Informed Sampling and Adaptive Monitoring using Sparse Gaussian Processes

Rajat Mishra*, Mandar Chitre† and Sanjay Swarup‡
*NUS Graduate School of Integrative Sciences and Engineering, National University of Singapore, Singapore
†Department of Electrical and Computer Engineering, Faculty of Engineering, National University of Singapore, Singapore
‡Department of Biological Sciences, Faculty of Science, National University of Singapore, Singapore
e-mail: rajat.mishra@u.nus.edu, {mandar, sanjay}@nus.edu.sg

***Abstract*—Physical sample collection from the hotspots of an environmental field is of interest to scientists, environmental protection agencies as well as public utility authorities. Such a sample collection task requires prior knowledge of the locations of the hotspots, which is generally not available. In this paper, we suggest an algorithm, Sampling and Adaptive Monitoring (SAM), to perform both the tasks of approximating the environmental field and sampling from the hotspots, simultaneously. We test our performance using a user defined utility function and show that we are able to obtain a good approximation of the field and collect samples from the hotspots within the stipulated time.**

I. Introduction and Related Work

Sampling of water hotspots is an important task in many environmental monitoring activities, providing a means for us to understand the various processes taking place in our water resources. Environmental fields such as temperature or chlorophyll content, are scalar quantities which exhibit spatio-temporal variations [1]. Moreover, to achieve the objective of sampling the hotspots, prior knowledge of the field's spatial and temporal variations is necessary. Such prior information regarding the variability is generally not available but by using some recent adaptive frameworks [2, 3, 4, 5, 6], it is possible to obtain a good approximation of the environmental field. We define such frameworks as monitoring or exploration-only frameworks.

In general monitoring frameworks impose a temporal constraint on the field estimation task. These temporal constraints are used to exploit the quasi-static characteristic of the environmental fields. In simple terms, the fields are assumed to be temporally constant for short durations of time, spanning anywhere between minutes to several hours [1, 7]. For the sample collection task, the monitoring frameworks can be used to obtain an approximation of the field, which can be followed by sampling from the hotspots of the approximated field. Such an approach will require the monitoring task to be completed before the sampling of hotspots can begin. This sequential order of tasks will introduce a temporal gap between the field estimation and the physical sample collection and ideally this gap should be as low as possible.

We are interested in combining both the tasks of field estimation and sample collection into a single on-line framework. This can be achieved by collecting representative data such as temperature measurements to get a good approximation of the field, and simultaneously use this approximation to collect physical samples. Similar to monitoring frameworks, the combined approach can make use of the temporal constraints on field estimation and sampling task, which helps in reducing the model complexity of the environmental field. It is important to note that our problem definition falls into the category of exploration-vs-exploitation dilemma. The exploration characteristics come from the objective to get a good approximation of the field, while the exploitation features come from collecting the samples from the estimated hotspots.

Fig. 1. Autonomous Surface Vehicle being deployed for monitoring and water sample collection in a freshwater lake in Singapore.

We briefly review some of the frameworks capable of simultaneous sampling and monitoring. The framework in [8] suggests an on-line multi-choice hiring algorithm. This framework presents an algorithm for making irrevocable sample selection decisions from a set of possible sampling candidates. Following this work, the algorithm in [9] uses a formulation of secretary-hiring problem [10] and probabilistic model for field estimation. In both of these frameworks, the field estimation and sample collection occur in an on-line manner but only for predefined paths, such as lawn mower or yo-yo motion. Such an approach, which predefines the paths, reduces the adaptive nature in the task of obtaining a good approximation of the field, and introduces a bias in the sampling task.

Recently, [11] suggested another framework for simultaneous monitoring and sampling. In this framework, the sampling and monitoring tasks are distributed between the two members of a team, consisting of an explorer robot and a sampler

robot. The sampling is based on a look-back secretary-hiring formulation, which controls the sampler robot's motion. The formulation of this framework is suitable for the scenarios where a robot can either be configured as an explorer or a sampler, within a team consisting of two or more robots. However, our objective is to develop a single robot framework for both sampling and monitoring tasks without the need for prior information about the environmental field. Such a single robot framework can then be easily extended to a multi-robot system, where the sampling and monitoring roles can be shared across members of the multi-robot team.

We recently proposed an online IPP framework, *AdaPP*, for estimating scalar environmental fields [6]. At each step, *AdaPP* evaluates the potential of all the neighbouring locations and selects the one which reduces maximum variance, thus making it an exploration only algorithm. In this paper, we present a single-robot framework, which we call SAM, to address the task of simultaneous sampling and monitoring of an environmental field. This framework aims to collect samples from the hotspots, while using a sparse variant of Gaussian Process (GPs) to estimate the environmental fields. Moreover, the sampling decisions are guided by an Upper Confidence Bound (UCB) formulation and a utility function defined by the user, thus offering scientists the freedom to customize the sampling criteria as per the experiment's requirement. Some background information about GPs and UCB algorithm is presented in Section II. Following this background information, Section III presents our technical approach. In Section IV, we test the performance of our framework using sea surface temperature dataset [12] and present the results. Finally, Section V provides the conclusion of our work.

II. Background Information

A. Sparse Gaussian Process for estimating environmental fields

GP models are generally used for non-parametric regression problems [13], such as spatial data modeling [14]. In a standard GP model with zero mean prior, kernel functions are solely responsible for modeling the relationship between output and input variables. Moreover, the training complexity of a standard model is $\mathcal{O}(N^3)$, where N is the total number of training samples. Such a training complexity poses a problem in learning GP models for large datasets. Unfortunately, most environmental sensors have high frequency rates and the training dataset can easily reach few thousands of data points for an hour-long mission, making standard GP models not ideal for on-line learning.

A good approach to mitigate the problem with large datasets is to use Sparse GPs. Such an approach uses M data points to summarize the information present in N data points, where $M \ll N$. The training complexity for Sparse GPs also reduces to $\mathcal{O}(NM^2)$. Interestingly, the prediction complexity for mean and variance becomes $\mathcal{O}(M)$ and $\mathcal{O}(M^2)$, respectively, which is completely independent of N. Such computation complexities give sparse GPs an edge over full GP models and make on-line learning of a GP model feasible.

B. Upper Confidence Bound Algorithm

A good approach for solving a sequential decision making problem is the Upper Confidence Bound (UCB) algorithm. In a general UCB formulation, the goal is to find the action that provides the maximum reward from the set of possible actions. The UCB algorithm uses the combination of mean and variance of the rewards for each action and simply selects the action which provides maximum possible reward. The definition of the maximum possible reward defines the characteristics of the decision process and has been used for a variety of applications [15, 16].

III. Technical Approach

We are interested in estimating a scalar environmental field and collecting samples from $\mathcal{K}$ locations, where the criteria for sampling is given by:

$$\mathcal{S} = \underset{x_i, x_j \in \mathcal{X}}{\arg\max} \sum_{i=1}^{\mathcal{K}} \Big(f(x_i) \times \prod_{j=1, j\neq i}^{\mathcal{K}} g(x_i, x_j) \Big) \tag{1}$$

such that

$$\mathcal{T}(\mathcal{S} \cup \{x_r\}) = T$$

where S is a set containing $\mathcal{K}$ spatial locations from the survey area given by $\mathcal{X}$, $f(x_i)$ is the value of the environmental parameter at a spatial location i, $g(x_i, x_j)$ provides a discounting at spatial location i given a sampled location j, x_i and x_j are the cartesian coordinates representing two spatial locations i and j respectively, T is the total mission time, x_r is the robot's current position and $\mathcal{T}(\mathcal{S} \cup \{x_r\})$ provides the time taken for a path containing $\mathcal{S} \cup \{x_r\}$ locations. In short, (1) represents the selection behavior for a set of spatial locations, which is controlled by the definition of the utility function $g(\cdot, \cdot)$ and field's distribution $f(\cdot)$.

Broadly, we have two optimization problems to solve, one is the estimation of the scalar environmental field and the other is the selection of samples satisfying (1). To have a standard terminology, we consider the task of field estimation as a monitoring or exploration problem and the task of physical sample collection as a sampling or exploitation problem. As discussed before, solving these two problems simultaneously represents an exploration-vs-exploitation dilemma. The algorithm in [6] explains our exploration-only framework, which aims at reducing the overall variance of the estimated field, and shows that we are able to get a good approximation of the ground truth. We use this idea of reducing the variance to address our exploration problem. At the same time we also leverage some of the algorithm's features such as spatial decomposition and sparse GP [17] for estimating fields. However, to address the exploration-vs-exploitation dilemma we have added certain elements to our framework and these are discussed here in detail.

A. Using Kernel Information in Planning

At each planning iteration, our exploration-only framework simulates future paths through all the neighboring cells and

selects the path that reduces the maximum variance in the model. This maximum variance for each future path is calculated by adding up the variances of all the cells to be visited while traversing any future path. These calculations are simple and fast, however, it decouples the planning from the available modeling information, such as the kernel fuction. More details about the framework are presented in [6]

An interesting characteristic of using Sparse GPs, or even standard GPs, is the predicted variance, which can be estimated using just the kernel model. In our field estimation task, this means that the change in model variance for any spatial location can be approximated without physically traversing to it. Moreover, the prediction complexity for estimating the variance is completely independent of N, which means the time required to estimate the variance is fixed, given by $\mathcal{O}(M^2)$, and does not suffer from increasing number of training points. The combination of these two characteristics, where kernel model is used to estimate change in variance, and where the prediction time is fixed, makes sparse GPs an ideal tool for quickly approximating the change in variance for an unvisited spatial location.

The exploration framework in *SAM* makes use of the kernel properties of Sparse GP for evaluating different paths through the neighboring cells. The exploration framework simulates candidate paths through each neighboring cell but path selection is based on the resulting variance, which is different from the criteria used in *AdaPP*. We estimate the resulting variance for a candidate path using the kernel model and the set of spatial locations in the future path, without actually traversing it.

For an exploration-only framework, the estimated resulting variance can be used to evaluate different candidate paths and find a path P which satisfies the following:

$$P = \arg\min_{P_i \in \mathcal{P}} \int_{\mathcal{X}} V(P_i) \tag{2}$$

where $\mathcal{P}$ is a set of all the candidate paths, $V(P_i)$ is the resulting variance after taking the path P_i and $\mathcal{X}$ is the entire survey area. However, our objective is to also collect samples from the hotspots of the estimated field and therefore, merely satisfying (2) will not solve our overall problem. We discuss this problem in Section III.C in detail. It is important to note that the path selection is not directly based on the individual rewards of the path's waypoints but on the resulting variance, thus making it more directed towards the overall performance of the path.

B. Selecting Sampling Locations from a Path

In an ideal scenario, we need to select the best sampling candidates that satisfy (1) using complete information about of $f(\cdot)$. However, we do not assume any prior information and the exploration happens while selecting the samples, therefore knowing the actual value of $f(\cdot)$ in advance is not possible. However, an estimate of $f(\cdot)$ can be obtained using the sparse GP's mean ($\mu(\cdot)$) and variance ($\sigma^2(\cdot)$), which are continuously updated using the data collected during exploration. In addition to this data collection, the exploration framework provides future candidate paths and the resulting variance as discussed in Section III.A.

Algorithm 1: s-*Sampler* UCB - Selecting sample locations from a path

Data: Sampling Canditates $(\mu(\mathcal{P}_i) + \sigma^2(\mathcal{P}_i))$, Remaining no. of samples (s), Previously collected samples $(\mathcal{S}_t)$, Utility function $g(\cdot,\cdot)$, Path through the ith neighbouring location $\mathcal{P}_i$

Result: Greedy subset for remaining locations $(\mathcal{S}_{T-t})$, Expected Reward (κ)

`/* Initialization */`

Initial bound $C = \mu(\mathcal{P}_i) + \sigma^2(\mathcal{P}_i)$;

for *j = 1:len($\mathcal{S}_t$)* **do**

  $C = C * g(\mathcal{P}_i, x_j)$;

$\mathcal{S}_{T-t} = \emptyset$;

`/* Algorithm Loop */`

while $s > 0$ **do**

  Find the maximum reward $x_c = \max(C)$;

  Add x_c to $\mathcal{S}_{T-t}$;

  Update no. of samples $s = s - 1$;

  Update the rewards $C = C * g(\mathcal{P}_i, x_c)$;

  Update expected reward $\kappa = \kappa + \mu(x_c) + \sigma^2(x_c)$

Given the estimated field and the candidate paths, the sampling problem can be divided into two sub-problems. First problem is to select the best $\mathcal{K}$ samples for each future path and the second problem is to select an option given all the $\mathcal{K}$ samples (exploitation) and the resulting variances (exploration) for each candidate path. The first problem follows the general formulation of a sequential decision making problem, whereas, the second problem represents an exploration-vs-exploitation dilemma. In this sub-section, we discuss the solution for the first problem and in Section III.C we discuss the solution for the second problem.

Given a future path P_i and the sparse GP's mean $\mu(\cdot)$ and variance $\sigma^2(\cdot)$, we define the spatial locations in the path P_i as the candidates for taking a sample and we define the reward for taking a sample as the sum of mean and variance of the selected location, $\mu(\cdot)+\sigma^2(\cdot)$. With this formulation, a simple approach to find the $\mathcal{K}$ sampling locations is to sort $\mu(P_i) + \sigma^2(P_i)$ in descending order and select the top $\mathcal{K}$ candidates. Such a selection process can produce all samples from the same hotspots, which is a valid solution but may be irrelevant to the scientific experimental objective. Therefore, we use a utility function $g(\cdot)$ to introduce the experimental objectives into our sample selection problem and adjust the sampling rewards on-the-fly. This utility function is used to control the sampling behaviour, such as minimum distance between all samples or adaptively control a minimum threshold on the sampling rewards.

Our approach to select the best sampling point from a given path is presented in Algorithm 1. At each sample selection step, the utility function is used to weigh the rewards for the

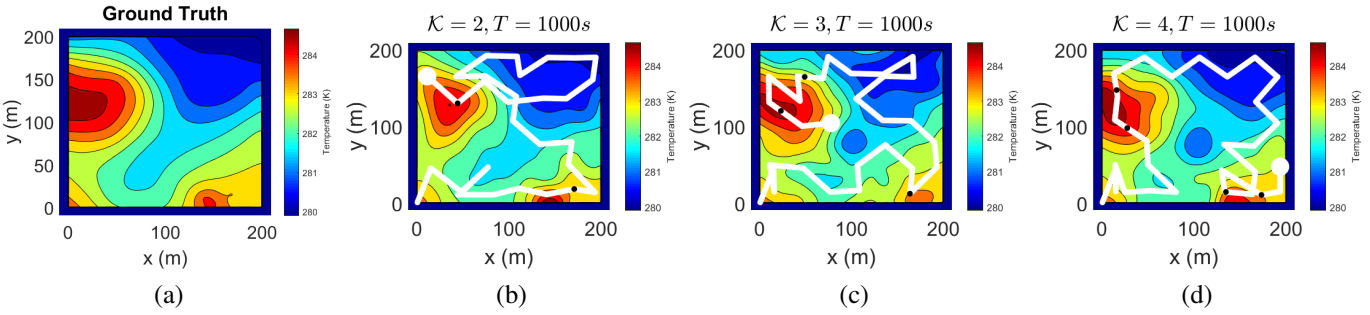

Fig. 2. Simulation results for a 200 m $\times$ 200 m sea surface temperature field. (a) is the ground truth obtained using the MUR dataset [12]. (b)-(d) represents the field predicted using *SAM* for T = 1000 s and different $\mathcal{K} = \{2, 3, 4\}$ respectively. The white lines represent the paths taken by the robot and the black dots represent the locations where the samples were taken.

Algorithm 2: SAM - Sampling and Adaptive Monitoring Framework

Data: Path Canditates $\mathcal{P}$, Total Rewards for each $R(\mathcal{P})$, Resulting Variance for each path $V(\mathcal{P})$, Threshold ϵ

Result: Selected path P_*

/* Initialization */

Find the maximum reward $R_m = \max R(\mathcal{P})$;

Generate the subset $\mathcal{P}' = \arg_{P_i \forall \mathcal{P}} R(P_i) > R_m - \epsilon$;

if $|\mathcal{P}'| == 1$ **then**

  Select the path with maximum reward $P_* = P_m$;

else

  Select the path with minimum variance $P_* = \arg\min_{P_j \forall \mathcal{P}'} V(P_j)$;

remaining sampling candidates. This is followed by selecting the next sampling location with the maximum weighted reward and this process continues until all the $\mathcal{K}$ sampling locations are selected for the given path. At the end, the total reward for a path is given by the sum of all weighted rewards for all the selected sampling locations.

C. Unifying Monitoring and Sampling tasks

In Section III.A and III.B, we have described the approaches for generating candidate paths for exploration and selecting best locations to sample from a given path. However, unifying these two components into one framework is not straight forward. The approach that we have used is given in Algorithm 2.

At each planning iteration, a set of candidate paths are available with information about the resulting variance and the total reward for collecting samples. This is where the exploration-vs-exploitation dilemma arises and to overcome it, we use the statistics of the resulting variance and total reward of each candidate path. We sort the candidate paths according to the total reward of the candidate paths and form a smaller set of paths $\mathcal{P}'$, whose total reward lies within a threshold ϵ of the maximum possible reward. The idea of forming a smaller subset is similar to the Branch-and-Bound method used for efficiently narrowing down a search problem. If the set $\mathcal{P}'$ has a single path, which is the candidate with the maximum reward, we straight away select it and take the first step according to this path. On the other hand, if the set has more than one path, we select the path with the minimum resulting variance. Such an approach of forming a subset based on the rewards and selecting paths using the resulting variance gives the balance required for the exploration-vs-exploitation dilemma.

IV. Simulation Results

For our simulations, we used the real field Sea Surface Temperature (SST) data provided by the Jet Propulsion Laboratory [12]. We extracted the temperature data for an area of 200 km $\times$ 200 km each and mapped it to an area of 200 m $\times$ 200 m. This mapping was done to conserve the features of the environmental field and have an area that can be explored within a practical value of T.

Moreover, we used the following utility function for our simulations:

$$g(x_i, x_j) = \min(\frac{||x_i - x_j||}{\sigma_l}, 1) \tag{3}$$

where σ_l is 50 m. The mission time is set to 1000 s and the vehicle is assumed to be traveling at constant speed of 1 m/s. We implemented our algorithm using MATLAB.

Fig. 2 shows the estimated fields and sampled locations using our algorithm for $\mathcal{K} = \{2, 3, 4\}$ sampling locations, with the value for ϵ set to 0.25. It can be observed that our algorithm was able to capture samples from the hotspots and provide a good approximation of the environmental field in each case. Moreover, we used the root mean square error (RMSE) to measure the performance of our exploration framework and these results are presented in Fig. 3a. This figure shows similar transient drops in RMSE for different values of $\mathcal{K}$, which demonstrates that our exploration strategy works efficiently.

We further examined the exploration characteristics over 10 runs for each sampling scenario and these results are presented as box plots in Fig. 3b. This figure shows the RMSE values at end of T seconds for 10 different simulations and for different values of $\mathcal{K}$. Although, the box plots are not completely identical, the mean RMSE for each value of $\mathcal{K}$ is similar. This provides more evidence to our observations from Fig. 3a that

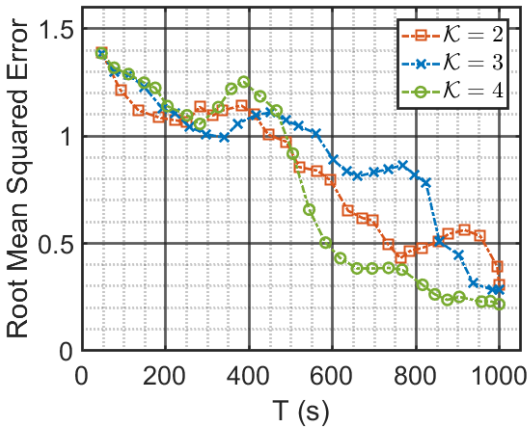

(a)

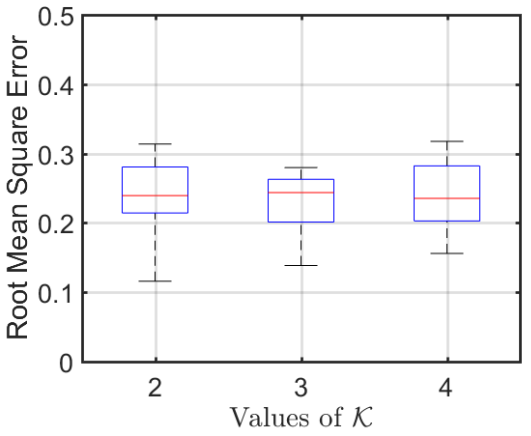

(b)

Fig. 3. Root Mean Squared Error in predicting the ground truth (a) with respect to time and (b) with respect to different $\mathcal{K} = \{2, 3, 4\}$ for 10 simulations each.

our framework provides good estimates of the field, even on increasing number of samples.

V. Conclusion

In this work, we presented a framework for simultaneously sampling and estimating a scalar environmental field using a single robot with the constraints on mission time. The environmental field was modeled with a sparse GP framework, SPGP, and the sampling decision process was handled using a combination of UCB and Branch-and-Bound approach. The simulations were conducted using real world temperature data and the results show that our algorithm is able to provide a good approximation of the environmental field, while simultaneously collecting samples from the hotspots of the estimated field.

Acknowledgment

The authors would like to thank Ahmed Mahmood, Hari Vishnu and Bharath Kalyan for the productive discussions on GP models and exploration-vs-exploitation dilemmas. The research work of Rajat Mishra was supported by the Singapore Center for Environmental Life Sciences Engineering (SCELSE), Singapore.

References

[1] N. D. Le and J. V. Zidek, *Statistical analysis of environmental space-time processes*. Springer Science & Business Media, 2006.

[2] B. Zhang and G. S. Sukhatme, "Adaptive sampling for estimating a scalar field using a robotic boat and a sensor network," in *Robotics and Automation, 2007 IEEE International Conference on*, IEEE, 2007, pp. 3673–3680.

[3] J. Yu, M. Schwager, and D. Rus, "Correlated orienteering problem and its application to persistent monitoring tasks," *IEEE Transactions on Robotics*, vol. 32, no. 5, pp. 1106–1118, 2016.

[4] N. Cao, K. H. Low, and J. M. Dolan, "Multi-robot informative path planning for active sensing of environmental phenomena: A tale of two algorithms," in *AAMAS 2013*, 2013, pp. 7–14.

[5] K. H. Low, J. Chen, J. M. Dolan, S. Chien, and D. R. Thompson, "Decentralized active robotic exploration and mapping for probabilistic field classification in environmental sensing," in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, International Foundation for Autonomous Agents and Multiagent Systems, 2012, pp. 105–112.

[6] R. Mishra, M. Chitre, and S. Swarup, "Online informative path planning using sparse gaussian processes," in *OCEANS 2018 MTS/IEEE Conference and Exhibition*, IEEE, 2018.

[7] D. Higdon, "A process-convolution approach to modelling temperatures in the north atlantic ocean," *Environmental and Ecological Statistics*, vol. 5, no. 2, pp. 173–190, 1998.

[8] Y. Girdhar and G. Dudek, "Optimal online data sampling or how to hire the best secretaries," in *2009 Canadian Conference on Computer and Robot Vision*, IEEE, 2009, pp. 292–298.

[9] J. Das, F. Py, J. B. Harvey, J. P. Ryan, A. Gellene, R. Graham, D. A. Caron, K. Rajan, and G. S. Sukhatme, "Data-driven robotic sampling for marine ecosystem monitoring," *The International Journal of Robotics Research*, vol. 34, no. 12, pp. 1435–1452, 2015.

[10] M. Bateni, M. Hajiaghayi, and M. Zadimoghaddam, "Submodular secretary problem and extensions," in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, Springer, 2010, pp. 39–52.

[11] S. Manjanna, A. Q. Li, R. N. Smith, I. Rekleitis, and G. Dudek, "Adaptive exploration and sampling by heterogeneous robotic team," in *Robotics and Automation (ICRA), 2018 IEEE International Conference on*, IEEE, 2018.

[12] N/A, "Jpl mur measures project, ghrsst level 4 mur global foundation sea surface temperature analysis," vol. 2, 2010. DOI: http://dx.doi.org/10.5067/GHGMR-4FJ01.

[13] C. E. Rasmussen, "Gaussian processes in machine learning," in *Advanced lectures on machine learning*, Springer, 2004, pp. 63–71.

[14] M. L. Stein, *Interpolation of spatial data: some theory for kriging*. Springer Science & Business Media, 2012.

[15] N. Srinivas, A. Krause, S. Kakade, and M. Seeger, "Gaussian process optimization in the bandit setting: No regret and experimental design," in *Proceedings of the 27th International Conference on International Conference on Machine Learning*, Omnipress, 2010, pp. 1015–1022.

[16] A. Garivier and E. Moulines, "On upper-confidence bound policies for switching bandit problems," in *International Conference on Algorithmic Learning Theory*, Springer, 2011, pp. 174–188.

[17] E. Snelson and Z. Ghahramani, "Sparse gaussian processes using pseudo-inputs," in *Advances in Neural Information Processing Systems*, 2006, pp. 1257–1264.