation was designed accordingly. The model can be applied in any kind of building, taking into account different concentrations of indoor pollution.




ACRONYMS

CFD:	Computational Fluid Dynamics
DA:	Data Assimilation
ELBO:	Evidence Lower BOund
GP:	Gaussian Process
IAQ:	Indoor Air Quality
KL:	Kullback-Leibler
LES:	Large Eddy Simulation
LSBU:	London South Bank University
MADE:	Masked Autoencoder for Distribution Estimation
MAF:	Masked Autoregressive Flows
MCMC:	Markov-Chain Monte Carlo
MI:	Mutual Information
NF:	Normalising Flow
SGP:	Sparse Variational Process
SVI:	Stochastic Variational Inference
TF:	TensorFlow
VGP:	Variational Gaussian Process
VGPosp:	Variational Gaussian Process optimal sensor placement

References

- [1] *K. Abhishek, M. P. Singh, S. Ghosh, A. Anand*: Weather forecasting model using artificial neural network. *Procedia Technology* 4 (2012), 311–318. doi
- [2] *Applied Modelling and Computation Group*: Fluidity manual (Version 4.1). Available at https://figshare.com/articles/Fluidity_Manual/1387713 (2015), 329 pages.
- [3] *R. Arcucci, L. D'Amore, J. Pistoia, R. Toumi, A. Murlì*: On the variational data assimilation problem solving and sensitivity analysis. *J. Comput. Phys.* 335 (2017), 311–326. zbl MR doi
- [4] *R. Arcucci, D. McLwraith, Y.-K. Guo*: Scalable weak constraint Gaussian processes. *Computational Science – ICCS 2019. Lecture Notes in Computer Science* 11539. Springer, Cham, 2019, pp. 111–125. MR doi
- [5] *R. Arcucci, L. Mottet, C. Pain, Y.-K. Guo*: Optimal reduced space for variational data assimilation. *J. Comput. Phys.* 379 (2019), 51–69. MR doi
- [6] *E. Aristodemou, R. Arcucci, L. Mottet, A. Robins, C. Pain, Y.-K. Guo*: Enhancing CFD-LES air pollution prediction accuracy using data assimilation. *Building and Environment* 165 (2019), Article ID 106383, 15 pages. doi
- [7] *M. J. Beal*: Variational Algorithms for Approximate Bayesian Inference: A Thesis Submitted for the Degree of Doctor of Philosophy of the University of London. University of London, London, 2003.
- [8] *J. H. T. Bentham*: Microscale Modelling of Air Flow and Pollutant Dispersion in the Urban Environment: Doctoral Thesis. University of London, London, 2004.
- [9] *D. M. Blei, A. Kucukelbir, J. D. McAuliffe*: Variational inference: A review for statisticians. *J. Am. Stat. Assoc.* 112 (2017), 859–877. MR doi
- [10] *B. Bócsi, P. Hennig, L. Csató, J. Peters*: Learning tracking control with forward models. *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, New York, 2012, pp. 259–264. doi
- [11] *D. Cornford, I. T. Nabney, C. K. I. Williams*: Adding constrained discontinuities to Gaussian process models of wind fields. *Advances in Neural Information Processing Systems* 11 (NIPS 1998). MIT Press, Cambridge, 1999, pp. 861–867.
- [12] *N. Cressie*: Statistics for spatial data. *Terra Nova* 4 (1992), 613–617. doi
- [13] *L. D'Amore, R. Arcucci, L. Marcellino, A. Murlì*: A parallel three-dimensional variational data assimilation scheme. *Numerical Analysis and Applied Mathematics, ICNAAM 2011. AIP Conference Proceedings* 1389. AIP, Melville, 2011, pp. 1829–1831. zbl doi
- [14] *C. Doersch*: Tutorial on variational autoencoders. Available at <https://arxiv.org/abs/1606.05908> (2016), 23 pages.
- [15] *T. H. Dur, R. Arcucci, L. Mottet, M. Molina Solana, C. Pain, Y.-K. Guo*: Weak constraint Gaussian processes for optimal sensor placement. *J. Comput. Sci.* 42 (2020), Article ID 101110, 12 pages. MR doi
- [16] *M. Germain, K. Gregor, I. Murray, H. Larochelle*: MADE: Masked Autoencoder for Distribution Estimation. *Proc. Mach. Learn. Res.* 37 (2015), 881–889.
- [17] *H. González-Banos*: A randomized art-gallery algorithm for sensor placement. *SCG'01: Proceedings of the 17th Annual Symposium on Computational Geometry*. ACM, New York, 2001, pp. 232–240. zbl doi
- [18] *I. Goodfellow, Y. Bengio, A. Courville*: *Deep Learning*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, 2016. zbl MR
- [19] *Google Brain Team*: TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. Available at <https://www.tensorflow.org/> (2015). sw
- [20] *C. Guestrin, A. Krause, A. P. Singh*: Near-optimal sensor placements in Gaussian processes. *ICML'05: Proceedings of the 22nd International Conference on Machine Learning*. ACM, New York, 2005, pp. 265–272. doi

- [21] *J. Hagan, A. R. Gillis, J. Chan*: Explaining official delinquency: A spatial study of class, conflict and control. *Sociological Quarterly* 19 (1978), 386–398. doi
- [22] *J. Hensman, N. Fusi, N. D. Lawrence*: Gaussian processes for big data. Available at <https://arxiv.org/abs/1309.6835> (2013), 9 pages.
- [23] *N. Jarrin, S. Benhamadouche, D. Laurence, R. Prosser*: A synthetic-eddy-method for generating inflow conditions for large-eddy simulations. *Int. J. Heat Fluid Flow* 27 (2006), 585–593. doi
- [24] *F. J. Kelly, J. C. Fussell*: Improving indoor air quality, health and performance within environments where people live, travel, learn and work. *Atmospheric Environment* 200 (2019), 90–109. doi
- [25] *D. P. Kingma, M. Welling*: Auto-encoding variational Bayes. Available at <https://arxiv.org/abs/1312.6114> (2013), 14 pages.
- [26] *A. Krause, A. Singh, C. Guestrin*: Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. *J. Mach. Learn. Res.* 9 (2008), 235–284. zbl
- [27] *S. Kullback, R. A. Leibler*: On information and sufficiency. *Ann. Math. Stat.* 22 (1951), 79–86. zbl MR doi
- [28] *C.-C. Lin, L. L. Wang*: Forecasting simulations of indoor environment using data assimilation via an ensemble Kalman filter. *Building and Environment* 64 (2013), 169–176. doi
- [29] *H. Liu, Y.-S. Ong, X. Shen, J. Cai*: When Gaussian process meets big data: A review of scalable GPs. Available at <https://arxiv.org/abs/1807.01065> (2018), 20 pages.
- [30] *D. J. C. MacKay*: Introduction to Gaussian processes. *Neural Networks and Machine Learning*. NATO ASI Series F Computer and Systems Sciences 168. Springer, Berlin, 1998, pp. 133–166.
- [31] *M. I. Mead, O. A. M. Popoola, G. B. Stewart, P. Landshoff, M. Calleja, M. Hayes, J. J. Baldovi, M. W. McLeod, T. F. Hodgson, J. Dicks, A. Lewis, J. Cohen, R. Baron, J. R. Saffell, R. L. Jones*: The use of electrochemical sensors for monitoring urban air quality in low-cost, high-density networks. *Atmospheric Environment* 70 (2013), 186–203. doi
- [32] *C. C. Pain, A. P. Umpheby, C. R. E. de Oliveira, A. J. H. Goddard*: Tetrahedral mesh optimisation and adaptivity for steady-state and transient finite element calculations. *Comput. Methods Appl. Mech. Eng.* 190 (2001), 3771–3796. zbl doi
- [33] *G. Papamakarios, T. Pavlakou, I. Murray*: Masked autoregressive flow for density estimation. *Advances in Neural Information Processing Systems* 30 (NIPS 2017). MIT Press, Cambridge, 2017, pp. 2338–2347.
- [34] *D. Pavlidis, G. J. Gorman, J. L. M. A. Gomes, C. C. Pain, H. ApSimon*: Synthetic-eddy method for urban atmospheric flow modelling. *Boundary-Layer Meteorology* 136 (2010), 285–299. doi
- [35] *J. Quiñonero-Candela, C. E. Rasmussen*: A unifying view of sparse approximate Gaussian process regression. *J. Mach. Learn. Res.* 6 (2005), 1939–1959. zbl MR
- [36] *N. Ramakrishnan, C. Bailey-Kellogg, S. Tadepalliy, V. N. Pandey*: Gaussian processes for active data mining of spatial aggregates. *Proceedings of the 2005 SIAM International Conference on Data Mining*. SIAM, Philadelphia, 2005, pp. 427–438. doi
- [37] *C. E. Rasmussen*: Gaussian processes in machine learning. *Advanced Lectures on Machine Learning*. Lecture Notes in Computer Science 3176. Springer, Berlin, 2003, pp. 63–71. doi
- [38] *D. J. Rezende, S. Mohamed*: Variational inference with normalizing flows. Available at <https://arxiv.org/abs/1505.05770> (2015), 10 pages.
- [39] *J. Smagorinsky*: General circulation experiments with the primitive equations I. The basic experiment. *Mon. Wea. Rev.* 91 (1963), 99–164. doi

- [40] *J. Song, S. Fan, W. Lin, L. Mottet, H. Woodward, M. Davies Wykes, R. Arcucci, D. Xiao, J.-E. Debay, H. ApSimon, E. Aristodenou, D. Birch, M. Carpentieri, F. Fang, M. Herzog, G. R. Hunt, R. L. Jones, C. Pain, D. Pavlidis, A. G. Robins, C. A. Short, P. F. Linden*: Natural ventilation in cities: The implications of fluid mechanics. *Building Research & Information* 46 (2018), 809–828. 
- [41] *M. K. Titsias*: Variational learning of inducing variables in sparse Gaussian processes. *Proc. Mach. Learn. Res.* 5 (2009), 567–574.
- [42] *M. K. Titsias*: Variational Model Selection for Sparse Gaussian Process Regression. Technical report, University of Manchester, Manchester, 2009.
- [43] *V. H. Tran*: Copula variational Bayes inference via information geometry. Available at <https://arxiv.org/abs/1803.10998> (2018), 23 pages.
- [44] *D. Tran, R. Ranganath, D. M. Blei*: The variational Gaussian process. Available at <https://arxiv.org/abs/1511.06499> (2015), 14 pages.
- [45] *H. Wickham*: *ggplot2: Elegant Graphics for Data Analysis*. Use R! Springer, Cham, 2016.  

Authors' addresses: *Gabor Tajnafoi, Rossella Arcucci* (corresponding author), Data Science Institute, Department of Computing, Imperial College London, South Kensington Campus, London SW7 2AZ, United Kingdom, e-mail: gabor.tajnafoi18@imperial.ac.uk, gtajnafoi@gmail.com, r.arcucci@imperial.ac.uk; *Laetitia Mottet*, Applied Modelling and Computation Group, Department of Earth Science & Engineering, Imperial College London, South Kensington Campus, London SW7 2AZ, United Kingdom, e-mail: l.mottet@imperial.ac.uk; *Carolanne Vouriot*, Department of Civil Engineering, Imperial College London, 58 Princes Gate, Kensington, London SW7 1AL, United Kingdom, e-mail: carolanne.vouriot12@imperial.ac.uk; *Miguel Molina-Solana*, Department of Computer Science and AI, Universidad de Granada, 18071 Granada, Spain, and Data Science Institute, Department of Computing, Imperial College London, South Kensington Campus, London SW7 2AZ, United Kingdom, e-mail: miguelmolina@ugr.es; *Christopher Pain*, Department of Earth Science & Engineering, Imperial College London, South Kensington Campus, London SW7 2AZ, United Kingdom, e-mail: c.pain@imperial.ac.uk; *Yi-Ke Guo*, Data Science Institute, Department of Computing, Imperial College London, South Kensington Campus, London SW7 2AZ, United Kingdom, e-mail: y.guo@imperial.ac.uk.