Analyzing the balance of Exploration and Exploitation in Adaptive Information Sampling of Unknown Spatial Fields

Aiman Munir Ramviyas Parasuraman

***Abstract*— Adaptive sampling and informative path planning approaches enable efficient selection of the mobile robot's waypoints to obtain accurate sensing and mapping of a physical process, such as the radiation or field intensity. However, it is not clear to what extent the parameters of the adaptive sampling function play a role in balancing the exploration of new information and exploitation of existing information combined with the robot's energy consumption perspective. This paper provides this necessary perspective and uniquely analyzes the impact of the adaptive sampling algorithm's information function used in exploration and exploitation to achieve a trade-off between balancing the mapping, localization, and energy efficiency objectives. In addition, we propose a new time-varying parameter to balance the sampling objectives dynamically during a mission. We use Gaussian process regression (GPR) to predict and estimate confidence bounds, thereby determining each point's informativeness. Through extensive experimental data, we provide a deeper and holistic perspective on the effect of information function parameters on the prediction map's accuracy (RMSE), confidence bound (variance), energy consumption (distance), and time spent (sample count). The results provide meaningful insights into choosing appropriate information function parameters based on sampling objectives (e.g., source localization or mapping).**

I. Introduction

Mobile robot-aided mapping of unknown spatial fields (environmental processes), such as information sampling [1], [2], informative path planning [3], [4], sensor coverage [5], localization of signal sources [6], and monitoring of environmental phenomena [7], [8], have been well investigated. Particularly, sensor coverage with multiple robots involves positioning robots in an optimal manner to maximize overall performance in terms of sensing the environmental phenomenon. Monitoring is a persistent process for identifying anomalies in physical processes by efficiently collecting the most informative samples. For all of these objectives, modeling of the underlying processes is required or involved, which can be obtained via data-driven sampling and predictive mapping (extrapolation) of the sensed quantity.

Further, modeling physical processes plays an important role in robot decision-making. Robots need to create a dynamic predictive model of the environmental phenomenon to accomplish mapping tasks, especially when the environment is underexplored or unexplored. Mapping spatial distribution enables the robots to work autonomously in search and rescue missions and make decisions without human intervention (e.g., for rescuing targets in areas with high radiation exposure). Similarly, the robot requires knowing the areas with higher risks so that the robot can always choose a path to connect to the network and maintain communication, etc. [9]. Therefore, robotics researchers are actively investigating different strategies for mapping physical processes, such as radiation, Wi-Fi signal strength, gas distribution, radio signal strength, etc., using unmanned vehicles [6], [10], [11].

In order to map the environmental phenomenon, it is crucial to measure the value of a physical process throughout the environment. However, not all the locations can provide useful information about the change in the process itself. Primary considerations involved in mapping such processes include the degree of autonomy, accuracy, and efficiency. Measuring intensity at every location is deemed impractical; hence, dense sampling is not viable in mapping. Instead, an accurate, time- and cost-effective process model can be obtained by gathering samples from the points that contain the most significant amount of information.

Exploration refers to accumulating samples from previously unexplored areas to reduce uncertainty in the predicted/extrapolated map, while **exploitation** implies determining the next sampling point based on the best information from the current estimates (to localize the sources, for example). Mapping algorithms and techniques in the literature either use exploration or exploitation, or a combination of both. For example, an active control law for a mobile robot is proposed in [6] to shift between exploration and exploitation objectives. Similarly, [5] uses a utility function for adjusting exploration and exploitation. On the other hand, the parallel strip route (pure exploratory approach) is used in [12], [13] to explore the environment and map the spatial distribution. These two techniques are fundamental to adaptive sampling.

The purpose of the study in this paper is to compare how well various exploration and exploitation techniques perform, as well as to offer an exclusive analysis by looking at the impact on all the key metrics, including completion time, mapping accuracy, variance, and robot energy consumption. Even though the exploitation coefficients' impact on completion time and mapping accuracy is described in a recent study [14], the effect of energy consumption and confidence bounds are not discussed. We have to consider practical limitations, such as that robots have limited energy (battery) and we do not always have ground-truth data to compare and estimate real-time mapping performance in realistic settings. In this case, the model prediction variance can inform us how confident we can be in our predictions. Thus, these metrics are critical to the mapping process.

This paper makes the following **contributions**:

- We comprehensively compare the effect of the trade-

School of Computing, University of Georgia, Athens GA 30602, USA, email: {aiman.munir, ramviyas}@uga.edu

off between exploration and exploitation objectives for different information function coefficients in adaptive sensor sampling. We measure the performance metrics related to mapping accuracy, prediction confidence (variance), exploration time, and energy (distance).
- We systematically analyze the impact of different source locations on this trade-off.
- Based on the above analysis, we propose an *TimeVariant* algorithm that adapts with respect to the time evolution of the mission and demonstrates the effect of balancing the parameters on the performance metrics.

The results from this study provide meaningful insights into choosing appropriate sampling functions and parameters for the map prediction, ultimately balancing the exploration objectives (minimizing uncertainty) and exploitation objectives (maximizing information value) combined with energy efficiency through adaptive sampling with mobile robots. A comprehensive page containing significant additional data to support our work is available[1].

II. Related Work

Both adaptive (taking informativeness of the sampled data) and non-adaptive (without informativeness) sampling approaches have been previously reported in the literature. Non-adaptive sampling methods focus on sampling the whole environment [13], [15], [16]. Non-adaptive methods are time-consuming, and in such methods, it is hard to achieve the desired threshold of information certainty. Alternatively, adaptive sampling methods provide convergence to an objective (threshold), and sampling can be stopped as soon as the desired threshold is reached. Several objective functions have been used in coordination with Gaussian Process Regression [17], [18] to map the physical process (i.e., to predict the samples at unexplored (unvisited) locations with confidence bounds). A comprehensive comparison of these works is available as a tabular form in the resource webpage[1].

The approaches to planning a robot path that contains the most-informative samples have been well studied. In [1], a technique that maximizes the mean entropy information metric is used to search a station. Then using an RRT-based Informative Planner algorithm, selects the path that provides maximum utility, i.e., trades off the informativeness and cost to reach that station. Authors in [7] use entropy as an information criterion over Sparse Gaussian Process to find the most informative locations to monitor salinity in the ocean persistently. Similarly, [19] uses wireless signals for robot localization in an indoor GPS-denied environment. It learns a path loss model from the data and then trains the Gaussian process with the mismatches between the models and the data with a focus on better model variance prediction. In [10], the authors focused on mapping in structured environments. The algorithm partitions the environment for each robot and uses differential entropy as an information theory metric on top of Gaussian Process predictions to determine the next sampling point. The authors in [20] proposed a Hexagonal Tree (HexTree) based sampling algorithm, which takes samples over a set of hexagonal grid points and builds a tree of possible trajectories by extending candidate trajectories toward the sampled points.

A related study in [14] presented a decentralized informative path-planning algorithm for a swarm of robots to balance exploration and exploitation. Additionally, it compares the exploitation coefficients in terms of completion time and mapping accuracy. But, the robot's starting position influences the efficiency of the algorithm, which is not studied. Our study differs from [14] in that we holistically consider the energy consumption and confidence bound into account for the trade-off analysis of different sampling objectives. We also report the effect of the information sampling parameter (coefficient) at different time instances during the exploration task. Compared to the literature, our work is novel and unique in the following ways: 1) we extensively analyze several variants of adaptive sampling parameters through simulation experiments with multiple source locations and distributions; 2) we present the balance between sampling and energy demands considering practical resource constraints; 3) we propose a new time-varying adaptive sampling coefficient and holistically analyze several variants.

III. GPR-Aided Robot Sampling

Gaussian Process Regression (GPR) has been widely utilized for modeling spatial processes. For example, [21] used GPR to model spatial functions for mobile wireless sensor networks and to generate a likelihood model for signal strength measurements. To obtain a spatial map in an environment with limited communication, [22] used GPR to model algae bloom occurrence. In [6], GPR is used to obtain a model of radio signal strength, which in turn is used to get the maximum likelihood of the source location. Therefore, we use GPR in our work to dynamically predict and continuously update the sensor data estimates at the whole map region (including unexplored locations) using the available sampled data so far from previously-visited locations by the mobile robot.

A Gaussian Process (GP) is a non-parametric continuous function that defines a probability distribution over functions. It assumes that every point has a normal distribution and there is a correlation between values at these points. Let $q \in Q \subset R^2$ be the location from where the signal strength is measured, and z be the measurement. The value of z at any location q can be related to a function $f(q)$ using the Gaussian noise model for the observed location q as:

$$z = f(q) + \varepsilon, \tag{1}$$

where ε is additive Gaussian noise. We are interested in calculating a posterior function $f_\star$ that makes predictions for given test locations $q_\star \in Q_\star \subset R^2$. A GPR model also known as Kriging assumes a GP prior that can be completely defined by mean and covariance. The joint Gaussian distribution on the test set $Q_* = Q \setminus D$ (D is the dataset containing

[1] https://hero.uga.edu/research/adaptivesampling/

sampled locations), assuming noisy observation z is

$$\begin{bmatrix} z \\ f_* \end{bmatrix} \sim \mathcal{N}\left(\mu(q), \begin{bmatrix} K+\sigma_n^2 I & k_* \\ k_*^T & k_{**} \end{bmatrix}\right), \quad (2)$$

where K is the covariance matrix between the training points, k_* is the covariance matrix between the training points and test points, and k_{**} is the covariance between only the test points. The posterior mean and variance for any testing location $q_\star$ learned by GPR are as follows:

$$\mu[f_\star] = m(q) + k_\star^T \left(K + \sigma_n^2 I\right)^{-1} (y - m(q)) \quad (3)$$

$$\sigma^2[f_\star] = k_{\star\star} - k_\star^T \left(K + \sigma_n^2 I\right)^{-1} k_\star \quad (4)$$

In our experiments, we employ a widely used kernel, i.e., the squared exponential [23].

$$k(q, q') = \sigma_f^2 \exp\left(-\frac{1}{2l^2}|q - q'|^2\right) \quad (5)$$

where both q and q' are training points, σ_f are l are hyperparameters and are called variance and length, respectively. We continuously learn these hyperparameters by maximizing the log marginal likelihood of the observations [24], [6]. The variance and mean in Eq. (3) and (4) are used to calculate the informativeness of every point.

IV. Information Sampling

A method for determining the next sampling point using a utility function during the model creation process is termed adaptive information sampling [3], [10]. Depending on the criteria for selecting the next sampling location, adaptive sampling can be exploration-based, exploitation-based, or a mixture of both types. In this paper, we have used two non-adaptive sampling approaches that do not use an information function as baselines: a predefined S-shaped sweep trajectory (ordered zigzag pattern) (FS); and a random walk (RW). We have combined a random walk baseline with different variants of the adaptive information sampling to get an idea of the environmental process around the robot and then choose the next waypoint based on the information. Here, the first few samples are done using RW, and then the adaptive sample follows. We used Gaussian Process Upper Confidence Bound (GP-UCB) [18] Information (utility) function to calculate the informativeness at every point q:

$$I(q) = \alpha\mu_q + \beta\sigma_q^2. \quad (6)$$

Here, $I(q)$, μ_q and σ_q^2 denote the informativeness, mean and variance of the point q. While α and β represent importance factors for mean and variance, respectively, and determine the weights given to mean and variance to calculate the informativeness of point q. The adaptive sampling strategy is then to optimize the path (next waypoint) for the robot,

$$x^*(t+1) = \arg\max_{q \in Q_*(t)} I(q). \quad (7)$$

In [18], the bounds β are presented in the context of the probability of achieving no regrets in the information gain. It is trivial to observe that high α can bias towards exploitation and high β towards exploration. However, we study how these bounds are connected to the robot's sampling objectives considering its limited resources (energy, time, etc.). Therefore, we set a constraint that $\alpha + \beta = 1$ so as to use a singular parameter in analyzing the influence of the adaptive sampling method on several metrics. Based on different weighted combinations of mean and variance in the informative function, we used the following approaches:

1) MaxMean ($\alpha = 1$) - MaxMean approach chooses the point with maximum intensity value, i.e., max mean location, as the next sampling point.
2) Alpha0.75/Alpha0.5/Alpha0.25 - These approaches selects the location with the highest information value given $\alpha = 0.75$ or $\alpha = 0.5$ or $\alpha = 0.25$, respectively.
3) MaxVar ($\alpha = 0$) - MaxVar chooses the location with the lowest confidence value as the target location.

To support our analysis, we present a new **TimeVariant** algorithm (Eq. (8)) that dynamically varies the α value with respect to the time evolution of the mission, gradually moving from exploration to exploitation. This is expected to provide a balanced performance on multiple objectives: mapping accuracy, source localization, and energy consumed.

$$\alpha(t) = \begin{cases} 0 \text{ (RW)} & \text{until } t_{min} \text{ to collect min. \# of samples} \\ \frac{t - t_{min}}{t_{max}} & \text{at any time t} \geq t_{min} \\ 1 & \text{at } t_{max} - \text{end of mission duration} \end{cases} \quad (8)$$

The idea for this time-varying α is to generalize the dependence of α w.r.t. robot limitations such as energy, communication, task requirements, etc. For instance, energy-limited robots like UAVs can choose α based on their current energy level. Higher energy means the robot can explore better reaching farther regions, while lower energy can let it exploit the signal variations to find the peaks (sources).

Table I enlists the values of alpha and beta for different variants used in this study. As per the literature, most utility functions in informative sampling rely on the mean, variance, or some combination of the two. Because of diversity, we have not used the same weights as in the literature but instead used a range of weights that can give a general idea of how an increase in weights can affect mapping performance. Nevertheless, this analysis can also be used to shed light on a specific scenario of exploration and exploitation, as well as generalizable to multi-robot settings.

TABLE I
VALUES OF *alpha* AND β FOR DIFFERENT $I(q)$.

Approach	α	β
MaxMean	1	0
Alpha0.75	0.75	0.25
Alpha0.5	0.5	0.5
Alpha0.25	0.25	0.75
MaxVar	0	1
TimeVariant	$\propto$ t	1 - α

V. Experiment Design and Implementation

We have developed the simulations using the Robot Operating Systems (ROS [25]) Gazebo simulation framework, built on top of the open-source code base from [26]. We considered a 10 m $\times$ 15 m simulated area free of obstacles (to avoid bias in the analysis due to collision avoidance algorithms). Until the robot's battery is depleted, it takes

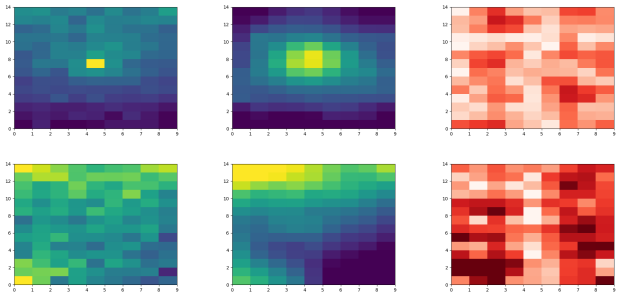

Fig. 1. An example of the Wi-Fi signal distribution with (top) one source at location (4,7) and (bottom) three sources at locations (0,0), (9,14), and (0,14). It shows the ground truth (left), predicted GPR mean (center), and the GPR variance (right) from a simple random walk exploration strategy.

the Wi-Fi device's Received Signal Strength (RSS) samples based on the informative function, navigates to the location, and collects more samples. With each new sample, the GPR is trained, the intensity values of the whole environment (Wi-Fi map) are predicted, and the informativeness of each location is updated based on one of the variants, which includes six adaptive sampling variants in Table. I and one of the two baseline non-adaptive sampling variants based on the scenario. The robot's starting location and battery decay timing were kept fixed for all scenarios.

The ground truth for the Wi-Fi signal map (see Fig. 1 for examples) was generated as per the relation $RSS = RSS_{d0} - 10\eta \ln(d) + \chi_g$, where $RSS_{d0} = TX_{power} - 20\ln(\frac{3}{4\pi * f})$ is the signal reference power at $d0 = 1m$, f is the signal frequency (2.4GHz), η is the path loss factor ($\eta = 3$ in our experiments), d is the distance between the signal source and robot, χ_g is a Gaussian distribution with zero mean and a variance ($0.65dBm^2$) to represent noise in signals, similar to the settings in [26].

A. Experiment Setup

For our experiments, we deployed a hector UAV robot (an aerial robot) with a battery capacity to sustain 500 ROS seconds (with a real-time factor close to 1) and a starting position of (4.5,0). For simplicity, we only restrict the exploration map to a 2D planar region, but the analysis can be extended to 3D domains as the sampling function in Eq. (7) is generalizable. The UAV randomly first explores the surrounding region with an RW approach by randomly moving up, down, left, or right by 3 points from its current position. This is to accumulate a minimum number of samples (15 in our case, corresponding to $\approx$ 75 ROS seconds) to start the GPR process. After t_min, the next sampling location is chosen based on one of the information function variants. We simulate two distributions: 1) **single-modal** source experiments, where there is one signal source 2) **multi-modal** source experiments, where there are three signal sources, as depicted in Fig. 1. We ran five trials per variant in each scenario. Further, the experiments were repeated for five different Wi-Fi source locations.

TABLE II

MAPPING PERFORMANCE (MEAN AND STD) OF THE DIFFERENT SAMPLING FUNCTIONS IN OUR STUDY.

Single modal source experiments				
Variant	# of Samples	RMSE (dBm)	Variance (dBm^2)	Cumulative Distance (m)
Alpha0.75	205 ± 14.66	4.33 ± 0.35	26.27 ± 14.27	175.2 ± 75.87
Alpha0.5	194 ± 12.61	4.23 ± 0.1	16.13 ± 10.27	284.8 ± 91.59
Alpha0.25	168 ± 7.68	4.11 ± 0.07	5.73 ± 0.57	494.11 ± 70.57
MaxVar	145 ± 3.46	4.18 ± 0.1	2.49 ± 0.29	696.28 ± 16.86
MaxMean	226 ± 1	6.06 ± 1.1	37.84 ± 44.57	61.44 ± 4.94
TimeVariant	189 ± 5.29	4.25 ± 0.17	12.85 ± 4.99	323.52 ± 34.59
RW	150 ± 1	4.45 ± 0.22	3.1 ± 0.94	520.94 ± 6.34
FS	36 ± 1	4.82 ± 0.66	7.12 ± 2.13	68.32 ± 0.64
Multi-modal source experiments				
Variant	# of Samples	RMSE (dBm)	Variance (dBm^2)	Cumulative Distance (m)
Alpha0.75	207 ± 9.64	7.52 ± 0.7	16.57 ± 0.7	143.29 ± 27.92
Alpha0.5	175 ± 19.75	6.97 ± 0.16	11.36 ± 1.16	382.4 ± 78.27
Alpha0.25	144 ± 5	6.79 ± 0.09	6.20 ± 0.41	692.47 ± 28.28
MaxVar	139 ± 10.86	6.92 ± 0.16	5.3 ± 0.7	695.6 ± 47.25
MaxMean	227 ± 0	9.11 ± 0.81	22.93 ± 5.31	59.44 ± 1.66
TimeVariant	176 ± 5.29	7.23 ± 0.91	8.29 ± 0.91	422.36 ± 43.7
RW	150 ± 1	7.13 ± 0.33	5.62 ± 0.33	521.03 ± 5.91
FS	36 ± 0	8.63 ± 0.17	47.74 ± 0.17	68.12 ± 0.25

B. Performance Metrics

Samples: The number of RSS samples taken by the robot using its Wi-Fi device.

RMSE: The root mean squared error between the predicted signal (mean) information (Wi-Fi signal strength) through the GPR and the ground truth information. The aim is to get predictions as close as possible to the ground truth, i.e., lower RMSE. The RMSE values shown in the tables and figures represent the average RMSE over the whole map.

Variance: The confidence bounds of the predicted values given by the GPR. The goal is to be confident about the predicted mean value, i.e., lower variance. The variance values shown in the tables and figures represent the average variance over the entire map.

Cumulative Distance: This refers to the total distance traveled by the robot(s). The shorter the distance traveled, the lesser power consumption. We used the cumulative distance metric as an indicative of the robot's energy cost.

Source localization accuracy: If the location at where the maximum mean value of the predicted GP map lies within $1m$ of the actual source location, then that is classified as the correct localization, else incorrect localization. The localization accuracy is then the percentage of correct localization of all trials of all source location experiments combined.

VI. RESULTS AND ANALYSIS

Table II provides a summary of the performance metrics results obtained by averaging the data collected over all trials with different source locations. Fig. 2 presents a summarized view of the results from the adaptive sampling coefficient perspective. Figs. 3 and 4 provide a comprehensive review of the single source and multi-modal experimental results, respectively. It is generally understood that the exploration objective seeks to minimize the variance (uncertainty) of the predicted information, while exploitation seeks to minimize the RMSE of the predicted map (information accuracy). In both cases, we need to make use of the predictions as soon as possible. For instance, in the case of exploitation, we need to identify the signal source location, i.e., the place with the maximum signal intensity. An informative function can be

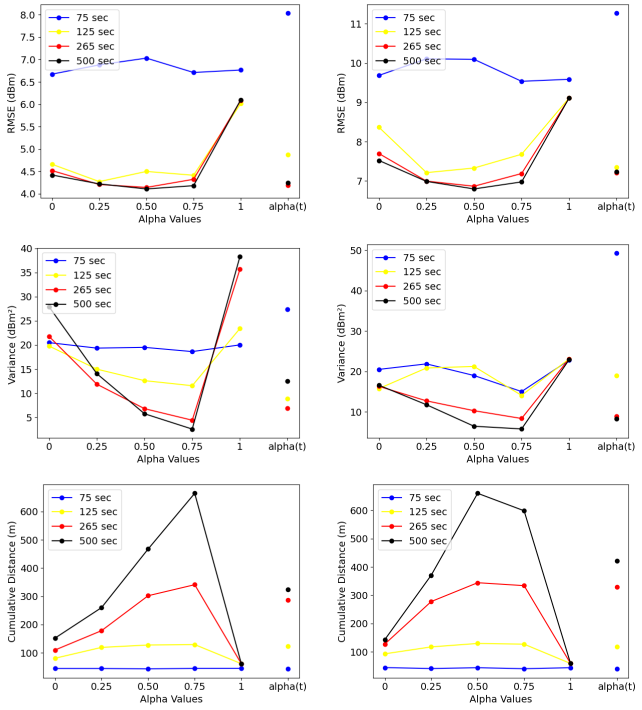

Fig. 2. Average RMSE, Variance, and Cumulative Distance metrics of all the variants analyzed. (Left) Single source experiments and (Right) Multi-modal distribution experiments. Here, alpha(t) is the TimeVariant algorithm that gradually increases α with respect to time.

either exploitation-based, exploration-based, or a weighted combination of both. In our variants, the MaxMean strategy does pure exploitation, whereas MaxVar, Fixed Sweep, and Random Walk are pure exploration strategies. The rest of the variants combine exploration and exploitation.

Our findings suggested that varying α can help strike a balance between the number of samples, the energy incurred, and the prediction accuracy while maintaining a high level of confidence. Specifically, in the earlier stages of the task, having a low value of α can be beneficial to obtain an accurate map (exploration) quickly. Then, gradually increasing the α value will help achieve an energy-efficient exploitation objective with a lower uncertainty in the predicted information. This would produce an optimal balanced exploration and exploitation method by which robots can model a physical process efficiently and accurately by intelligently adjusting the weight for exploration and exploitation coefficients based on the variance level, the size of the environment, etc., while fulfilling their energy requirements. Therefore, the TimeVariant approach varies the α based on the current time instance, which could be helpful as an exemplar strategy that dynamically prioritizes the objectives.

A. Exploration Perspective

The exploration objective is to obtain accurate predictions of the sampled environmental process with the highest confidence bounds and lowest variance in all areas of the map. We compare the performance metrics with respect to the variation of the α values in Eq. 6 at different instances during the exploration task. As can be seen from Figs. 2, 3 and 4, the higher the alpha values, the lesser the cumulative distance will be. Furthermore, the number of samples required for RMSE and variance saturation also increases with alpha. MaxVar, i.e., $\alpha = 0$ yielded good convergence results for all cases but at the cost of increased distance and number of samples that the robot had to collect. For $\alpha \geq 0.5$, the values for variance are not stable as the robot gets stuck in the local optima and keeps taking samples from the same location; therefore, new sampling data does not improve the map variance.

Comparing the baselines with $\alpha < 0.5$, we deduce that RW provided better results for variance and RMSE. The range of alpha values greater than or equal to 0.2 and less than 0.5 achieves the best convergence and balances all performance metrics. For a more accurate selection of alpha values, the energy consumption and variance must be taken into account depending on the specific mapping scenario. The proposed TimeVariant approach outperformed other methods considering both RMSE and distance together but did reasonably well in terms of reducing the overall uncertainty (variance) of the map.

B. Exploitation Perspective

Table III shows the localization accuracy of all variants based on the number of times the resultant GP map can be used to correctly identify the source locations after collecting different numbers of samples during the experiments.

Fig. 2 shows the results of the experiments in terms of RMSE, Variance, and cumulative distance, with respect to different alpha values. Surprisingly, the MaxMean did not perform well, especially in the early stages of the experiment. With FS, fewer measurements are taken, while with MaxMean, the measurements are taken repetitively at the same location (local maxima) since the information function (with $\alpha = 1$) is only dependent on the predicted GP mean. The pure RW approach had improved performance, but it was slower than the other approaches. The proposed TimeVariant approach outperformed others in terms of RMSE and source localization accuracy. Alpha0.25 and MaxVar, however, do not represent cost-effective exploitation since both involve very long distances. As a result, Alpha0.5, TimeVariant, and Alpha0.75 converged faster and are cost-effective.

From Figs 3 and 4, we can find that the alpha range between 0.5 and 0.75 works when minimizing distance cost is the first priority (e.g., if energy availability is heavily limited [27]). However, the increased variance continues to be a concern. In general, Alpha0.25 works best in situations where the source localization has to be even more accurate, and cost-effectiveness is less of a concern. When cost is not a concern at all, then surprisingly, the MaxVar approach ($\alpha = 0$) provided the best exploitation performance in satisfying both mapping accuracy and confidence. Overall, the TimeVariant provided the best and most balanced results.

C. Further Discussion

Additional results and figures such as the comparison of the obtained 2D Wi-Fi map's mean and variance plots for all

TABLE III

SOURCE LOCALIZATION ACCURACY (%) AT DIFFERENT TIME INSTANTS IN SINGLE SOURCE EXPERIMENTS.

Samples	Alpha0.75	Alpha0.5	Alpha0.25	MaxVar	MaxMean	RW	FS	TimeVariant-$\alpha(t)$
10	56	48	44	60	52	40	24	28
25	96	96	96	100	56	64	28	96
35	100	96	96	100	56	80	-	96
45	100	96	100	100	56	80	-	96
50	100	100	100	100	56	80	-	96
After half samples	100	100	100	100	56	80	48	100
After last sample	100	100	100	100	56	100	68	100

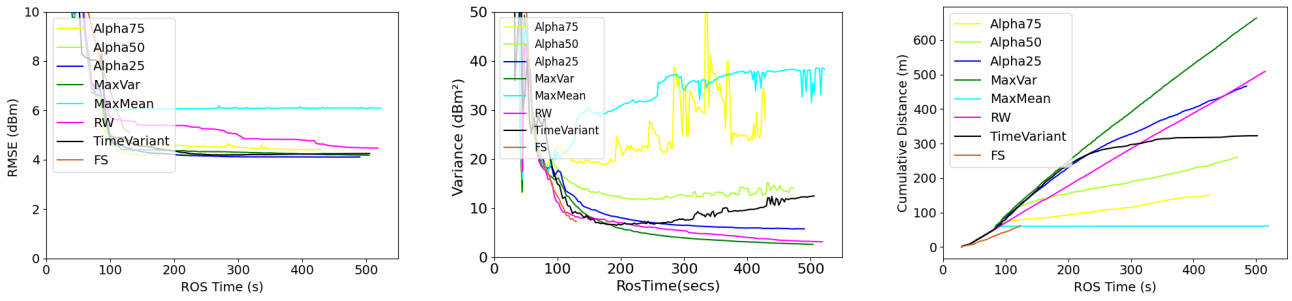

(a) RMSE (b) Variance (c) Cumulative Distance

Fig. 3. RMSE, Variance, and Cumulative Distance metrics obtained in our experiments for all variants for single source experiments.

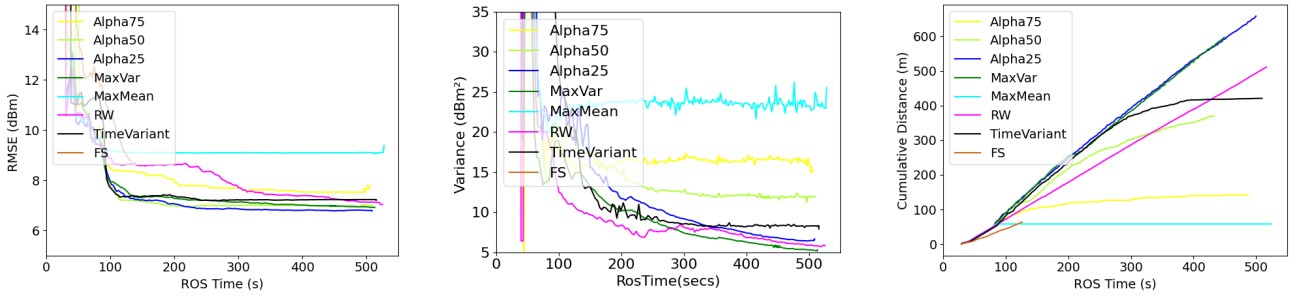

(a) RMSE (b) Variance (c) Cumulative Distance

Fig. 4. RMSE, Variance, and Cumulative Distance metrics obtained in our experiments for multi-modal experiments.

the approaches as well as high-resolution figures are available as an appendix to this paper at the resource webpage[1].

Impact of source locations: We also analyzed the impact of different source locations of single source experiments on the sampling performance. We found that there was almost no impact on the results across all sources, especially when the β value (i.e., the weight towards confidence bounds) is higher. However, for variants where α the value is higher (MaxMean, Alpha0.75, and Alpha0.5), they gave significantly different results for the furthest source locations bottom-right (0,14) and top-right (9,14) of the map area. This could be attributed to the fact that when $alpha$ is higher, exploitation is more preferred, and therefore localizing a source that is much farther could be difficult to accomplish. In summary, the effect of source locations was not observed for informative functions with greater weights for exploration.

Impact of signal distribution: Further, we analyzed the impact of the signal distribution itself on the results by repeating all the variants on single source experiments with a different path loss factor ($\eta = 2$). It was observed that all approaches with $\eta = 2$ performed similarly to the cases of $\eta = 3$, and we found that the change in variance and RMSE metrics was smoother for all variants when $\eta = 2$ than the same approaches for $\eta = 3$, as expected.

VII. CONCLUSION

This study provided a novel perspective on balancing the exploration and exploitation objectives for different adaptive sampling functions of informative path planning in mobile robots. We conduct an in-depth analysis of various variants of informative functions in an adaptive sampling of the environmental phenomenon. Our results show that the optimum approach varies depending on the objective of sampling. Both the energy efficiency and mapping/localization objectives can be met with an optimum balance based on the specific objective and application domain. The unique data generated through extensive simulations and time-consuming experiments provide insights into choosing the best range of weight values for the exploration and exploitation coefficients. This holistic perspective will be useful in achieving realistic robot tasks such as exploration, search, and rescue situations, for example, mapping the radio signal distribu-

tion and identifying victims' locations through their radio beacons.

REFERENCES

[1] A. Viseras, D. Shutin, and L. Merino, "Online information gathering using sampling-based planners and gps: An information theoretic approach," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 123–130.

[2] M. Malencia, S. Manjanna, M. A. Hsieh, G. Pappas, and V. Kumar, "Adaptive sampling of latent phenomena using heterogeneous robot teams (aslap-hr)," *arXiv preprint arXiv:2208.06053*, 2022.

[3] A. Singh, A. Krause, and W. J. Kaiser, "Nonmyopic adaptive informative path planning for multiple robots," *IJCAI International Joint Conference on Artificial Intelligence*, pp. 1843–1850, 2009.

[4] G. Hitz, E. Galceran, M. Ève Garneau, F. Pomerleau, and R. Siegwart, "Adaptive continuous-space informative path planning for online environmental monitoring," *Journal of Field Robotics*, vol. 34, pp. 1427–1449, 2017.

[5] W. Luo and K. Sycara, "Adaptive sampling and online learning in multi-robot sensor coverage with mixture of gaussian processes," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 6359–6364.

[6] J. Fink and V. Kumar, "Online methods for radio signal mapping with mobile robots," in *2010 IEEE International Conference on Robotics and Automation*, 2010, pp. 1940–1945.

[7] K. C. Ma, L. Liu, and G. S. Sukhatme, "Informative planning and online learning with sparse gaussian processes," *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 4292–4298, 2017.

[8] W. Luo, C. Nam, G. Kantor, and K. Sycara, "Distributed environmental modeling and adaptive sampling for multi-robot sensor coverage," in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 2019, pp. 1488–1496.

[9] S. Caccamo, R. Parasuraman, L. Freda, M. Gianni, and P. Ögren, "Rcamp: A resilient communication-aware motion planner for mobile robots with autonomous repair of wireless connectivity," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 2010–2017.

[10] N. Fung, J. Rogers, C. Nieto, H. I. Christensen, S. Kemna, and G. Sukhatme, "Coordinating multi-robot systems through environment partitioning for adaptive informative sampling," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 3231–3237.

[11] C. Stachniss, C. Plagemann, and A. J. Lilienthal, "Learning gas distribution models using sparse gaussian process mixtures," *Autonomous Robots*, vol. 26, pp. 187–202, 2009.

[12] A. H. Zakaria, Y. M. Mustafah, J. Abdullah, N. Khair, and T. Abdullah, "Development of autonomous radiation mapping robot," in *Procedia Computer Science*, vol. 105. Elsevier B.V., 2017, pp. 81–86.

[13] P. Gabrlik, T. Lazna, T. Jilek, P. Sladek, and L. Zalud, "An automated heterogeneous robotic system for radiation surveys: Design and field testing," *Journal of Field Robotics*, vol. 38, no. 5, pp. 657–683, 2021.

[14] P. Ghassemi and S. Chowdhury, "Decentralized informative path planning with exploration-exploitation balance for swarm robotic search," *arXiv preprint arXiv:1905.09988*, 2019.

[15] N. A. A. Rahman, K. S. M. Sahari, M. F. A. Jalal, A. A. Rahman, M. I. A. Adziz, and M. Z. Hassan, "Mobile robot for radiation mapping in indoor environment," in *IOP Conference Series: Materials Science and Engineering*, vol. 785. Institute of Physics Publishing, 5 2020.

[16] R. A. Cortez, X. Papageorgiou, H. G. Tanner, A. V. Klimenko, K. N. Borozdin, R. Lumia, and W. C. Priedhorsky, "Smart radiation sensor management," *IEEE Robotics and Automation Magazine*, vol. 15, pp. 85–93, 2008.

[17] C. Williams and C. Rasmussen, "Gaussian processes for regression," *Advances in neural information processing systems*, vol. 8, 1995.

[18] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger, "Information-theoretic regret bounds for gaussian process optimization in the bandit setting," *IEEE Transactions on Information Theory*, vol. 58, no. 5, pp. 3250–3265, 2012.

[19] R. Miyagusuku, A. Yamashita, and H. Asama, "Improving gaussian processes based mapping of wireless signals using path loss models," in *IEEE International Conference on Intelligent Robots and Systems*, vol. 2016-November. Institute of Electrical and Electronics Engineers Inc., 11 2016, pp. 4610–4615.

[20] A. Al Redwan Newaz, S. Jeong, and N. Y. Chong, "Fast radioactive hotspot localization using a uav," in *2016 IEEE International Conference on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAR)*, 2016, pp. 9–15.

[21] D. Gu and H. Hu, "Spatial gaussian process regression with mobile sensor networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 23, pp. 1279–1290, 2012.

[22] S. Kemna, J. G. Rogers, C. Nieto-Granda, S. Young, and G. S. Sukhatme, "Multi-robot coordination through dynamic voronoi partitioning for informative adaptive sampling in communication-constrained environments," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 2124–2130.

[23] N. Ulapane, K. Thiyagarajan, and S. Kodagoda, "Hyper-parameter initialization for squared exponential kernel-based gaussian process regression," in *2020 15th IEEE Conference on Industrial Electronics and Applications (ICIEA)*. IEEE, 2020, pp. 1154–1159.

[24] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2, no. 3.

[25] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, A. Y. Ng *et al.*, "Ros: an open-source robot operating system," in *ICRA workshop on open source software*, vol. 3, no. 3.2. Kobe, Japan, 2009, p. 5.

[26] Y. Shi, N. Wang, J. Zheng, Y. Zhang, S. Yi, W. Luo, and K. Sycara, "Adaptive informative sampling with environment partitioning for heterogeneous multi-robot systems," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 11 718–11 723.

[27] R. Parasuraman, P. Pagala, K. Kershaw, and M. Ferre, "Energy management module for mobile robots in hostile environments," in *Conference Towards Autonomous Robotic Systems*. Springer, 2012, pp. 430–431.